

Chapter 9

Algorithms for Auctions and Games

Vincenzo Bonifaci and Stefano Leonardi

Abstract Economics is one of the diverse aspects of our life in which algorithms play a – perhaps subtle – role. Whether we are buying an item through an eBay auction, or choosing a pricing scheme from our telephone company, we are participants in distributed decision-making processes having consequences that may affect our personal “welfare”. And with the advent of the Internet and the Web, there is an increasing demand for algorithms that run, manage or analyze such economic transactions. In this chapter, we discuss some of the basic notions that underlie an algorithmic view of economic and strategic interactions.

9.1 Introduction

The exchange of ideas between the theory of algorithms and the economic theory of games is one of the most fascinating facets of computer science in the last decade. Such an encounter originated from the development of the Internet and the World-Wide Web (see Chap. 5) as organic, self-organizing systems, whose evolution is guided by a multitude of independent agents that operate according to economic principles. Computer science has taken inspiration from the economic theory of games in order to understand the economic mechanisms that motivate the agents of the network; in fact, the strategic interaction between agents is now considered an essential aspect of algorithm design in such contexts.

V. Bonifaci (✉)

Istituto di Analisi dei Sistemi ed Informatica “Antonio Ruberti”,
Consiglio Nazionale delle Ricerche, viale Manzoni 30, 00185 Roma, Italy
e-mail: vincenzo.bonifaci@iasi.cnr.it

S. Leonardi

Dipartimento di Ingegneria Informatica, Automatica e Gestionale,
Sapienza Università di Roma, via Ariosto 25, 00185 Roma, Italy
e-mail: leon@dis.uniroma1.it

The Internet is a decentralized system in which agents are independent and able to take individual decisions. In such scenarios it is important to foresee the states that the system is going to enter, since it is normally impossible to enforce some particular, favorable state. Consider, for example, the selection by the network agents of the path and transmission speed when transferring data from some origin to some destination. The strategy of each agent can consist in selecting, for each transmission, the path with the smallest estimated delay. Then we may, for example, ask whether the system will reach a state in which congestion is much larger than that of a state in which routing strategies are centrally planned. Or whether the system can oscillate from state to state due to the continual change of the strategies of the agents. The interaction between agents makes the adoption of the points of view of the theory of games unavoidable, and in particular calls for the notion of equilibrium states, those states of the system from which no agent has an incentive to deviate.

In recent years, several methodological contributions have been provided by algorithmics to economics and to game theory in particular. First of all, we remark that the theory of algorithms, mathematical optimization and the theory of computational complexity are tackling the delicate questions of the existence of equilibria and of the hardness of computing them. Secondly, a study has begun of the inefficiency of the solutions determined by the equilibrium states, in other words a quantitative comparison between the equilibrium solutions and the best solutions that could be imposed by a centralized “enlightened dictatorship”. Finally, given a view of the Internet as an open and democratic system aiming to guide the agents towards solutions that are more favorable towards the overall community, algorithms and protocols have been proposed that include incentives and penalties.

Computer science and network science also bring an important contribution to economic sciences, since the ever larger and pervasive diffusion of the Internet determined the migration of several economic and commercial activities on the net and created more of them, including ones that were unthinkable before, things like electronic commerce and online computerized auctions with a multitude of participants (eBay, Amazon, etc.). This determines the need to design algorithms and software that manage the commercial transactions electronically. Another important aspect has been the advent of digital goods in unlimited supply and the availability of commercial spaces on digital media such as the Web, blogs, forums and social networks.

We would like to stress one more aspect of the encounter between computer science and economics. The Internet and the World-Wide Web can be viewed as formed and governed by the action of a multitude of autonomous, rational agents that operate with the goal of optimizing their own individual “payoff functions”. Network algorithms and protocols have to operate on data supplied by these agents. But such data could be manipulated by the agents in order to optimize their own utility. For example, an algorithm searching for shortest routes along the Internet will have to acquire as input the costs for traversing (sub)networks, costs that are estimated and declared by autonomous systems, the independent economic entities that manage the networks. The traditional approach to algorithm design should then

be rebuilt to accommodate the issue of input data manipulation from the part of rational agents. The fundamental question there is whether it is possible to design algorithms that efficiently solve some algorithmic problem and at the same time give incentives to the agents to reliably report that part of the input data that is their private information. *Mechanism design* offers in this direction an essential methodological reference and an important contribution from economic sciences to computer science.

The arguments that we discussed should have convinced the readers that the encounter between algorithms and economic mechanisms is a necessity imposed by the Internet, an information system that in the last two decades revolutionized the modes of production and communication. The fact remains that the fundamental principles and the quality properties that are required by algorithms and economic mechanisms are deeply different, and it is not at all clear that for one of the two disciplines it should be possible to relax them while adopting those from the other discipline. In the case of algorithms, such requisites are usually expressed in terms of the computational resources necessary for execution or in terms of the quality of the computed solution. In the case of economic mechanisms, these requisites are expressed in terms of objectives such as the equilibrium between demand and supply, the maximization of the agents' profit or utility and, finally, the impartiality of a recognized authority with respect to all the participants of some economic game. Such an integration requires, then, a deep knowledge of both fields and the development of new models, algorithms and mechanisms, as well as methods for their analysis.

In this chapter we intend to present some fundamental concepts of the theory of games and some examples that illustrate how algorithms inspired by the theory of games are essential for many economic activities that happen daily on the Internet.

9.2 Games and Solution Concepts

9.2.1 Prisoner's Dilemma

To better understand the approach of game theory to the analysis of conflict situations, consider the following famous scenario, proposed in 1950 by American mathematicians Merrill Flood, Melvin Drescher and Albert W. Tucker. Two crime suspects, R. and C., have been captured by the police and have been charged with a major offense. In order to acquire the necessary evidence, the prosecutor is trying to separately persuade each of the suspects to confess his crime. The suspects are held in separate cells, and no communication can occur between them. The prosecutor proposes the following to each: if the prisoner confesses, his jail sentence will be cut in half – but only if his accomplice does not also confess, otherwise there would be no choice but to convict both. If the prisoner does not confess, he will get the maximum possible time in jail if the prosecutor is successful; otherwise, both suspects will be convicted anyway for some other minor offense.

Table 9.1 Prisoner's Dilemma

	Confess	Silent	
Confess	-5	-5	-10
Silent	-10	0	-1

We reach the situation summarized in Table 9.1. The action selected by R. determines the row of the table, the one selected by C. the column; this is why R. is also called the Row player and C. the Column player. In each of the four entries of the table, the bottom left number denotes the payoff (also called *utility*) of R. and the top right number the payoff of C. Most numbers are negative since they denote the penalty due to the number of years to be spent in prison: serving 10 years is worse than serving 5. A zero means that the criminal was let free.

What will each player do? Imagine taking R.'s role. Observe that if C. confesses, it is better to confess, since you'll have to serve 5 years instead of 10. And if C. doesn't confess, it is better to confess too: you'd be let free instead of spending a year in jail. So no matter what C. does, confessing is a better option than not confessing! An action with such a property – the property of being preferable to all other actions, independently of the other players' choices – is called a dominant strategy.

