AVIAD HEIFETZ

Game Theory

Interactive Strategies in Economics and Management

Game Theory

Game theory is concerned with strategic interaction among several decision-makers. In such strategic encounters, all players are aware of the fact that their actions affect the other players. Game theory analyzes how these strategic, interactive considerations may affect the players' decisions and influence the final outcome. This textbook focuses on applications of complete-information games in economics and management, as well as in other fields such as political science, law, and biology. It guides students through the fundamentals of game theory by letting examples lead the way to the concepts needed to solve them. It provides opportunities for self-study and self-testing through an extensive pedagogical apparatus of examples, questions, and answers. The book also includes more advanced material suitable as a basis for seminar papers or elective topics, including rationalizability, stability of equilibria (with discrete-time dynamics), games and evolution, equilibrium selection, and global games.

Aviad Heifetz is Professor of Economics at the Open University of Israel, where he served as chair of the Economics and Management Department from 2006–2009. He was previously Visiting Professor of Managerial Economics at Kellogg School of Management, Northwestern University (2009–2011).

Game Theory

Interactive Strategies in Economics and Management

Aviad Heifetz

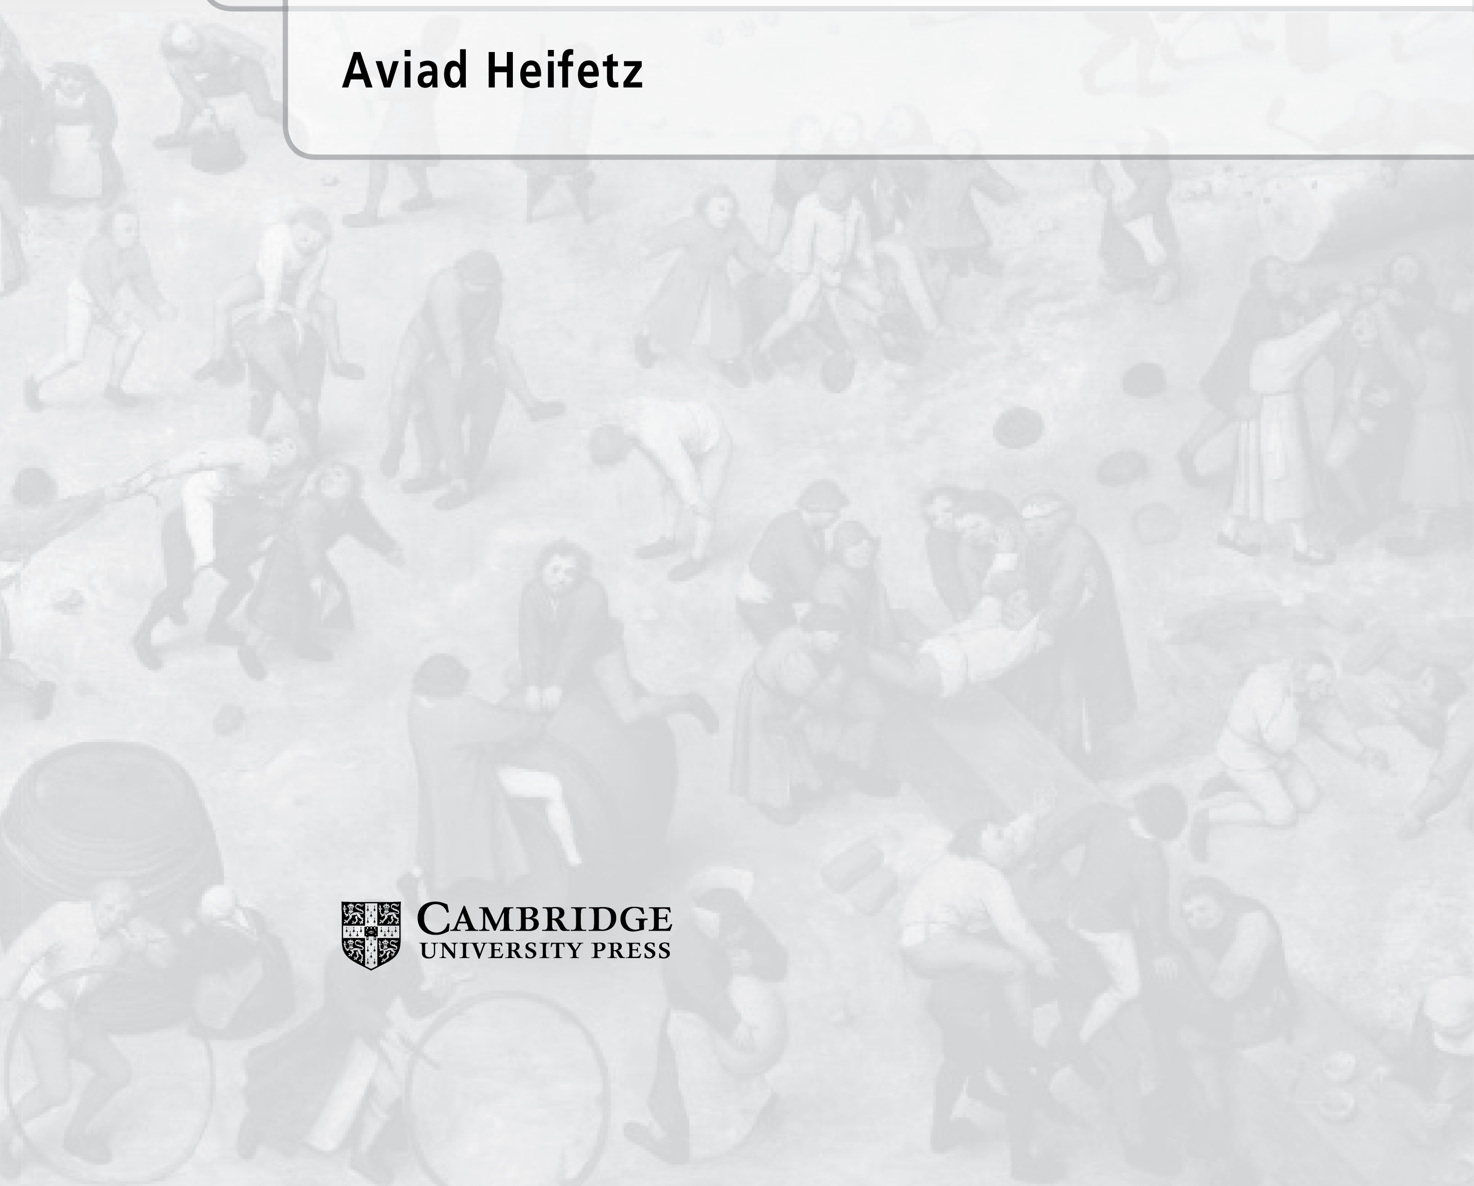

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town,
Singapore, São Paulo, New Delhi, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521176040

© Open University of Israel, 2012

This publication is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without the written
permission of Cambridge University Press.

