CHAPTER 1

Introduction and Preview

This book synthesizes what one may refer to as *contest theory*, understood in a broad sense to encompass scientific methods and theories for the better understanding and informed design of contests. Its goal is to provide a contest designer with a set of theoretical results and methods that can be used for the design of contests. An ambitious aspiration is to provide a toolkit for a contest designer of a similar kind to what control theory offers to engineers for the design of control systems. This is, undoubtedly, a challenging task, primarily because of the complexity of user behavior and incentives that play a key role in most of the systems of concern. This book covers a wide range of models developed in different areas of science including computer science, economics, and statistics.

Generally speaking, we refer to *contests* as situations in which individuals invest efforts toward winning one or more prizes, those investments of efforts are costly and irreversible, and prizes are allocated based on the relative values of efforts. A prize is understood in a broad sense to refer to a notion of value that is general enough to include not only monetary prizes but also social reputation and gratitude. How to allocate a prize purse to competitors in a contest was studied as early as 1902 by Galton, who reasoned about the question, "what is the most suitable proportion between the values of first and second prizes?" assuming a statistical model according to which individual production outputs are independent and identically distributed random variables with a given distribution. An economist's approach is to assume that contestants are rational players who strategically invest efforts with a selfish goal of maximizing their individual payoffs, which combine in some way the value of winning a prize and the cost of production. The study of a contest as a game using the framework of *game theory* allows us to reason about properties that arise in a strategic equilibrium. The design of a contest needs to ensure that proper incentives are put in place to achieve a desired objective. Commonly studied objectives include the total effort invested by contestants, the maximum individual effort over all contestants, and the social welfare defined as the value of the prizes to those who win them.

A canonical auction theory model of a contest is that of an *all-pay auction*. In an all-pay auction, each player incurs a cost equal to his or her invested bid, and a prize is

allocated according to an allocation mechanism that specifies the winning probabilities for every given vector of bids. This simple model of a contest offers plenty of interesting equilibrium properties and allows one to study how the choice of a prize allocation mechanism affects a quantity of interest in a strategic equilibrium. A natural contest design is to offer several placement prizes. Contest architectures can also have more sophisticated structures. For example, a contest owner may award prizes only if the individual performance outputs meet a certain level of quality. Other examples include contests that are structured as a composition of simple contests; for example, a system of *simultaneous contests* where each player strategically decides in which contests to invest efforts; a system of *sequential contests* where the competition proceeds through multiple rounds, each consisting of one or more simple contests; or a *single-elimination tournament* that is yet another instance of a composition of simple contests.

In this book, we cover a general theory of contests that provides models for a wide range of situations. We use the term *online contests* to refer to a broad range of systems that arise in the context of Internet online services whose design is based on some elements of contests. Many Internet services can be modeled in this way.

A general question that serves as a motivation for much of this book is as follows:

How should a contest be designed so as to achieve a given objective?

In designing a contest, one may need to address a broad range of design choices. In this book, we focus on theoretical foundations to provide a better understanding of the implications of various contest design choices. Specific questions of interest concern properties of production outputs and are naturally studied in the framework of game theory, drawing from the areas of auction theory, rent-seeking, algorithmic mechanism design, and social choice. Other questions may have little to do with the strategic behavior of players and require methodologies from other scientific disciplines. For example, the estimation of skills of players based on observed outcomes in contests requires methods from statistics. Another example of a contest design problem that is not necessarily studied as a strategic game is that of seeding of a single-elimination tournament, which under certain conditions boils down to a sorting problem with noisy observations. Figure 1.1 indicates some of the well-established areas from which one may draw methodologies to reason about contests.

This book puts together in one place a wide range of theoretical results that have been developed over many years by different scientific communities, including various branches of economic theory such as political economy; theory of games, and in particular auction theory; computer science; and statistics. The original studies were motivated by applications as diverse as animal conflicts, auctions for selling various kinds of goods, political lobbying, research and development races, production in firms, and Internet online services, to name a few.

Three particular fields are of special interest in the theory of contests: economics, computer science, and statistics. Here we highlight some of the main contributions in these areas.

Economic theory established much of the theory of contests, particularly, in the context of game theory, auction theory, and studies of rent-seeking activities in public choice theory. It introduced the area of *mechanism design*, which is a field in game theory that studies solution concepts for games with private information where the goal

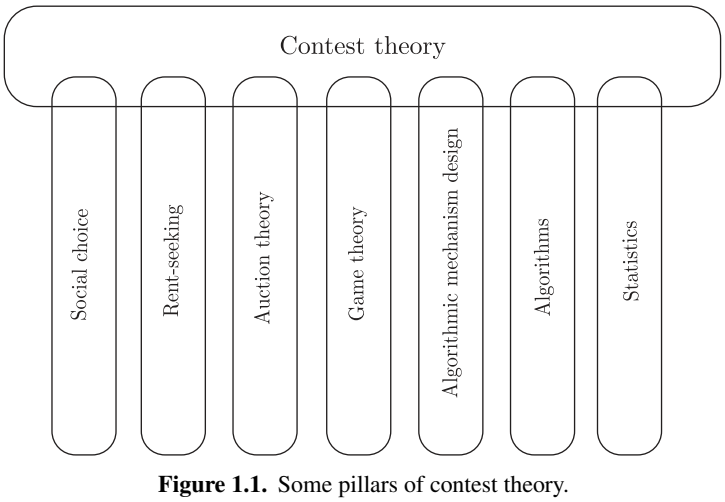

Figure 1.1. Some pillars of contest theory.

is to design a mechanism to achieve a desired objective. The foundations of mechanism design theory were laid down in a series of works starting in the 1960s with major contributions by Leonid Hurwicz, Eric Maskin, and Rodger Myerson, who together won the 2007 Nobel Memorial Prize in Economic Sciences for this line of work. The optimum auction design with respect to social welfare was established by Vickrey (1961) and with respect to revenue by Myerson (1981). A complete characterization of the equilibrium of the all-pay auction with complete information was established by Baye et al. (1996). A model of rent-seeking, commonly referred to as *the Tullock contest*, was formulated by Tullock (1980). This model formally corresponds to an all-pay auction with a specific form of a prize allocation mechanism.