Looking at the table, one sees that the roles of R. and C. are symmetrical, therefore confessing is a dominant strategy for C. too. Thus, the overall outcome is that both prisoners will have to spend 5 years in jail, more than each of them would have if they both chose not to confess. This result may sound paradoxical, but several experiments in psychology and sociology show that outcomes of this kind are actually not uncommon in practice.

It should be intuitively clear that the Prisoner's Dilemma is not limited to the simple scenario outlined above. It has been applied in the most diverse contexts and several examples can be found also in telecommunication and computer networks. A more general situation is the setup and usage of some public good. Say that a public good is worth 3€ to each of the two players, Row and Column. Some investment is necessary to create the good, so the good will be available only if 4€ total are invested. Each player may volunteer to invest or not: if he is the only volunteer, he will pay 4€, otherwise the cost is shared equally. The payoff of a player is the value of the good (or zero, if the good is not created), minus the investment cost. The resulting payoff matrix is given in Table 9.2.

The situation of each player is similar to that in the Prisoner's Dilemma: there are two choices, one of which is better from an opportunistic point of view, but that damages the other player. Although numerically different, the payoff table has in fact the same structural properties of the payoff table for the Prisoner's Dilemma: namely that the more aggressive option is, for both players, a dominant strategy. Therefore, the theory predicts that no investment will be made.

Table 9.2 Provision of a public good

	Free ride	Invest	
Free ride	0	3	-1
Invest	-1	1	1

Table 9.3 Battle of the Sexes

	Movies	Concert	
Movies	3	1	4
Concert	2	4	3

A final remark on the pessimistic outcome that game theory predicts in the Prisoner’s Dilemma. The result is based on several assumptions that are not necessarily realistic: that the players are rational; that they care exclusively about their own good; and that the game is played exactly once. In the absence of one or more of these conditions, it is possible to have outcomes that are different and more optimistic. In fact, Merrill Flood, one of the discoverers of the Prisoner’s Dilemma, once made the following proposal to a secretary of the institute he was working in: she could either immediately accept 100 dollars, or be given 150, but in the second case she would need to find an agreement on how to split the money with a second secretary that had been kept unaware of the proposal. The secretary chose the second option and agreed with her colleague to get 75 dollars each, even though she could easily have kept the 100 dollars without her colleague ever getting to know anything.

9.2.2 Coordination Games

In some strategic situations, conflicts may arise simply due to the impossibility for the players of the game to coordinate their choices. Consider another scenario, known in game theory as the “Battle of the Sexes”. A man and a woman are deciding how to spend the evening out: the man would like to go to the movies, while the woman would prefer a concert. It is too late to consult and each of them should reach the chosen location directly from his or her office. Both of them prefer to spend the evening together rather than alone, but, given the choice, the man would prefer the movie while the woman would prefer the concert. A possible payoff table for the game is reported in Table 9.3, where the values denote again the utility of each action combination.

What does game theory predict in such a case? The answer is less clear-cut than that we saw for the Prisoner’s Dilemma. If the woman knew for certain that the

Table 9.4 Rock–Paper–Scissors

	Rock	Paper	Scissors
Rock	0	−1	1
Paper	1	0	−1
Scissors	−1	1	0

man would go to the movies, she would choose the movies over the concert; if she knew that the man would go to the concert, she would be happy to join him there. In other words the best response to the man’s action of going to the movies is for the woman to go to the movies, and the best response to the man’s action of going to the concert is for the woman to go to the concert. A similar reasoning holds of course for the man. The two solutions (Movies, Movies) and (Concert, Concert) are therefore “stable” in the sense that in each of them the action chosen by each player constitutes a best response to the other player’s choice.

Such “mutual best response” outcomes are called Nash equilibria, from the name of the American mathematician and Nobel prize for Economics winner John F. Nash, who introduced and studied them. In the Battle of the Sexes, the only equilibria are those in which the woman and the man spend the evening together: for example, the outcome in which the woman goes to the concert and the man goes to the movies is not a Nash equilibrium, since the man could improve his payoff by “changing his mind” and going to the concert instead.

A small weakness of the Nash equilibrium as a solution concept of a game is that it is not necessarily unique, as can be seen from the example of the Battle of Sexes, in which there are in fact two Nash equilibria. On the other hand, this ambiguity captures the uncertainty of outcomes in a conflict situation.

9.2.3 *Randomized Strategies*

The game known as Rock–Paper–Scissors is a simple and popular two player game. The players simultaneously represent, by a gesture of their hand, an object among paper, scissors and rock. Rock beats scissors, scissors beats paper, and paper beats rock. If the players choose the same object, the game is tied. A player gets payoff 1 if she wins, −1 if she loses, and 0 if there is a tie. Therefore, the sum of the two payoffs of the players, whatever the outcome of the game, is always zero; a player can win only as much as her opponent loses – something that was not the case with other games such as the Prisoner’s Dilemma or the Battle of the Sexes. Such a game is called *zero-sum*. In a two-player zero-sum game it is sufficient to specify the payoff of one of the two players, say the Row player, as in Table 9.4.

The Rock–Paper–Scissors game has an interesting property: it admits no Nash equilibria. The reason is that the best response of a player to a given action of the opponent renders the action of the opponent a “bad response”: if we knew our opponent would play paper, we would play scissors, but then our opponent would play rock, etc. Therefore there is no Nash equilibrium in the sense we previously discussed.

What can game theory say about such cases? The key idea lies in generalizing the notion of strategy to allow probabilistic or *mixed* strategies (see the box “Events, Probabilities and Expected Values”). An example of mixed strategy is: play paper with probability 70 %, scissors with probability 20 % and rock with probability 10 %.

As a consequence of this generalization, the concept of payoff of a player is substituted by that of expected payoff, which is simply the average utility obtained by choosing the actions according to the prescribed probabilities. If, for example, we play 70 % paper, 20 % scissors, 10 % rock, and our adversary answers with 100 % scissors, our expected utility will be

$$0.7 \times (-1) + 0.2 \times 0 + 0.1 \times (+1) = -0.6.$$

To remark on the difference with such mixed strategies, the strategies in which each player selects a single action are called pure strategies. A pure strategy can be seen as a very special case of mixed strategy.

Events, Probabilities and Expected Values

Mathematically, the probability of an event A is a real number between 0 and 1, often denoted as $\Pr[A]$. When A is an impossible event, $\Pr[A] = 0$, while when A is certain, $\Pr[A] = 1$.

The *complement* event to A (which occurs if and only if A does not occur) has probability $1 - \Pr[A]$.

When A and B are *independent* events, the probability that they simultaneously occur is equal to $\Pr[A] \cdot \Pr[B]$. For example, the probability of obtaining two heads when tossing two coins is equal to $1/2 \times 1/2 = 1/4$.

When A and B are *mutually exclusive* events, the probability that at least one of them holds is equal to $\Pr[A] + \Pr[B]$. For example, the probability to obtain a 5 or a 6 after throwing a standard six-faced die is equal to $1/6 + 1/6 = 1/3$.