First published in Hebrew by Open University of Israel Press 2008
First published in English by Cambridge University Press 2012

Translation: Judith Yalon Fortus

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data
Heifetz, Aviad, 1965–
Game theory : interactive strategies in economics and management / Aviad Heifetz.
pages cm
ISBN 978-0-521-76449-0
1. Game theory. 2. Economics. 3. Management. I. Title.
HB144.H455 2012
330.01′5193–dc23 2012000086

ISBN 978-0-521-76449-0 Hardback
ISBN 978-0-521-17604-0 Paperback

Cambridge University Press has no responsibility for the persistence or
accuracy of URLs for external or third-party Internet websites referred to
in this publication, and does not guarantee that any content on such
websites is, or will remain, accurate or appropriate.

In loving memory of my father
Gutman Heifetz, 1933–2010

CONTENTS

9 Coordination games and strategic uncertainty

In 1753, the Academy of Dijon in France announced an essay contest on the topic of: "*What is the origin of inequality among men, and is it authorized by the natural law?*"

Jean-Jacques Rousseau took up the challenge, and wrote his discourse on "The Origin and the Foundation of Inequality among Men".[1] This essay (together with his later book, "*The Social Contract, or Principles of Political Right*,"[2] published in 1762) became one of the cornerstones of the social sciences and political philosophy.

Rousseau secluded himself for a week in the Forest of St Germain to muse about what human life was like at the dawn of civilization, and completed the discourse upon returning to Paris. In his discourse, Rousseau describes human evolution from the primordial era of the "savage" to the age of the social order of "civilized man."

In the first part of his discourse, Rousseau describes the "savage" or natural man who lived a life of instinct in Nature, innocent of alienation and inequality. In the second part, he goes on to describe the gradual process whereby social and political organization-forming activity took shape, providing man with security and technological progress, but also bringing inequality, wars and alienation in its wake. Rousseau describes the beginning of social cohesion as follows:

> Taught by experience that the love of well-being is the sole motive of human actions, he found himself in a position to distinguish the few cases, in which mutual interest might justify him in relying upon the assistance of his fellows; and also the still fewer cases in which a conflict of interests might give cause to suspect them. In the former case, he joined in the same herd with them, or at most in some kind of loose association, that laid no restraint on its members, and lasted no longer than the transitory occasion that formed it. In the latter case, every one sought his own private advantage, either by open force, if he thought himself strong enough, or by address and cunning, if he felt himself the weaker.

[1] *The Social Contract and Discourses by Jean-Jacques Rousseau*, translated with an Introduction by G. D. H. Cole (London and Toronto: J. M. Dent and Sons, 1923).

[2] Ibid.

> In this manner, men may have insensibly acquired some gross ideas of mutual undertakings, and of the advantages of fulfilling them: that is, just so far as their present and apparent interest was concerned: for they were perfect strangers to foresight, and were so far from troubling themselves about the distant future, that they hardly thought of the morrow. If a deer was to be taken, every one saw that, in order to succeed, he must abide faithfully by his post: but if a hare happened to come within the reach of any one of them, it is not to be doubted that he pursued it without scruple, and, having seized his prey, cared very little, if by so doing he caused his companions to miss theirs.[3]

The archetypal scenario Rousseau describes here is particularly appropriate for describing using the tools of game theory. First, man perceives himself as separate from other humans; and social conventions dictating the categories that will help him understand himself and the world around him have yet to come into existence. Wherefore, “the love of well-being is the sole motive of human actions.” The extent to which the individual enjoys well-being depends, of course, on the actions he and other people take, but there is no interdependence in the definition of different people’s well-being.

Indeed, it is precisely this notion of preferences that we ascribe to players in game theory. We assume that every player can be characterized by the way she ranks the possible action profiles of all the players (including her own). A preferred action profile will procure her a greater measure of well-being, which may accordingly be represented by a higher level of utility according to the utility function that represents her preferences. The utility function is one of the building blocks with the aid of which the game is defined. Thus, in a game, we may replace the utility function of one of the players by another function (which will define the payoff to that player for any action profile) without modifying the definition of the utility function of the other players, and obtain a well-defined new game.

Second, Rousseau describes interactions, each of which is unique, and is perceived and analyzed on its own account. People “were so far from troubling themselves about the distant future, that they hardly thought of the morrow.” In other words, Rousseau here describes interactions that are amenable to description with the aid of strategic form games.

To begin with, Rousseau notes the few opportunities that a person encounters “in which mutual interest might justify him in relying upon the assistance of his fellows.” In the language of game theory, these are situations in which man perceives that each of those surrounding him has a **dominant strategy**.

Rousseau adds, however, that not all opportunities are of this kind, and promptly cites a pertinent example. The specific game that Rousseau describes has been designated the Stag Hunt game. It may be described as follows.

[3] Ibid.

9.1 The Stag Hunt game

A band of hunters is trying to catch a stag. If the hunting band consists of two hunters, the payoffs in the game can be described as follows:

		Hunter 2	
		Stag	Hare
Hunter 1	Stag	3,3	0,2
	Hare	2,0	2,2

Why does this game reflect Rousseau's description of the situation? The utility from the hunting of the hare is here represented by the payoff 2. It is easier for each hunter to catch a hare on his own. Therefore, if a hunter chooses to set out to catch a hare, he will ensure himself of a payoff of 2, regardless of what the other hunter does. Success in a stag hunt, however, necessitates coordination between the two hunters – they must ambush the stag from two different points in the forest. Given such coordination, the stag hunt will be crowned with success, the hunters will share the kill between them, and each hunter will obtain a larger hunk of meat (represented by the payoff 3) than he would have obtained by hunting a hare on his own. But if one of the hunters should desert his post during the stag hunt in order to catch a hare he has happened to catch sight of, the remaining hunter will be unable to capture the stag on his own, and will be left empty-handed (with a payoff of 0).

There are two Nash equilibria in this game. In one of the equilibria each hunter traps a hare on his own and gets a payoff of 2. In the other, the two hunters collaborate in the hunting of the stag, and get a payoff of 3. For both hunters, the second equilibrium is preferable to the first equilibrium – the stag hunt yields a higher payoff for each one of them.

Definition

When one of two equilibria

1. is deemed at least as preferable as the other equilibrium **by all players**, and
2. is strictly preferred over the other equilibrium by at least one of the players,

we say that it is **more efficient** than the other equilibrium. An equilibrium is called **payoff dominant** if it is more efficient than all other equilibria of the game.

Definition A game with several Nash equilibria, any two of which are comparable in terms of their efficiency, is called **a coordination game**.

It is not possible to rate the equilibria in terms of efficiency in every game with several equilibria. In other words, not every game is a coordination game. In games such as the Battle of the Sexes (in section 6.1) and Divvying up the Jackpot (in section 6.3.1), for example, if equilibrium A is preferable to equilibrium B in the eyes of a particular player, then in the eyes of the other player, by contrast, equilibrium B is preferable to equilibrium A. Accordingly, in these games, no one equilibrium is more efficient than the other, since the players cannot reach unanimous agreement as to which equilibrium is preferable.

Therefore, in a game such as the Stag Hunt, in which a particular equilibrium can be indicated as being the most efficient of all, one might reasonably assume that the players would naturally focus on the efficient equilibrium as a focal point equilibrium, and that this equilibrium would be the one to be brought into play.

As we have seen, Rousseau does not share this view. In his opinion, "if a hare happened to come within the reach of any one of them, it is not to be doubted that he pursued it without scruple." What negative feature can be found, therefore, in the efficient equilibrium?

When each hunter sets out to bag a hare on his own, he is entirely independent of anyone else's cooperation. Had the hunters reached prior agreement to adhere to the non-efficient equilibrium and set out on a hare hunt, no hunter would have suffered any damage if another hunter had, nevertheless, tried (unsuccessfully!) to catch a stag. In other words, this non-efficient equilibrium is safe for every player, since no player suffers damage if the other player nonetheless deviates from the agreement.

By contrast, the efficient equilibrium, in which the players cooperate in hunting the stag, is risky for every one of them. Of course, being a Nash equilibrium, such cooperation is a stable agreement: if every player believes that the other will play his part in the agreement, he will prefer to adhere to it likewise. However, if the player suspects that for any reason the other player will not do his bit – either because he has decided to hunt a hare as described by Rousseau, or because he has been taken ill in mid-hunt and is unable to continue, or for any other reason – the player will begin to doubt whether it is worth his while to continue in pursuit of the stag. After all, if he withdraws from the agreement and catches himself a hare, he can save himself the uncertainty and the risk of going dinner-less. Thus, every hunter cooperating in the stag hunt faces **strategic uncertainty**.

As we saw in Rousseau's description, the criterion of strategic uncertainty prevails over the criterion of efficiency in the eyes of the hunters. What, in your opinion, is the leading criterion in situations of this sort? What does it depend on?

Actually, Rousseau describes a more complex scenario. The hunter who betrayed his comrades' trust "having seized his prey, cared very little, if by so doing he caused his companions to miss theirs." Which is to say, there are more than two hunters in the "herd," but the withdrawal of even a single hunter suffices for the stag hunt to fail. In other words, the "weakest link" in the band is liable to bring about the failure of the group as a whole.

Let us assume, therefore, that the band consists of n hunters and that, for the stag hunt to succeed, all must cooperate. As before, success in the stag hunt secures each hunter a payoff of 3, while failure incurs a payoff of 0 for every hunter who has participated in the stag hunt.[4] Alternatively, every hunter who succeeds in bagging himself a hare will thereby ensure himself a payoff of 2 (while dooming to failure all the other hunters who are trying to hunt a stag).

What happens when the band of hunters sets out to hunt a stag? Let us assume that a certain hunter believes that every one of the other $n-1$ hunters is liable, with a slight chance of $\varepsilon > 0$, to withdraw from the stag hunt of his own volition, and that the chances of withdrawal of the various hunters are independent of one another. In the view of that hunter, the probability that all the other hunters will adhere to the common effort is only $(1-\varepsilon)^{n-1}$. The greater the number of hunters, n, the chance $(1-\varepsilon)^{n-1}$ tends to 0. If, for instance, the prospect of abandonment by each hunter is $\varepsilon = \frac{1}{10}$, then the likelihood that not one of the ten hunters will opt out is $\left(\frac{9}{10}\right)^{10} \cong 0.347$, and the prospect that 100 hunters will all, to a man, stick to the job in hand is only $\left(\frac{9}{10}\right)^{100} \cong 0.0000266$.

Therefore, the more hunters there are in the band, the more each hunter ought to fear that the mission will not succeed, and the greater the temptation he faces to withdraw from the stag hunt in order to assure himself of a hare for dinner. In this sort of scenario, the intuition presented by Rousseau comes more sharply into focus.

More generally, less extreme situations could occur in which it suffices that a part α of the hunting band will stick to the job in hand in order for the stag hunt to be crowned with success. In such a case, when at least a proportion α out of the band of hunters set out to hunt a stag, they pull it off successfully, winning a payoff of 3. By contrast, if fewer than α of the hunters band together to hunt the stag, these hunters fail and return home empty-handed (with a payoff of 0). As before, each hunter who elects to hunt a hare assures himself of a payoff of 2.

[4] This is a reasonable situation if a large band of hunters can hunt a large herd of deer, such that the overall kill is proportionate to the number of hunters, and the kill per hunter remains constant. In order to simplify the terminology, we will nevertheless continue to speak of "a stag" rather than a "herd of deer."

9.1.1 Whale hunting

A hunting game of this type actually takes place in a whale-hunting village on the island of Lembata in Indonesia.[5,6] In the dry season, between May and September, boats powered by oars and palm-fronds sails put out to sea every morning, to a distance of up to about 13 kilometers from the shore, to find and hunt whales. Whale hunting is a complex and dangerous task that calls for at least eight crew members – the "captain," the harpooner and his assistant, the helmsman, and others. When the boatmen spot a whale, they usually lower the sail and row powerfully in the direction of the prey. The moment the boat comes within suitable range, the harpooner, standing in the boat on a small platform especially designed for the purpose, casts his spear at the whale. The whale then plunges or drags the boat with him until it is exhausted. The danger, of course, is that the whale will drag the boat far out to sea, or that the boat will capsize.

Thus, early every morning, the "captain" of each wooden whaling boat (called a *téna*) must recruit a crew of at least eight men for the job. The villagers must decide whether to join the boat, or, alternatively, to go fishing for themselves, either alone or in pairs, seeking smaller fry near the shore (or doing other onshore jobs, such as tending to their livestock, which will include a few goats, some poultry, and some pigs). Each crew member taking part in the whale hunt will obtain, on average, a larger hunk of meat than he could get by fishing close to the shore. In addition, the crew member earns the gratitude of his extended family, since he shares the catch with them.