Finally, when a numerical variable X can take the values x_1, x_2, \dots, x_n , with probability p_1, p_2, \dots, p_n , respectively, its *expected value* is given by the formula

$$p_1 \cdot x_1 + p_2 \cdot x_2 + \dots + p_n \cdot x_n.$$

The conceptual leap from pure strategies to mixed strategies has an interesting and somewhat unexpected consequence: any game that is finite, in other words any game that has a finite number of players and a finite number of strategies, always admits at least one mixed strategy equilibrium. Such a fundamental property was proved by John Nash in 1950 and has since been known as Nash’s theorem.

What is, then, a Nash equilibrium in mixed strategies for the Rock–Paper–Scissors game? Assume we play the three actions at random, each with probability $1/3$. If our opponent plays paper with probability p , scissors with probability s , and rock with probability r , his expected utility will be

$$\begin{aligned} & (1/3) \times 0 \times p + (1/3) \times 1 \times p + (1/3) \times (-1) \times p \\ & + (1/3) \times (-1) \times s + (1/3) \times 0 \times s + (1/3) \times 1 \times s \\ & + (1/3) \times 1 \times r + (1/3) \times (-1) \times r + (1/3) \times 0 \times r \\ & = 0. \end{aligned}$$

Therefore the expected payoff of our opponent will be zero, independently from the mixed strategy that he chooses. If he, too, is choosing with the same probabilities rock, paper and scissors, then the expected payoff of both players will be zero and neither of them will have a way to obtain an expected payoff larger than zero by adapting their strategy. In other words, we will have a situation of mutual best response, exactly as required by the definition of Nash equilibrium. In fact, in this case the outcome we described is the only Nash equilibrium of the game, although this uniqueness property does not hold in general, as we saw in the Battle of the Sexes example.

The reader may wonder whether actual Rock–Paper–Scissors games are played as the theory prescribes. It is certainly true that if one follows the optimal strategy, then it is impossible to achieve a negative expected utility (although it is always possible to be particularly unlucky!). However, many “real-world” players, due to limitations of several kinds, such as not being able to choose the actions in a perfectly random fashion, may not be playing an optimal strategy but some slightly different mixed strategy. In this case we might obtain an expected payoff larger than zero by adopting a mixed strategy different than the perfectly uniform one. For example, if our opponent plays scissors and paper with the same probability, but never plays rock, then we can increase our expected utility by always playing scissors. Our expected payoff then becomes positive, since

$$1 \times 0.5 \times 0 + 1 \times 0.5 \times 1 = 0.5.$$

For this reason, in games such as Rock–Paper–Scissors, the players that more often succeed are those that try to learn the adversary’s strategy, at the same time adapting their own strategy to it.

9.2.4 *Hawks and Doves*

So far, we exemplified potential conflict situations through some simple but paradigmatic two-player games: the Prisoner’s Dilemma, the Battle of the Sexes

Table 9.5 Chicken

	Swerve	Keep
Swerve	0	1
Keep	-1	-100

and Rock–Paper–Scissors. We add an equally interesting game to the list, called “Chicken”. The scenario is as follows: two contestants start driving cars at each other at high speed. The first of the two that, by swerving, deviates from the collision trajectory will be the loser, in other words the “chicken”. The player that resists swerving until the end will instead be the winner. Clearly, a participant in Chicken should try to resist as much as possible in order to overcome his opponent, but if both players act out this idea, the result is a disaster!

Analyzing the game (Table 9.5) in light of the concepts discussed so far, we observe that there are only two Nash equilibria in pure strategies: the one in which the first player swerves and the second keeps the trajectory, and the one in which the opposite happens. Is there by any chance some Nash equilibrium in which the strategies are all probabilistic? Indeed there is. Denote by R_s and R_k the probability with which the first player (the Row player) swerves or keeps the trajectory, respectively. If we consider the configuration in which $R_s = 99\%$ and $R_k = 1\%$, then we notice that

$$0 \times R_s + (-1) \times R_k = 1 \times R_s + (-100) \times R_k = -1.$$

This implies that the expected payoff for the second player is the same (-1) independently of the action he chooses. Similarly, if $C_s = 99\%$ and $C_k = 1\%$ are the probabilities with which the second player swerves or keeps the trajectory respectively, we have

$$0 \times C_s + (-1) \times C_k = 1 \times C_s + (-100) \times C_k,$$

so that also the expected payoff of the first player is the same independently of his action. Such a combination of mixed strategies therefore constitutes a Nash equilibrium.

Games akin to Chicken are those in which the players have to decide whether they want to be aggressive “hawks” or peaceful “doves”. The Cuban missile crisis of 1962 was analyzed by the advisors of US President J.F. Kennedy and classified exactly as a situation of this kind, where the catastrophic outcome could have been a nuclear war. Kennedy decided to let USSR Chairman N. Khrushchev know that the United States would not have played the dove’s role, even to the point of having to enter the war. Luckily for everybody, Khrushchev chose to concede and be a “dove”.

The Chicken game also allows us to introduce an important extension of the notion of Nash equilibrium, proposed by the Israeli mathematician Robert

J. Aumann (also a Nobel prize for Economics) in 1974. In this extension, known as correlated equilibrium, we consider acceptable equilibria all outcomes of the game in which players do not have an incentive to change their own strategy, assuming that the strategy that the players follow is suggested in some way by a trusted third party. For example, consider a third party that with probability $1/2$ suggests to the first player to swerve and to the second player to stay on track, and with probability $1/2$ gives the opposite suggestions. In this case, if one of the two players assumes that the other one behaves according to the suggestion of the third party, he will have, in fact, an incentive to behave exactly as suggested; in other words, the suggestion is self-enforcing. This is not at all a phenomenon that is far removed from daily experience: a traffic light is nothing but such a trusted third party in the Chicken-type game that any two car drivers play when they have to go through the same crossing from different directions.

9.3 Computational Aspects of Game Theory

We have seen through the previous example how rigorous theorems such as Nash's theorem ensure, for several games, the existence of the equilibria we are interested in. We can ask ourselves whether these equilibria can actually be computed via some algorithm: knowing that an optimal strategy for a given game exists does not help us much if we do not know how to determine it. For this reason the computational aspect of game theory is crucial. We may even go further and say that if a given type of equilibrium is hard to compute, then most likely that type of equilibrium does not fully capture the realistic outcomes of a game, since after all the participants in a game, be they individuals or companies or specifically designed software, always have a limited computational power. The computational properties of equilibria therefore help us understand what equilibrium concepts are indeed more realistic.

9.3.1 *Zero-Sum Games and Linear Optimization*

In the case of games where players move simultaneously, the computation of equilibria is more or less complicated depending on the type of equilibrium that is sought. For example, in the case in which one seeks a Nash equilibrium in pure strategies, it is indeed sufficient to consider all combinations of the players' strategies and for each combination verify whether both players are playing a best response action. If the first player has m actions available and the second player has n , we have to consider $m \times n$ cases, and for each of those we must compare the payoff with other $m + n - 2$ values of alternative payoffs. As we have already seen, though, an equilibrium in pure strategies might not exist.

More interesting is the case of Nash equilibria in mixed strategies. If we consider arbitrary games (even with only two players) things get complicated pretty quickly.