However, over a period of years in which it has gradually become increasingly apparent that schools of whales in the fishing zone are small and rare, the "captains" face, day by day, a tougher job in manning their boats. On the whole, crew members whose whaling boats, on a particular day, do not go to sea, don't go fishing near the shore either because to do that they need boats of a different type (coracles), nets instead of harpoons, and so forth. Thus the "strategic uncertainty" facing the whale hunters increases, and the number of boats putting out to sea every morning progressively decreases throughout the hunting season in such difficult years.

9.1.2 Laboratory experiments of the Stag Hunt game

Various laboratory experiments have been devised to examine the Stag Hunt game. In one such experiment[7] the participants repeatedly played the game with the following payoffs:

[5] Alvard, M. S. and D. A. Nolin (2002), "Rousseau's Whale Hunt? Coordination among Big-Game Hunters," *Current Anthropology* 43 (4), 533–559.

[6] Indonesia is not a signatory to the international convention for the prevention of whale hunting, but in any event, that convention exempts from restrictions natives who engage in whale hunting for their own subsistence.

[7] Cooper, R., D. DeJong, B. Forsythe, and T. Ross (1990), "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review* 80, 218–233.

		Player 2: S	Player 2: H
Player 1	S	100,100	0,80
	H	80,0	80,80

where the figures represent the percentage prospect of obtaining $1 from the experimenters. In this game, each player can assure himself of an 80 percent prospect of winning the dollar if he chooses strategy A. In order to close the small remaining gap to 100 percent, the two players must successfully coordinate between themselves the choice of the strategy S. But the failure of such coordination will leave the player who chooses S with no prospect of winning the dollar, and therefore the risk involved in this strategy may be perceived as high in comparison with the difference in rewards. Sure enough, in the final eleven rounds of the experiment, the players, in 97 percent of cases, chose the profile (H, H).

The picture changed when the participants were permitted to send each other messages. When player 1 was permitted, prior to each round of the game, to announce the strategy he was about to play, the players were able to coordinate the most worthwhile profile (S, S) in 53 percent of cases; but a lack of coordination occurred in 31 percent of cases, giving (S, H) or (H, S), while in the remaining 16 percent of cases the participants played (H, H). Thus, in some cases in which player 1 announced his intention of playing S, he nevertheless feared lack of cooperation on the part of player 2, and therefore, ultimately, played the safe strategy H after all. Correspondingly (or consequently), player 2 did not in fact cooperate in choosing S, even when player 1 announced his intention of choosing S. This resulted in instances of non-coordination, or in the choice of the safe strategy H by both players.

The picture improved dramatically when the experimenters permitted two-way communication, meaning that they permitted player 2 to respond and likewise to declare the strategy he intended to choose. In the last eleven rounds of the game, the participants played the efficient profile (S, S) in 91 percent of cases, while the remaining cases suffered from lack of coordination.

The Stag Hunt game is the archetype of a large group of coordination games. A coordination game, it will be recalled, is a game in which there are several Nash equilibria that can be ranked according to their efficiency.

In another coordination game experiment, there was a hierarchy of seven possible levels of cooperation between the players, where the payoff to each player depended on his choice and on the minimal level of cooperation within the group of players as a whole.[8] The choice of "7" by all players ensures the highest payoff to all, but also entails the heaviest "fines" if one player deviates and chooses a lower level of cooperation. Universal choice of a lower level of cooperation yields a lower payoff to each player, but also reduces the "fines" imposed in case of deviation on the part of any of the players.

Expressly, in this game the payoffs to each player were as follows:

		The lowest number selected by the players						
		7	6	5	4	3	2	1
	7	1.3	1.1	0.9	0.7	0.5	0.3	0.1
	6	–	1.2	1.0	0.8	0.6	0.4	0.2
	5	–	–	1.1	0.9	0.7	0.5	0.3
The player's choice	4	–	–	–	1.0	0.8	0.6	0.4
	3	–	–	–	–	0.9	0.7	0.5
	2	–	–	–	–	–	0.8	0.6
	1	–	–	–	–	–	–	0.7

In this game, there are seven equilibria – every profile of actions in which all players select the same level X of cooperation is a Nash equilibrium. How did the participants actually play?

Each participant took part in 7–10 rounds of the game. When a large number of players participated, a gradual convergence took place in the course of the game rounds, toward a players' choice of low (and "safe") levels of cooperation – to the lowest level of "1" in 77 percent of instances in the final round of the game, and to a level of "2" in another 17 percent of final-round instances.

By contrast, when there were only two players in a set, the great majority of playing pairs (twenty-one out of twenty-four pairs in the experiment) converged to coordinating on the highest level "7" of cooperation in the final round of the game. It transpires that when there are only two participants in a game and one of them is initially fearful and starts off his game rounds by selecting a low level, his fellow player is frequently prepared to "wait" for him by repeatedly selecting the level "7" until, in most instances, the first player in fact quickly overcomes his misgivings and joins the efficient "7" choice. Such a "waiting period" was not observed when the

[8] Van Huyck, J. B., R. C. Battalio, and R. Beil (1990), "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review* 80, 234–248.

pair was randomly swapped in each game round. In this configuration, insistence on a "7" level of cooperation cannot serve as a signal of preparedness for effective cooperation, and the cooperation level mostly deteriorated to the lowest common denominator, namely "1."

In a similar experiment, the participants played in trios.[9] The choice of S by everybody would have secured a payoff of 90 for all, but a deviation by any one of them to H would have caused the payoff to those adhering to S to plummet to 10. Meanwhile, the universal choice of H would have only slightly reduced the reward to 80, while at the same time greatly reducing the risk inhering in deviation on the part of one of the players; a player choosing H where one of the others chooses S would get 60.

In this experiment, seven out of eight trios of participants were successful in learning to mutually coordinate on the efficient action S in the course of twenty game rounds. In practice, the participants chose S three-fourths of the time, even in the early rounds of the experiment.

The results were completely different when eight participants were (virtually) arranged in a circle, each playing the game with whoever was next to him in the circle (which is to say that as before, his payoffs depended on his choice and the choices made by his neighbors on either side; the same choice of a player pertained to his interaction with both his neighbors). Here the participants chose S from the outset only half the time, and the frequency of the selection S deteriorated progressively during the twenty game rounds. In the last round, none of the participants chose S; rather, the unanimous choice fell on H.

What is the origin of the difference in outcomes? Where each trio is isolated unto itself, mutual trust can evolve. In the circle game, by contrast, each participant is indirectly dependent also on his remote neighbors: they affect their neighbors' choices, which in turn affect those of their own neighbors, and so on and so forth; and ultimately, they also affect the participant's immediate neighbors. Thus, a reluctance to choose S is liable to spread like a plague through the whole circle, ultimately "infecting" everybody, and causing them to choose H.

Stag Hunt games, or more general games of coordination, succinctly describe many realistic situations unrelated to the world of hunting. We will now describe a few examples.

9.2 Keyboard arrangement

The generally accepted arrangement of English letter keys on the computer keyboard is called QWERTY, after the first six characters running from left to

[9] Berninghaus, S. K., K. M. Erhart, and C. Keser (2002), "Conventions and Local Interaction Structures," *Games and Economic behavior* 39, 177–205.

right on the top row of letters. This arrangement has nothing to do with typing convenience but originates from a period when mechanical typewriters were used for typing.

In mechanical typewriters, striking a keyboard key activated a lever, the end of which featured a relief of the corresponding letter. When activated, the lever was thrown forward to impact an inked ribbon that was stretched close across the sheet of paper held in the rollers, and the letter mould was thus imprinted on the page. A common problem in mechanical typewriters was that two levers could become entangled when two adjacent letters on the keyboard were struck one after another. The keyboard was therefore arranged in such a way that letters frequently succeeding one another in English words would be kept apart.

But preventing the typewriter levers from getting crossed was not the only criterion for the efficiency of the alpha-numerical arrangement of the keyboard. Typing speed was another important issue, and the lever problem is in any event no longer relevant in the age of personal computers. In the 1930s, August Dvorak and William Dealy therefore invented a different arrangement of keyboard characters, which could significantly improve typing speeds. The arrangement is designated DVORAK.[10]

The preferable keyboard layout, however, was not popular and the QWERTY arrangement remains dominant in practice. Since most of the world's keyboards conform to that layout, there is little point in learning and adapting to a different design. If a majority of computer users worldwide were to switch simultaneously to using the DVORAK-style keyboard, typing would become quicker and easier, and the improvement would justify the costs and the effort involved. But the existing situation, in which almost everybody is accustomed to QWERTY typing, itself likewise constitutes a Nash equilibrium among computer users.

Comprehension check

Describe a game between computer users that corresponds to the above description, and indicate its payoffs.

[10] In Microsoft Windows operating systems, for example, you can avail yourself of the option of using DVORAK by choosing Start – > Settings – > Control Panel – > Regional and Language Options – > Text Services and Input Languages and add the DVORAK layout.

9.3 Video cassette recording technology

In the 1980s, two rival technologies, VHS and Betamax, competed in the video cassette recording (VCR) and viewing market. In terms of recording quality, Betamax was considered preferable. The Japanese firm Sony produced VCRs using this technology and drove it to market. Almost all other electronics manufacturers, however, produced VHS-compatible equipment. In order to encourage consumers to purchase its VCRs, Sony promised that it would continue producing them by means of that technology, and that it would set up Betamax movie lending libraries worldwide.

For several years, the two technologies coexisted in the market. Gradually, the VHS technology bit off a larger market share and, commencing from a certain stage, the vast majority of demand was channeled to VHS-type cassette recorders. Sony thereupon also started producing VHS, relinquishing its Betamax line.

9.4 Consumer network externalities[11]

The utility consumers derive from products often depends on the number of other consumers using the same product. The facsimile machine is a classic example of this rule. If nobody, anywhere, had a fax machine, it would not be worth anybody's while to get one, because there would be no one to send faxes to or receive from. For every potential user, the utility to be got from the device progressively increases with the number of people worldwide who likewise use a fax machine. Even so, given a certain distribution of fax machines among the population, various users will gain a different level of utility from purchasing one for themselves. For example, various business owners obtaining a fax machine for their business will be able to boost their income to different extents. Accordingly, the maximum price each will be willing to pay for the machine will vary.

The simplest way to model such a state of affairs is to assume that every potential fax consumer is characterized by her **type**, $\tau > 0$. If the number of fax users in the population is $n \geq 2$, then a type τ consumer is prepared to pay at most $n\tau$ to purchase a fax machine. But if nobody but the potential consumer herself purchases a fax machine, then a type τ consumer will not shell out any cash for the device, because she will have no use for it.

[11] The model described here is based on Rohlfs, J. (1974), "A Theory of Interdependent Demand for a Communications Service," *The Bell Journal of Economics and Management Science* 5, 16–37.

To simplify the problem even further, let's assume that in a particular country there are A potential fax consumers, of the types $\tau = 1,2, \ldots, A$. We will assume that fax machines are now being initially offered for sale in the market at price p. Who of the consumers will purchase fax machines at a Nash equilibrium?

First, a situation could come about in which everyone believes that nobody else is going to purchase a fax machine. Thus they will all refrain from purchasing the device. In other words, a state of affairs in which nobody buys a fax machine is a Nash equilibrium.

Is there also an equilibrium in which fax machines are in demand? To check this out, we will assume that at a Nash equilibrium, $n^* \geq 2$ devices are sold. In other words, some types in the population find good reason to purchase a fax machine at the price p. Obviously, if an individual of type τ decides to purchase a fax machine (which is to say $n^* \tau \geq p$) then everyone of a higher type $\tau' > \tau$ will also wish to purchase a fax (since then $n^* \tau' > p$). In other words, we will be able to identify the minimal type τ^* who will wish to purchase a fax at price p given that n^* consumers altogether start using the fax. To keep it simple, let's assume that for the type τ^*, the equation:

$$p = n^* \tau^*$$

is satisfied. This means that the price p is the maximal price that type τ^* is prepared to pay for the device, given that n^* consumers altogether buy and use fax machines.