Table 9.6 An example of a zero-sum game

	C	D
A	2	-1
B	1	3

We thus start from the relatively simpler case of zero-sum games; this case was analyzed for the first time in 1928 by John von Neumann. As we have seen with the Rock–Paper–Scissors game, zero-sum games can be described by a single matrix in which the entries indicate at the same time the payoff for the first player and the cost for the second player, as in Table 9.6 where, for example, if the first player chooses strategy *A* and the second strategy *C*, then the first player gains 2 and the second pays 2.

The basic idea behind von Neumann’s result is that an equilibrium strategy should be stable even when it has been revealed to one’s opponent. If the mixed strategy of the Row player is known, the Column player will select a mixed strategy that minimizes the payoff of the Row player. By foreseeing this, the Row should then select a strategy that minimizes his own maximum loss, or in other words that maximizes his minimum gain. Such a strategy his called a *minimax strategy*: some mathematical details are given in the box “Minimax Strategy”. The main point is that in this case the equilibrium strategies can be computed efficiently.

Minimax Strategy


A minimax strategy can be expressed by a system of inequalities, where the variables *a* and *b* represent the probability with which the player chooses his strategy *A* or *B*, respectively, and where variable *v* represents the gain obtained by the player.

$$\begin{aligned}
 \max \quad & v \\
 & 2a + b \geq v \\
 & -a + 3b \geq v \\
 & a + b = 1 \\
 & a, b \geq 0.
 \end{aligned}$$

The first two conditions ensure that no matter the answer of the Column player, the payoff of the Row player will be at least *v*. The other conditions simply ensure that a solution of the system represents a mixed strategy. By solving the system – for example graphically (see Fig. 9.1) – we obtain *a* = 2/5, *b* = 3/5, that is to say that the Row player should play action *A* with probability 40% and action *B* with probability 60%. That ensures an expected payoff *v* = 7/5 to the Row player.

(continued)

Fig. 9.1 Equilibrium computation for a simple zero-sum two-player game



(continued)

Clearly, a similar argument can work for the Column player too. In this case, if w represents the minimum loss by the Column player, we get $w = 7/5$. The fact that v equals w is a general property that is ensured by von Neumann's theorem: when both players follow a minimax strategy, the resulting outcome is an equilibrium. In general such systems of inequalities constitute instances of linear optimization problems, a well-studied field in mathematical optimization. Algorithms are known that solve such problems in polynomial time (see Chap. 2). The same algorithms can thus be directly applied to the computation of mixed strategy equilibria in zero-sum games with two players.

9.3.2 Fixed-Points: Nash's Theorem and Sperner's Lemma

What happens when we move from zero-sum games to the general case? Is it still possible to efficiently compute the Nash equilibria in mixed strategies? It turns out that no one has a definitive answer to such an apparently simple question.

A natural approach, when trying to understand how to find a Nash equilibrium in a non-zero sum game, is to take a step back and analyze the proof of Nash's theorem – in other words, to understand why an equilibrium in mixed strategies must exist. We will not give here the details of the proof, but we will hint at some of the ideas on which it is based.

Nash's theorem can be seen as a so-called fixed point theorem, that for a given function F asserts that, under suitable conditions, the equation $F(x) = x$ always admits a solution – in other words, that the function F has at least one fixed point.




Fig. 9.2 A valid coloring of a subdivided triangle ABC

In the case of Nash’s theorem, x represents a list of mixed strategies (one for each player), and the function $F(x)$ gives a new list of mixed strategies in which every player is adopting a best response to the configuration in which the other players choose their actions according to what is specified by x . The equation thus represents the fact that we seek a list of mixed strategies such that each is a best response to the other ones.

To illustrate one of the ideas on which the theorem relies, we will discuss a result used in the proof of the theorem and strongly linked to it, known as Sperner’s lemma. We will not formulate the lemma in its full generality. However, one of the consequences of this lemma is the following: take some triangle ABC and arbitrarily subdivide it into smaller triangles, as in Fig. 9.2. Then color the vertices of all small triangles with one of three colors (say white, black and gray), while obeying the following rules:

1. The three vertices A, B, C of the original triangle should be colored with three different colors;
2. The vertices that lie on a same edge of the original triangle should not be colored with the color assigned to the vertex opposite to that side; for example, vertices lying on the AB line cannot be colored with C ’s color;
3. The vertices in the interior can be arbitrarily colored.

Sperner’s lemma asserts that, independently of how we color the vertices, there will always exist a triangle of the subdivision whose vertices have three distinct colors. And, in fact, in the case of Fig. 9.3 such a triangle exists. Why is that?




Fig. 9.3 Illustration of the proof of Sperner’s lemma

To understand, mentally picture each triangle of the subdivision as a triangular room with three walls, one per side. Each wall has a door in it, if the corresponding side of the triangle has one white vertex and one black vertex. It is easy to check that such rooms may have zero, one or two doors, but never three. Moreover, if a room has a single door, it must correspond to a triangle with three differently colored vertices, because if that was not the case there should be an additional white–black side, and the doors would be two.

Observe that from the outside of the original triangle ABC there is an odd number of accessible doors (again, this can be seen to be always true). Take any one of this doors and follow the path through each door. Notice that we never have a choice, since the number of doors in a room on the path is one or two. The path will necessarily end either in a room with a single door, or again outside the ABC triangle. But in this last case the path has used two doors facing the outside of the triangle, and since there was an odd number of those doors, there must be another door from which we can proceed in a similar fashion. Thus the only way to end this process is to find a room with a single door, corresponding to the triangle with three distinct colors whose existence was claimed by the lemma. Figure 9.3 illustrates the argument.

The above argument may perhaps give a vague intuition of why finding Nash equilibria in mixed strategies appears to be difficult. The procedure to determine the triangle that verifies Sperner’s lemma is correct, but in general it may require very long paths, which in turn correspond to a long computation time. Nothing, however, prevents that in principle one could devise a smarter, more “direct” method of determining the correct triangle, even though that appears unlikely.

9.3.3 Mixed Nash Equilibria in Non-zero-Sum Games

There is an algorithm, although markedly inefficient, that allows us to determine a Nash equilibrium in mixed strategies for any two-player game. An idea at the basis of the algorithm is the notion of *support* of a mixed strategy. The support is simply the set of actions that in the mixed strategy considered are selected with probability larger than zero. In other words, it is the set of pure strategies that concur to form a given mixed strategy. Now, a property of equilibria is that a mixed strategy constitutes a best response if and only if all the pure strategies in its support are best responses. This fact is useful to determine a Nash equilibrium. In the box “Equilibrium Computation” a complete example is given, with all the corresponding calculations.

Equilibrium Computation

Consider the following example.