These n^* consumers are of the types:

$$\tau = \tau^*, \ldots, A$$

which is to say:

$$n^* = A - \tau^* + 1$$

Hence we may infer that:

$$p = n^* \tau^* = (A - \tau^* + 1)\tau^*$$

This quadratic equation in τ^* has two solutions (see Figure 9.1):[12]

$$\tau_1^* = \frac{(A+1) - \sqrt{(A+1)^2 - 4p}}{2}$$

$$\tau_2^* = \frac{(A+1) + \sqrt{(A+1)^2 - 4p}}{2}$$

Each of these solutions defines a Nash equilibrium. In the first equilibrium, all the types $\tau \geq \tau_1^*$ purchase fax machines, i.e. the number of faxes sold altogether is:

12 As long as $p < \frac{(A+1)^2}{4}$.

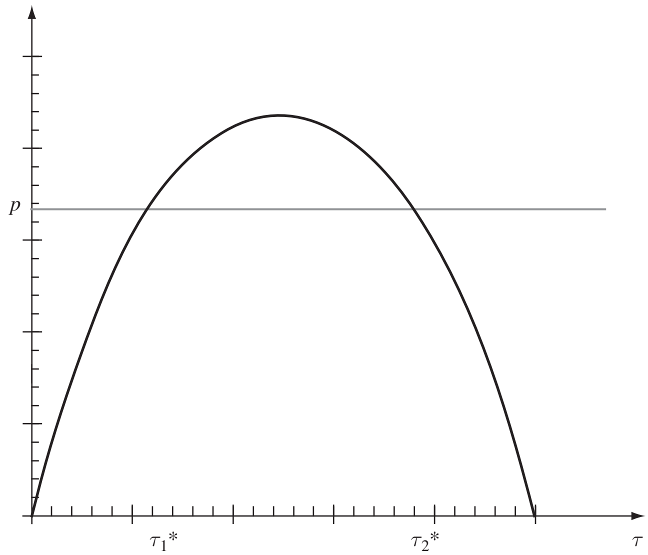

Figure 9.1

$$n_1^* = (A+1) - \tau_1^* = \frac{(A+1) + \sqrt{(A+1)^2 - 4p}}{2}$$

In the second equilibrium, only the types $\tau \geq \tau_2^*$ purchase faxes, and the overall number of devices sold is:

$$n_2^* = (A+1) - \tau_2^* = \frac{(A+1) - \sqrt{(A+1)^2 - 4p}}{2}$$

We therefore obtain that there are altogether three Nash equilibria in this game. The equilibrium τ_1^* is the most desirable of all – the one with the largest number of consumers who use the fax machines and benefit from them.

Comprehension check

The world's personal computers are currently divided into two principal types – the PCs with Microsoft Windows operating systems, and Macintosh computers. PCs are most widely used, but some consumers consider their performance to be inferior to that of the Mac, especially for graphic design applications. The greater popularity of the PC results in a situation in which a wider diversity of software programs is on offer for that type of computer.

Give a verbal and formal description of a coordination game between personal computer users, one of the equilibria of which corresponds to the state of affairs prevailing in reality in which a majority of consumers choose the PC and a minority the Macintosh.

Guidance: assume that there are two types of computer users. The first type prefers to work with Macintosh in any event. Users of the second type, constituting the majority of computer users in the population, prefer to work with that type of computer for which most software programs are written. Assume that the type of computer for which more software programs are written is the one used by the larger part of the population.

9.5 Job search and unemployment

Coordination games also have important application in describing macro-economic phenomena. We will now describe such an application, which is based on ideas from the model presented in an influential article by Diamond (1982).[13] For these and related ideas Diamond received the Nobel prize in Economics in 2010.

The inhabitants of a certain tropical island go out picking coconuts every morning. To pick a coconut, one must first climb the coconut tree. A coconut picker finding a tree that bears a coconut must decide whether to climb up and pluck it or whether to look for another tree where the coconut hangs lower, so that it will be easier to pluck. Islanders face a taboo against eating coconuts they have picked themselves. So in the afternoons, islanders who picked coconuts in the morning look for partners with whom to swap coconuts. Once the swap has taken place, each partner can eat the coconut he now has in his possession. Traditionally, coconuts, once picked, may not be kept for the next day. Therefore, an inhabitant who has picked a coconut but has not found a partner with whom to transact the swap must throw it away and does not get to eat a coconut that day.

This is, of course, "in a nutshell," a metaphor for the economic organization of modern society. The vast majority of people do not subsist solely on the products they produce themselves; they exchange most of their produce with others in consideration of different goods (by selling their produce or their labor in exchange for money, and using that money to purchase other goods and services). This is due to the sophistication and the expertise involved in manufacturing processes – most consumer goods are produced in multiple stages by very many people, using a large number of intermediate by-products of other manufacturing processes. (The tropical

[13] Diamond, P. (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90 (5), 881–894.

island of the fable has no form of professional specialization that necessitates barter trade, and therefore barter trade is anchored in a different assumption – that of the taboo.)

Most consumer goods become obsolete and spoil over the course of time. This assumption is represented in the fable by the (simplifying) assumption that coconuts won't keep from one day to the next. This is an extreme assumption, which is designed to facilitate the computations in the model we will present. (Alternatively, the value of a coconut might have been assumed to decrease gradually over time,[14] but this complication would not yield any new insights from the model for our purposes here.)

Job seeking or searching for a business opportunity is likened, in this story, to the search for a coconut that the individual will wish to pluck. While searching, the individual is "unemployed." Having picked his coconut, he is "employed." His pay packet or profit finds expression in his chance of finding a partner with whom to swap coconuts, and if that chance comes to fruition, he will be able to eat a coconut that day. The better chance he stands of finding a partner, the higher will be his profit from picking the coconut in the morning.

The more coconut pickers there are on the island on any given morning who are looking for an opportunity to exchange their fruits in the afternoon, the better the chance of finding a partner for the swap. If the prospect of finding a partner is high, each individual will be more motivated to make the effort to pick the higher-hung coconuts, too. Yet if there is only a slight chance of finding a partner for the swap and satisfying one's appetite for a coconut, each individual will prefer to preserve his strength and look for low-hanging coconuts.

Thus numerous Nash equilibria are possible in this model. If none of the islanders picks coconuts, it is worth nobody's while to do so on his own account, because he will never be able to exchange them and get to eat coconuts.

A different situation is possible in which people attempt to pick only low-hanging fruit. Only a small number of coconuts is picked in this sort of situation, and therefore the chance of finding a partner for the swap is low. Hence, the incentive for making an effort and picking coconuts is low to begin with, justifying the islanders' tendency not to exert themselves to pick high-hung coconuts.

In a more successful equilibrium, people climb to pick the higher coconuts, too, on the expectation that the great effort involved will justify itself by providing a higher chance of finding a partner for the swap. In this equilibrium, this is indeed a self-fulfilling expectation, because a large number of coconuts, both high and low hanging, are picked on the island and many coconut pickers roam around the island in the afternoon seeking a partner with whom to exchange the fruits of their labors.

[14] As, in fact, Diamond (1982 – see note 13) assumes.

The moral of this story is clear. A nation may fall into a "poverty trap" in which unemployment surges and few entrepreneurs establish new businesses. Potential entrepreneurs fear there will be no demand for new produce, because the low-income population cannot afford to buy it. So these potential entrepreneurs do not open new businesses, and no employment opportunities are created for the unemployed. The population as a whole remains mired in poverty and this state of affairs justifies the entrepreneurs' fears.

At an equilibrium of prosperity, entrepreneurs expect that employees will earn high salaries and will want to spend their money on purchasing numerous goods. The entrepreneurs therefore proceed to new business initiatives, offering employment to most inhabitants, and the latter do indeed step up their consumption accordingly.

The model of Diamond (1982 – see note 13) is formulated in continuous time. We will now present a simpler model of a game that describes some of the ideas in the article.

Every afternoon, $0 \le e \le 1$ is the ratio of islanders who are "employed," which is to say, they roam around carrying a coconut they picked that morning and try to find a partner in a similar situation with whom to swap coconuts. Their chance of finding a swap partner is given by $b(e)$, where:

$$b : [0,1] \to [0,1]$$

is an increasing function: the higher the number of employed persons, the better, too, the chances of finding a partner that day. Of course, $b(0) = 0$ is satisfied; if nobody seeks a partner, no partner can possibly be found. For simplicity's sake, we will assume in this discussion that the function b is given by:

$$b(e) = e$$

In the morning hours, all the islanders are out looking for fruit-bearing coconut palms. All palms are the same height, which is also the measurement unit used by the islanders. In other words, the height of every palm tree is "1." Every day, one coconut ripens on each tree. The height of the ripe coconuts h on the trees is uniformly distributed up the tree. A person climbing to a height h in order to pick a ripe coconut invests effort $c(h)$ for that purpose, while the function:

$$c : [0,1] \to R_+$$

is an increasing and convex[15] function that assumes positive values. The effort $c(h)$ is expressed in terms of the prospect $b(e)$ of finding a partner with whom to swap coconuts. The effort is worthwhile if:

[15] That is to say, the second derivative c'' is not negative. The convexity of the function expresses the assumption that the coconut picker gets tired as he climbs, each additional yard he has to climb being at least as hard for him to climb as the previous one he has already climbed.

$$c(h) \leq b(e)$$

but is not worthwhile, from the point of view of the islander, if:

$$c(h) > b(e)$$

We will look for a symmetric Nash equilibrium, in which on all days the proportion e^* of persons "employed" on the island of an afternoon remains constant, and all the morning's "job seekers" on the island adopt the following threshold strategy: they will climb the coconut palm they have found only if the (ripe) coconut on it hangs at a height of not more than h^*. In other words, h^* satisfies:

$$c(h^*) = b(e^*)$$

such that for every coconut at a lower height, $h \leq h^*$ the inequality $c(h) \leq b(e^*)$ is satisfied, and the islander considers his climbing effort to have paid off. Since we have assumed that $b(e) = e$,

$$c(h^*) = e^*$$

will be satisfied at equilibrium.

We will assume that every "job seeker" finds, in the course of his searches on a given morning, just one coconut palm. Since the height of the ripe nuts on the island's trees is uniformly distributed, the chance of finding a coconut at a height that does not exceed h^* is h^*, and therefore this will also be the ratio of "employed persons" e^* proffering coconuts in the afternoon:

$$e^* = h^*$$

If we add this to the previous equation, we will obtain that every equilibrium of the type we seek is bound to satisfy the equation:

$$c(h^*) = h^*$$

Thus, the number of equilibria and the nature thereof in this model depend on the properties of the effort function $c(h)$. For example, if:

$$c(h) = h$$

then the game has a continuum of different equilibria: for every $h^* \in [0,1]$ it may be that the convention on the island is that in the morning, one climbs to a height of h^* in order to pick coconuts. This convention justifies itself because it brings about a situation in which the chance of finding a partner for swapping coconuts in the afternoon is likewise h^*, and given that chance, every islander deems it not worth his while to make an effort to pick coconuts that are higher than h^*.

Of all these equilibria, the worst is the one in which $h^* = 0$. In this equilibrium, none of the islanders picks any coconuts, because every one of them (rightly) believes that he will not find a partner with whom to swap coconuts in the afternoon. The most efficient equilibrium is the one in which $h^* = 1$. In this equilibrium, every islander picks a coconut in the morning, at whatever height he finds it on the tree, because he is convinced that he will certainly be able to exchange his coconut in the afternoon. This belief is indeed justified, since all the islanders pick coconuts in the morning hours.

Assume now, alternatively, that the effort function is given by:

$$c(h) = 2h^2$$

We have seen that:

$$c(h^*) = h^*$$

must be satisfied at equilibrium, which is to say:

$$2(h^*)^2 = h^*$$

This equation has two solutions:

$$h_1^* = 0, \quad h_2^* = \frac{1}{2}$$

The equilibrium $h_2^* = \frac{1}{2}$ at which half the coconuts are picked every morning is, of course, more efficient than the equilibrium $h_1^* = 0$ at which no coconuts are picked at all.

Finally, we will assume that the distribution of ripe coconuts on the palm trees is non-uniform, but is given by the cumulative distribution function P, which may be any increasing function:

$$P : [0, 1] \to [0, 1]$$

which satisfies $P(0) = 0$ and $P(1) = 1$.

At an equilibrium in which coconuts are plucked every morning up to a height of h^*, in the afternoon a proportion $e^* = P(h^*)$ of the inhabitants are looking for swap partners, and this is also the prospect of their finding a swap partner at that time (since we have assumed that $b(e^*) = e^*$). Therefore, the equilibria h^* of the game are the solutions to the equation:

$$P(h^*) = c(h^*)$$

If the cumulative distribution function inflects several times as it increases in the range of [0,1], it can also intersect several times with the convex function $c(h^*)$ in this range, and each of these intersection points will be a Nash equilibrium. The highest intersection point will also be the most efficient equilibrium.