	C	D	
A	1	3	
	2	0	
Silent	2	1	
	1	4	

The pure strategies of the Row player are A and B , those of the Column player C and D . Let us ponder whether a Nash equilibrium exists in which the support of the first player is $\{A, B\}$ and that of the second player is $\{C, D\}$. Call a, b, c, d the probabilities assigned to the respective strategies. Then it should hold

$$\begin{aligned}
 a, b, c, d &> 0, \\
 a + b &= 1, \\
 c + d &= 1.
 \end{aligned}$$

Moreover, in order for (a, b) to be a best response to (c, d) , and since we are assuming that both A and B are in the support of (a, b) , A must be a best response to (c, d) and should thus give an expected payoff at least as large as that would be obtained with B :

$$2c \geq c + 4d,$$

(continued)

(continued)

but by a similar argument B should be a best response to (c, d) :

$$c + 4d \geq 2c.$$

We thus obtain $2c = c + 4d$, that combined with the previous equations gives $c = 4/5$, $d = 1/5$. Analogously we can deduce, from the fact that both C and D should be a best response to (a, b) , that $a + 2b = 3a + b$ and so $a = 2/3$, $b = 1/3$. Such an argument shows that the main difficulty in finding a Nash equilibrium is in the determination of the support sets of the two players. Once the supports are known, as we explained it is sufficient to verify that a given system of linear inequalities admits a solution. A possible algorithm to determine a Nash equilibrium then, however inefficient, consists in enumerating all possible pairs of supports and for each of them verifying if it gives rise to an equilibrium by solving the linear system. The running time of such an algorithm is dominated by the number of possible pairs of supports, that in the case of m actions for one player and n for the other one, is roughly 2^{m+n} .

9.4 Inefficiencies

9.4.1 *The Tragedy of the Commons*

In many of the examples discussed so far, we have seen how one can mathematically formalize the kind of individual behavior that emerges from the interaction of rational agents, each of whom is driven to his own goals. A natural question is: what happens to the overall system? What levels of “social welfare” are obtained when each agent is pursuing his own goal separately? For such questions to make sense we need to define what we mean by social welfare. There are many possible definitions that are equally valid. For example, with social welfare we could mean, in a utilitarian perspective, the sum of the payoff of all the players; or the payoff of the player with the smallest payoff (the “poorest”). In any case it should intuitively be clear that when the users of a system individually pursue their own goals, the resulting social welfare is not necessarily maximized. In other words, the solution determined by the agents will not in general be a globally optimal solution.

Consider, for example, a simple scenario in which a set of 100 users shares the access to the same Internet connection, by means of a certain finite transmission bandwidth B . Every user can regulate the amount of bandwidth that he intends to use. The strategy of the i th user consists in the fraction $x(i)$, with $0 \leq x(i) \leq 1$, of


the amount of bandwidth used (so that the user gets a bandwidth equal to $x(i) \cdot B$). Assume that the payoff of a user depends on the amount of bandwidth he consumes, but also on the amount of bandwidth left unused by the other players, according to the formula

$$u(i) = x(i) \times (1 - x(1) - x(2) - x(3) - \dots - x(100)).$$

The second term in this formula captures the fact that the latency of the communication channel is lower when the channel is less congested, while if the channel is almost saturated the latency is very high and thus the payoff quickly drops. (For simplicity we do not require that the sum of the $x(i)$ fractions be inferior to 1, so the second term can even become negative.) We thus obtain a multiple player game. Such a game is not finite – the number of players is finite, but the strategies are values from a continuous set and are thus infinite – so we cannot invoke Nash’s theorem directly.

Still, in this case a Nash equilibrium exists; we can find out that it corresponds to the solution in which each user has a fraction of the bandwidth equal to $1/101$. The communication channel in this solution is almost completely saturated, since it is used for a total fraction equal to $100/101$. The payoff of each user will then be $1/101 \times (1 - 100/101) = 1/(101)^2$. How large is the social welfare in this case? If by social welfare we mean the total utility of the players, we obtain a social welfare equal to $100/(101)^2$, roughly 0.01. However, if we could have forced each user to use a bandwidth fraction equal to $1/200$, half of the total bandwidth would have stayed unused, so that the payoff of each user would have been equal to $1/400$, and the total payoff equal to $1/4 = 0.25$. Such value is 25 times larger than the value obtained at the equilibrium (0.01), so that we can say that in this case the independence of the players caused a large decrease of the social welfare. Such a negative phenomenon is well-known in economics under the name of *tragedy of the commons*. It manifests itself every time that the individual interests of a group of users tend to destroy the advantages deriving from the use of a common resource. Clearly, not all economical interactions are of this kind and sometimes the individual interest gets close to the collective one, but such ‘tragedies’ are indeed frequent.

Although such phenomena were known since long ago, only more recently have researchers from the computer science community, starting with Elias Koutsoupias and Christos Papadimitriou, analyzed them in a deeper quantitative way, through the notion of *price of anarchy*: that is, the ratio between the globally optimal social welfare and the social welfare arising from the equilibrium (or between the social cost at the equilibrium and the optimal social cost). The closer such ratio is to one, the more reasonable it is to claim that the individual interest approximately coincides with the collective interest, while if the price of anarchy is very high – as in the case of the bandwidth sharing game – the outcomes in the two settings can be very different. The notion of a price of anarchy has been applied and continues to be applied in the study of many economical scenarios, in particular those related to networks.

Fig. 9.4 Pigou's example

9.4.2 Routing Games

Imagine a familiar scenario: we need to move by car from one end of the city to the other, and the choice of our itinerary will have deep impact on the travel time. Such time will also depend on how many other car drivers will choose the same itinerary, due to possible congestion effects. We thus see that the scenario can be modeled as a game in which the players are the drivers and the actions are the itineraries. What are the equilibria of the game, and how inefficient are they? The question is not futile, also because a similar model can be applied to computer networks, with the difference that the flows to be routed are flows of data instead of flows of vehicles. It should not be surprising then that the answer has been in part given by computer scientists, as well as by traffic engineers.


To understand it we first have to better specify our model. We represent the traffic network as a graph in which nodes are the locations of interest and the arcs represent the connections from one place to the others. Moreover, to describe the possible congestion effects, we assign to each arc a cost (or “latency”) function that gives the travel cost incurred for each level of traffic on that particular connection. For example, in the graph of Fig. 9.4 the upper arc has a constant cost function (function c_1), while for the lower arc the cost function is the identity (function c_2).

We finally select an origin point s , a destination point t and a flow value, that is, an amount of traffic to be routed from s to t . We can imagine that such flow is composed of an infinite number of “traffic particles”, each of which has no influence by itself, but that can behave independently (in this formulation we thus have a game with infinitely many players – the existence of an equilibrium is not ensured by Nash’s theorem, but can in fact be proved by similar means). The overall cost of a certain flow is obtained by summing, on each arc, the product between the quantity of flow traversing the arc and the latency of the arc. For example, if in Fig. 9.4 we had a flow equal to 0.2 on the upper arc and equal to 0.8 on the lower arc, the overall cost would be equal to $0.2 \times 1 + 0.8 \times 0.8 = 0.66$.

Suppose, for example, that the overall traffic flow is equal to 1 unit (say, 1,000 cars). If everybody used the upper link in the figure, the cost would be $1 \cdot c_1(1) = 1$. Such configuration is however not an equilibrium. Indeed, if nobody used the lower link, a car driver would find it convenient to leave the first route for the second, since the cost that he would find on the new link would be equal to $c_2(0.001) = 0.001$, which is much lower than 1.

In general, the traffic is in equilibrium if for each pair of paths P and P' from s to t in the network, such that the traffic along P is positive, it holds that the cost

Fig. 9.5 The two networks of Braess' Paradox



along P is not larger than the cost along P' . As a consequence, the costs of paths actually used in the equilibrium are all equal, and they are all less than or equal to the cost which a single driver would incur on any unused path.

If the traffic was equally split among the two links, the total cost would be $1/2 \times c_1(1/2) + 1/2 \times c_2(1/2) = 1/2 + 1/4 = 3/4$. But that would not be an equilibrium, since the flow along the upper link would be positive but the one along the lower link would have a cost smaller ($1/2$) than that of the upper link (1).

It can be checked that the only equilibrium in this example corresponds to the case in which the flow is only routed along the lower link. We thus obtain an overall cost equal to $1 \cdot c_2(1) = 1$. As seen above, this value is not optimal; in terms of the price of anarchy we have a ratio between cost at equilibrium and optimal cost equal to $4/3$.

The example just seen was discussed in the 1920s by the British economist Arthur C. Pigou. Afterwards it became clear that the individual behavior of the users can even give rise to some counterintuitive phenomena, as shown by the following example due to German mathematician Dietrich Braess (Fig. 9.5). Assume once again one unit of flow has to be routed. In the network on the left of the figure, the equilibrium flow is the one in which half of the traffic follows the path (s, v, t) and the other half the path (s, w, t) . The overall cost is then $(1/2) \times (1/2 + 1) + (1/2) \times (1 + 1/2) = 1.5$.

In the network on the right a 'superhighway' has been added that has zero cost and connects v and w . In the new network now, the only equilibrium is the one that routes all the flow along the path (s, v, w, t) . But this flow has overall cost equal to $1 \times (1 + 0 + 1) = 2$, larger than in the previous case! The apparently beneficial decision of adding a connection from v to w has thus given rise to a degradation of the system's performance. Such a phenomenon is now known as *Braess' Paradox*.

Anomalies notwithstanding, is it possible to give more encouraging results for such routing games? How large can, in general, be the price of anarchy of such games? In general, the answer depends not so much on the structure of the network, as on the type of cost functions. If such functions are linear, that is, of the form $c(x) = ax + b$, Tim Roughgarden and Éva Tardos have shown that the price of anarchy is never larger than $4/3$, so that things never get worse than in Pigou's example, even in networks that are much larger and more complex. On the other hand, if the cost functions have some "nonlinearities", then the price of anarchy can

be very high and it is not possible to bound it a priori. Informally, this corresponds to saying that things tend to get worse when roads have a certain maximum capacity and more easily act as bottlenecks.

9.5 Mechanism Design and Online Auctions

Mechanism design is the field of economics that is concerned with procedures for aggregating the preferences of rational agents among a set of possible economic choices, about the allocation of goods and the definition of prices. The goal of mechanism design is to define an algorithm that takes as input the agents' evaluations and returns as output the allocation of goods to the agents and the price that the agents should pay for the allocated good. Analogously, one can view computer science as interested in algorithms and protocols that, through the use of computationally limited resources, determine solutions of good quality to a problem for a given set of input data. Mechanism design is therefore an ideal meeting ground between economics and computer science. The algorithmic design of mechanisms is a research area originated by N. Nisan and A. Ronen that has flourished in recent years and that aims at formally defining the economic features that are algorithmically implementable and the computational issues that arise in the design of economic mechanisms. To introduce some of the main ideas in the algorithmic design of mechanisms we will refer to some simple examples in the area of online auctions.

9.5.1 The Vickrey Auction

Online auctions are a typical scenario in which the design of some economic mechanism is required. The goal of a mechanism for an online auction is the identification of an allocation of the goods to the participants and the definition of the price that each participant will pay. Such an algorithm has to be conceived with the goal of satisfying some desirable features both from the point of view of the auctioneer, such as the maximization of the profit obtained by selling the goods, and from the point of view of the users, such as the perception of the electronic commerce site as reliable and of the auction as fair.

Consider first a well-known example of auction, the Vickrey auction or second-price auction, that aims at selling one indivisible good, such as an art piece, to one of a set of n interested agents. The Vickrey auction assumes that the i th agent has his own evaluation v_i for the good being sold and that he communicates to the mechanism, by means of a sealed envelope, his offer b_i for the good. Observe that user i has some private information (the evaluation), while only the offer b_i is communicated to the mechanism, that will open the envelopes after all offers has been received from all participants. The mechanism has to determine the agent i that

will obtain the good and the price p_i that agent i should pay. The utility function of player j is $v_j - p_j$ in case he obtains the good at price p_j , or 0 if he does not obtain the good. We are interested in agents behaving rationally, that decide upon their strategy – the offer communicated to the mechanism – in a way that maximizes their own utility.

The strategy of an agent may in principle depend on the behavior of the other players, whose offer and valuations are however not known to the agent, and from the publicly known mechanism, deciding the allocation of the good and the selling price. What one aims to design is a mechanism for managing the auction that induces a dominant strategy for each agent, that is, a strategy that optimizes the payoff of each player independently from the behavior of other players.

An example of auction that does not induce a dominant behavior is the assignment of the good to the agent with the highest offer, for the price indicated in the offer. In such an auctions there may be several situations in which a lower offer from a player may lead to the allocation of the good for a lower price. Consider, for example, the case of two agents that offer for the good the values $b_1 = v_1 = 10\text{€}$ and $b_2 = v_2 = 20\text{€}$. Observe that in this example the offers coincide with the valuations. The strategy of the second player is not optimizing the agent's utility, since any offer above 10€ will allow to obtain the good for a lower price.

The Vickrey auction, or second-price auction, also allocates the good to the highest bidding player, but at the second highest price that is offered by any player. In the preceding example, player 2 would obtain the good at a price of 10€ . One can observe that in this case, independently of the offers of the other players, the dominating strategy is to communicate an offer equal to one's own valuation for the good being sold. The selling price, and thus the utility of the agents, are not affected if a different offer from an agent does not change the allocation of the good. A different agent could obtain the good by bidding an offer higher than his own valuation, but then he would incur a negative utility. An offer lower than the valuation might cause the loss of the good and therefore a reduction in utility. Agents have therefore no incentive in communicating a bid different from their own valuation of the good.

The Vickrey auction, for all its simplicity, allows us to reason about several important properties of mechanisms:

- *Social welfare*: Firstly, the good is being assigned to the user that had the highest valuation for it. The social welfare of an allocation is the sum of the valuations of the agents for that allocation. The Vickrey auction is maximizing the social welfare, since its maximum value is equal to the maximum valuation of a single agent. Such a goal is relevant, since it is usually desirable that goods are given to the users who value them most highly.
- *Implementability*: A second important property is the possibility of each user implementing his own strategy simply, without having to ponder the other players' choices. In other words, each agent needs an elementary computational power to determine his own optimal strategy, called dominant. In the case of the Vickrey auction the dominant strategy is very simple, since it corresponds

to communicating an offer that is equal to the user's valuation. We observe that in less elementary classes of games, such as those illustrated in the preceding sections, the optimal strategies, such as pure Nash equilibria, may be multiple or not even exist, and reaching those strategies might require a set of relatively complex interactions among the agents.

- *Truthfulness*: An additional result of the Vickrey auction is the revelation from each agent of his own evaluation, in other words of the private information that defines the agent and his contribution to the input data of the problem. This property is the prevailing fundamental solution concept in the design of mechanisms, since in its absence the algorithm allocating the goods has to base its decisions on misleading information about the features of the participating agents, and thus on input data that are different from the actual ones.