For example, the cumulative distribution function:

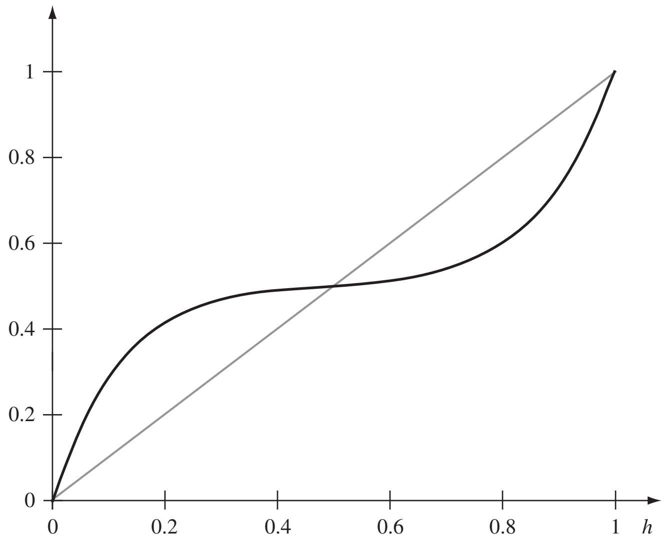

Figure 9.2

$$P(h) = \frac{1}{2} + 4\left(h - \frac{1}{2}\right)^3$$

intersects the effort function:

$$c(h) = h$$

at three points (see Figure 9.2):

$$h1^* = 0, \quad h_2^* = \frac{1}{2}, \quad h_3^* = 1$$

Each of these three points is a Nash equilibrium of the game.

PART IV

Uncertainty and mixed strategies

INTRODUCTION

In Chapter 10 we start addressing the issue of uncertainty. When a player is unsure what strategies his rivals will choose, we will assume that the player attaches a probability to each of the possible choice combinations. Each choice of a strategy of his own then defines the probability with which each profile of all the players' choices would be realized. Overall, each of the player's strategies defines a **lottery** over the strategy profiles in the game.

In order to decide which strategy to choose, the player then has to figure which of the lotteries induced by his strategies he prefers. We will assume that the player's preference over these lotteries is expressed by the **expected utility** accrued by this lottery – the weighted average of his utilities from his choice and the others' choices, weighted by the probabilities that he ascribes to the other players' choices. This assumption means that the utility levels now have a **cardinal** (rather than just **ordinal**) interpretation. Preferences over lotteries which can be represented by an expected utility over outcomes are named after von Neumann and Morgenstern, who isolated four axioms on the preference relation which obtain if and only if an expected-utility representation of the preferences is feasible.

These axioms do not always obtain. We bring in the example of the Allais Paradox for preferences that seem "reasonable" but which nevertheless cannot be represented by an expected utility.

We then apply the notion of expected utility to define **risk dominance** in 2×2 games. The risk-dominant strategy of a player in such a game is the one which yields him a higher expected payoff assuming that his rival chooses each of her two possible strategies with the same probability ½. A risk-dominant equilibrium in a 2×2 game is one in which both players choose their risk-dominant strategy. We observe that in the Stag Hunt game, it is the inefficient equilibrium which is risk dominant.

In Chapter 11 we launch the study of **mixed strategies**. A mixed strategy is a choice of the player among his strategies made using a lottery with specific probabilities. Extending the game by allowing the players to choose mixed strategies is called the **mixed extension of the game**. Within this extension, the original strategies (chosen with certainty) are called the **pure strategies**.

Matching Pennies is an example of a game which has no Nash equilibrium in pure strategies, but which has a Nash equilibrium in mixed strategies. In fact, by a theorem of Nash, every (mixed extension of a) matrix game has an equilibrium in pure or mixed strategies. In the appendix to the chapter we provide the outline of the proof, which relies on a fixed-point theorem.

If a player employs a mixed strategy at equilibrium, then she must be indifferent among all the pure strategies that she mixes (otherwise she would be better off choosing the one among them yielding her the highest expected payoff). This idea is somewhat counter-intuitive, and for some games this may be the source of failing to reproduce the mixed-strategy equilibrium behavior at the lab.

Mixed-strategy Nash equilibria have several potential interpretations. First, the mixed strategy of a player at equilibrium could be interpreted not as a mindful randomization but rather as the probabilistic prediction made by the player's rivals. Another interpretation would be to understand mixed strategies as simple, history-independent rules of thumb for behavior in repeated games. These could be relevant in games in which surprising the opponent is of value, such as service aces in tennis or penalty kicks in soccer. Indeed, both examples were empirically investigated to check for the use of mixed strategies by professional players, and partial positive evidence was discovered. Yet another interpretation is that for the strategic encounter each player is drawn at random from a large population, and knows only the characteristics of the average behavior of the population she is facing rather than that of the individual representative with whom she was matched. In this interpretation of a **population game**, each player chooses a determinate pure strategy, but different individuals in the population choose different pure strategies, and the mixed strategy represents the frequencies with which the different pure strategies are chosen. We offer an example in which this interpretation is particularly plausible.

In Chapter 12 we study **strictly competitive** two-player games, in which the interests of the players are diametrically opposed – whenever a player prefers a strategy profile over another, her opponent has the reverse preference. A particular instance of strictly competitive games are **zero sum** games, in which for each strategy profile the payoff of each player is just minus the payoff of her rival.

A **security** or **maxmin strategy** is one which maximizes the player's payoff under the assumption that for each strategy she may choose, her opponent will choose the strategy which will minimize her payoff. This notion is particularly relevant in strictly competitive games – in such games the opponent will indeed wish to minimize the player's payoff not due to mere cruelty but rather simply with the view of maximizing *his* own payoff.

In general, the payoff guaranteed to the player by her security strategy might be lower than her payoff at a Nash equilibrium. However, in the particular case of strictly competitive games, mixed-strategy equilibrium strategies are also security strategies, and Nash equilibrium payoffs are the same as the maxmin payoffs. Moreover, the **minimax theorem** asserts that in the mixed extension of a zero sum game, each player has a mixed security strategy, which guarantees her the same payoff she could get if she were to best reply to each particular mixed strategy of her opponent, while given this optimal behavior the opponent were to choose his strategy most spitefully to her.

Chapter 13 brings further elaborate examples of mixed strategies in general games. The first example is the Volunteer's Dilemma, in which out of a pool of potential volunteers each individual has to choose whether to volunteer for a costly mission (for which he would opt if he were the only potential volunteer) or wait and hope that somebody else will volunteer instead. Somewhat like a public good game, in the symmetric mixed-strategy equilibrium of this game each potential volunteer volunteers with a probability smaller than 1, which tends to zero as the number of potential volunteers increases; moreover, as the pool of volunteers gets larger, the overall probability that at least one person would volunteer decreases. However, laboratory experiments of this game show more optimistic outcomes than this theoretical prediction.

A further example, in which each player has more than two strategies, is the Rock-Paper-Scissors game, whose only equilibrium is in mixed strategies. Finally, the appendix to Chapter 13 elaborates an additional patent-race model, in which the competitors may choose one out of a finite number of R&D intensity levels using a mixed strategy. The mixed-strategy equilibrium is studied, and its properties compared with the findings of laboratory experiments.