9.5.2 Vickrey–Clarke–Groves Mechanisms

Some assumptions of our model of the Vickrey auction require further discussion. In particular, we assumed that the valuation of the good by an agent could be expressed in monetary terms. It follows that the utility of an agent is also expressible in monetary terms: the utility of agent i when receiving the good at price p_i is equal to $v_i - p_i$, or 0 if the good is not allocated to the agent. Such a utility model is called quasilinear. But other models have been extensively adopted during the history of economical sciences. For example, each agent could express an order of preference among all possible alternatives or communicate his own preference among each pair of allocations. However the quasilinear model of utilities yields important results in terms of dominant strategies that may not be achieved by other models.

We should also define in a slightly more formal way the notion of mechanism, in order to better appreciate the generality of this concept. The notion of mechanism in the economic sciences is the counterpart of the concept of algorithm in computer science. A mechanism for the allocation of goods has as input the description of the preferences of the agents among all possible allocations of the goods. The output of the mechanism is the selection of one allocation for the goods and of the payment that each agent should contribute. Somewhat more formally, the input to an economic mechanism is composed of a set I of n agents and a set A of possible alternative allocations of the goods. For example, in the case of an auction for a single good, the set of alternative allocations is composed of the n distinct assignments of the good to one of the n agents, the winner of the auction. The preferences of agent i on each alternative in A are represented by a valuation function v_i . The term $v_i(a)$ denotes then the valuation of agent i for alternative a . The function v_i is, however, a private information of agent i . It may be convenient for the agent i to communicate a different and not necessarily truthful valuation function v'_i . The result of the application of the mechanism, that is, the output of the mechanism, consists in the selection of an alternative in A and of a payment p_i for each agent i .

The key concept in mechanism design is that of truthfulness, as we already informally discussed in the case of the Vickrey auction. A mechanism is truthful if, for every agent i , reporting the true valuation function v_i is a dominant strategy when maximizing the utility and so is preferable to reporting any other function v'_i . If a and a' are, respectively, the alternatives selected by the mechanism when agent i declares valuations v_i and v'_i , and p_i, p'_i are the respective payments, then a mechanism is truthful if for every agent i , $v_i(a) - p_i \geq v'_i(a') - p'_i$, independently of the strategy of the other agents. Such mechanisms are also called *direct revelation* mechanisms since they are implementable by an algorithm that receives from the agents the real input data formed by the private valuation of the agents. This is not the only class of mechanisms that can implement a function as the dominant strategy equilibrium of a game. The so-called *Revelation Principle* ensures, however, that all mechanisms that can be implemented through dominant strategies can also be implemented in the form of truthful mechanisms.

The issue is then that of determining a selection mechanism for one of the alternatives and a set of payments that induce a truthful behavior from the agents. The fundamental result in this direction is the Vickrey–Clarke–Groves (VCG) mechanism, that we discuss in the related box for the interested reader.

Vickrey–Clarke–Groves Mechanism

We define the social welfare of an alternative a as the value $\sum_i v_i(a)$, that is the sum of the valuations of the agents for a . A mechanism is called a Vickrey–Clarke–Groves (VCG) mechanism (with Clarke’s rule for the definition of payments) if

- (i) It selects the alternative $a \in A$ maximizing the social welfare $\sum_i v_i(a)$;
- (ii) It defines $p_i = \max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a)$ as the payment of agent i , that is, the maximum possible reduction in social welfare of the other agents due to the existence of agent i , that by reporting valuation v_i has determined the selection of a by the mechanism.

Intuitively, the payment required from every agent is the compensation due to the maximal damage caused by agent i to the other agents because of his existence, when alternative a is selected. Such a mechanism is truthful since the agent maximizes the social welfare by revealing his own true preferences. In a way, the VCG mechanism carries out the task of decoupling the choices of the agents and determines as the dominant strategy the declaration of the true preferences of the agent. We also observe that payments are always positive and the utility of each agent equals $\max_{a \in A} \sum_i v_i(a) - \max_{b \in A} \sum_{j \neq i} v_j(b)$.

As an example of application of VCG with Clarke’s rule, we proceed to prove that the Vickrey auction is a VCG mechanism. The Vickrey auction assigns the good to the agent with the highest valuation and then chooses

(continued)

(continued)

the alternative maximizing the social welfare among the n possible ones. For the agent that obtains the good, $p_i = \max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a)$ is exactly equal to the second price, since the highest valuation of a different alternative is the second price offered by an agent, while the other agents have valuation 0 as they do not receive the good. For an agent i not obtaining the good, $p_i = \max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a)$ has in fact value 0 since both terms equal the maximum valuation of an agent.

Another example of application of VCG is given by a multiple items auction. In a multiple items auction, k identical goods are offered to a set of $n > k$ agents, each of which seeks to obtain a single unit of the k available goods. Each agent has the same fixed valuation for all alternatives in which one of the goods is allocated to him. Each agent submits his bid in a sealed envelope. VCG chooses the alternative that maximizes the social welfare, which in this case means assigning the k units to the k agents with the highest valuations. The payments computed by VCG with Clarke's rule for the k agents that receive the good are in this case set to the $(k + 1)$ th highest valuation of an agent, in other words the highest valuation of an agent that does not obtain the good. Indeed, for an agent i obtaining the good, the first term in $p_i = \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a)$ is exactly the sum of the $k + 1$ highest valuations minus the valuation of agent i , while the second term equals the sum of the k highest valuations minus the valuation of agent i . It can be promptly checked that the payment asked to the agents that do not obtain the good equals 0. The multiple items auction can be further generalized, for example, in the setting of the so-called combinatorial auctions, that, however, we will not consider in this discussion.

9.5.3 Computational Aspects of Mechanism Design

The VCG mechanism and the definition of truthful, direct revelation mechanisms are the fundamental concepts at the basis of the area of mechanism design and its algorithmic aspects. In the following, we consider a fundamental aspect that concerns the computational issues related to social welfare maximization.

VCG requires the maximization of an objective function, the social welfare, in the domain of all possible allocations. The maximization of such an objective function can be computationally difficult for many problems of high relevance (see also Chap. 3). It is customary in the field of optimization algorithms to circumvent the computational complexity of exactly optimizing a function by the use of an approximation algorithm, that is, an algorithm that always allows one to obtain, on any instance, a solution close to the optimal one, while incurring a polynomial

computational cost. In a strategic setting it is also required, however, that such functions be implementable through some dominant strategy equilibrium.

A fundamental question is thus the characterization of the social choice functions that are implementable. Are there simple properties that a function should have so that it is implementable by dominant strategies of the agents?

A relevant case is that in which the valuation of an agent is some numerical value in a subset of the alternatives $W_i \subseteq A$, called winning for agent i , while it is zero on all other (losing) alternatives. More precisely, the agent is characterized by a valuation t for all alternatives in W_i and 0 for all alternatives outside W_i . In this particular case it is possible to completely characterize the set of truthful mechanisms by relying on the monotonicity properties of the implementable functions. A single-parameter allocation function is called *monotone* in v_i if it is possible to define a single critical value corresponding to the minimum valuation that allows the agent to be in a winning alternative. For example, in the case of the Vickrey auction the critical value equals the largest of the valuations of the losing agents.

The definition of critical value plays a fundamental role in the definition of payments, since the agents' payments can be fixed exactly at the critical value. The importance of monotone functions is in the fact that a monotone social choice function and the respective critical values allow one to implement a truthful mechanism in which the winning agents pay exactly the critical value. Such a characterization is of great importance, since it allows the implementation of truthful mechanisms for the optimization of social welfare functions that are computationally hard and that can be approximated by monotone functions that are simple to compute.

Consider, for example, a multi-unit auction problem, in which m units of the same good are available to n agents. Each agent j desires some number q_j of units, known to the mechanism, and has a private valuation function v_j for obtaining such a set of units. This is a clear case in which agents are completely described by a single private value, their valuation. The problem of maximizing social welfare in the allocation of the m available units to the n agents corresponds to the well-known knapsack problem (see Sects. 2.3.2 and 5.3), for which an application of VCG would require a computational cost that is not polynomial in the size of the input. For the interested reader, the box "Monotone Mechanisms" reports an alternative approach to the same problem.

Monotone Mechanisms

We consider an alternative approach to the multi-unit auction problem, based on the adoption of monotone mechanisms that yield a good approximation to the optimal solution of the problem in question. In particular, consider these two algorithms:

1. VAL: Sort the agents non-increasingly by their valuation v_j ;

(continued)

(continued)

2. DENS: Sort the agents non-increasingly by their density (valuation per unit) v_j/q_j .

Both algorithms define a ranking among the agents and allocate the elements according to this ranking, until some agent requires more units than those that remain available. Let's look at the execution of both algorithms when there are 3 players ($n = 3$), 4 units of the good ($m = 4$), and the agents' data are the following:

$$v_1 = 5, q_1 = 2; v_2 = 3, q_2 = 1; v_3 = 4, q_3 = 2.$$

The algorithm that sorts by valuation will allocate 2 units to agent 1 and 2 units to agent 3, with an overall social welfare of 9. The algorithm that sorts by density will allocate 1 unit to agent 2 and 2 units to agent 1 for an overall social welfare of 8. Observe that in this last case one unit of the good is not allocated.

Firstly, we observe that both algorithms are monotone. Indeed, if an agent with a given valuation is selected, he is also selected when his valuation is increased.

Let us compute the payments for the agents in both cases. Such payment for each agent is equal to the minimum valuation that would allow the agent to be selected. In the case of ranking by valuation, the payments required from the two winning agents equal the valuation of the losing agent.

In the case of ranking by density, consider the first agent in the ranking that is not winning. Let this agent be j . The payment of winning agent i will be equal to $q_i(v_j/q_j)$. Observe that in case all agents are winning each payment equals 0, since the presence of agent i is not decreasing the social welfare of the remaining agents.

Each algorithm individually is not producing a good approximation of the optimum. Indeed, the ranking by valuation could accept an agent requiring all m units with a valuation of 2 and at the same time not accept m agents each requiring one unit with valuation 1. Analogously, the ranking by density could accept an agent requiring 1 unit with valuation 2 and not accept an agent requiring m units with valuation m . Notice, however, that if we define a third algorithm called MAX that outputs the best solution among those given by VAL and DENS, then we obtain a solution close to optimum. In our example, MAX returns the solution constructed by algorithm VAL. It can be easily shown that, in fact, algorithm MAX always constructs a solution with value at least half the optimum solution.

(continued)

(continued)

Unfortunately, the algorithm obtained by the combination of two monotone algorithms is not necessarily monotone. For this to hold, it is required that the two algorithms satisfy a condition stronger than monotonicity, called *bitonicity*. The combination of two bitonic algorithms is indeed monotone. An algorithm is bitonic if it is monotone and whenever a losing agent increases his valuation then one of the two following conditions is satisfied:

1. The agent becomes winning; or
2. The agent remains losing but the value of the constructed solution does not improve.

In the case in question, it is immediate to verify that any algorithm that defines the winning agents by traversing a ranking in which the priority of each agent can only increase if his valuation increases, is bitonic. Indeed, if the valuation has not increased enough to make the agent winning, then the constructed solution does not change.

9.6 Price-Setting Mechanisms and Competitive Equilibria

Mechanisms for defining prices of goods have played a central role in economic theory during the last two centuries. Price-setting mechanisms have been given the task of introducing fundamental notions such as efficiency and equilibrium in markets, the availability of goods and services, and the operation of the economy in stable conditions. The role of initiator of the mathematical theory of markets is attributed to Léon Walras, who in 1874 was the first to define the notion of *competitive equilibrium*: an equilibrium in which each agent obtains the good from which he extracts the highest utility and in which all goods left unsold have zero price. Mechanisms for price setting form a further fascinating field of application of algorithmic methods and provide important computational questions.

Consider the case in which we have a set I of n agents and a set J of m distinct goods. Each agent has a valuation function v_i for each set S formed by some of the m goods. The goal is to allocate a set of goods to each agent. A price-definition scheme defines a price p_j for each good $j \in J$. The *demand of agent i* is defined as the set S of goods preferred by agent i , the one that maximizes the utility $v_i(S) - \sum_{j \in S} p_j$ of the agent.

A *Walras equilibrium* is defined as a set of prices in which each agent receives his demand and all goods not allocated have price 0.

A first important result, also called *first welfare theorem*, proves that a Walras equilibrium determines an allocation maximizing the social welfare of the agents. In the following, we consider an example scenario.

Consider two agents, Alice and Bob, and two elements $\{a, b\}$. Alice has valuation 2€ on every nonempty set of elements, while Bob has valuation 3€ on the whole set $\{a, b\}$ and valuation 0€ for each single item. The solution that maximizes social welfare assigns the set $\{a, b\}$ to Bob. Thus, to obtain a Walras equilibrium, Alice should prefer the empty allocation to each of the individual items. For this to be true, the price of each good should be at least 2€ . But then the price of the set $\{a, b\}$ is 4€ and so Bob, too, will prefer the empty set as an allocation. There is therefore no Walras equilibrium in this case.

An important question in economic theory is the characterization of markets that admit a Walras equilibrium. The answer to such a question can be reduced, surprisingly, to the solution of some important algorithmic questions. It is in fact possible to relate the existence of a Walras equilibrium to the existence of an integer solution to a linear optimization problem in fractional variables. We stop here and do not embark on a field that would require a long discussion. The interested reader will find in Sect. 9.7 some references to the main monographs surveying the area.

9.7 Bibliographic Notes

The book that marked the birth of game theory is the classic text of von Neumann and Morgenstern [108]; it is remarkable that the book is largely due to John von Neumann, one of the fathers of the digital computer era, as a well as world-class mathematician.

Modern treatments of the theory of games and its applications to economics are given, for example, by Binmore [9], Osborne and Rubinstein [88] and Mas-Colell, Whinston and Green [76]. The encyclopedic work by Aumann and Hart [6] is an updated state of the art of the theory, and contains many further references to the scientific literature. Computational aspects and applications of game theory to computer science are discussed in a recent textbook by Nisan et al. [86], to which we point the interested reader for further investigation of many of the topics touched in this chapter, including algorithmic mechanism design.