Computer science introduced the use of analytical tools from theoretical computer science such as *worst-case complexity* and *approximation ratio* to study mechanism design questions. The concept of a worst-case equilibria was introduced by Koutsoupias and Papadimitriou (1999). This has been followed by many studies of the social efficiency of various games using the concept of *the price of anarchy*, defined as the ratio of the optimum social welfare and the social welfare in a worst-case equilibrium. The term *algorithmic mechanism design* was coined by Nisan and Ronen (1999, 2001) to refer to a theoretical computer science approach to mechanism design questions using concepts such as communication complexity, approximation algorithms, and worst-case analysis.

Statistics has been used to reason about contests as early as in the aforementioned study by Galton (1902). Statistical methods have played a major role in the study of tournament plans and, in particular, seeding procedures, with theoretical foundations laid down in the works of Chung and Hwang (1978), Hwang (1982), Israel (1982), and Horen and Riezman (1985) and continued developments being made since then. Statistical methods for ranking of alternatives based on pair comparisons were developed by Thurstone (1927) and Zermelo (1929) and subsequently studied by Bradley and Terry (1952, 1954) and many others. This work laid foundations for the design of rating systems that are in popular use today in sport competitions, crowdsourcing services, and online computer games.

The remainder of this chapter is structured as follows. Section 1.1 provides an overview of different situations in which contests arise and describes some concrete examples of contests that demonstrate different contest architectures and structures of prizes. Section 1.2 provides a game theory primer that defines various types of games, equilibrium concepts, and measures of efficiency of equilibria studied in this book. An informed reader may quickly skim through this section. Section 1.3 provides an overview of the topics covered in this book and highlights some of the interesting questions and results. Section 1.4 concludes the chapter with a discussion of bibliographical references.

1.1 A Survey of Examples

In this section we discuss various examples of contests starting from the more traditional ones including sport contests, rent-seeking, patent races, labor markets, and scientific research projects, and then go on to discuss online contests in the context of crowdsourcing services and other types of competitions.

1.1.1 Sport Contests

Sport competitions have a long and rich history and a wide range of contest architectures, structures of prizes, and methods for rating and ranking of players and teams of players. The ancient Olympic Games were held at Olympia as early as 776 BC and were abolished in 393 AD. The games organized at Olympia led to the development of the Panhellenic Games, which were held at Olympia and three other sites. At the Panhellenic Games there was only one winner, awarded a wreath or crown of leaves. The modern Olympic Games were inaugurated in 1896, with the first, second, and third placed athletes awarded with, respectively, gold, silver, and bronze medals. Finley and Pleket (2005) and Swaddling (2011) are good sources for these historical facts. Evidence of the use of prizes in ancient Greece can also be found in Homer's *Iliad* II. 23, 249–897, in the funeral games instituted by Achilles in honor of Patroclus: these games consisted of the chariot race, boxing, wrestling, foot race, single combat, discus, shooting with arrows, darting, and javelin. According to Papakonstantinou (2002), the winners were awarded various kinds of valuable objects: "the circulation of these valuable objects was an integral part of aristocratic gift-exchange and that therefore such prizes reaffirmed social hierarchies and consolidated networks of power relationships of the Homeric elites."

The use of prizes in various sports has prevailed to date, with a wide range in the number and values of prizes being awarded. For example, in the game of football in Europe (soccer), only one club wins the title of a national season champion, but this is not the only prize. The final outcome is a league table ranking where the position of a team in the ranking is associated with various kinds of prizes. A number of top-ranked teams often qualify to participate in international competitions such as the UEFA Championship League and UEFA Europa League. On the other hand, a number of teams at the bottom of the league table are relegated to a lower league. The league table ranking typically positively correlates with the monetary prizes awarded to the

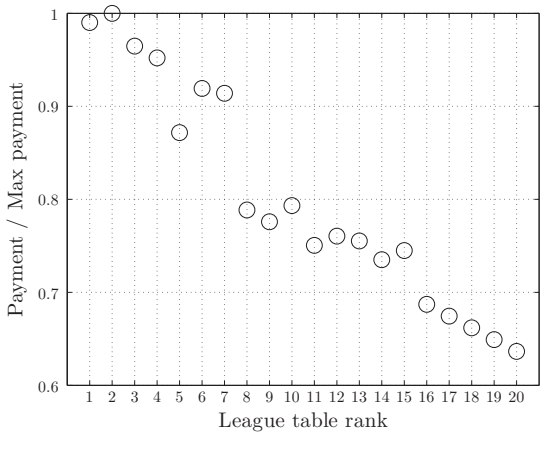

Figure 1.2. Broadcast payments made to the clubs in the premiership league season 2013/2014 versus the league table rank: the maximum payment is of value $97,544,336 (Barclays Premier League (2014)).

teams. For example, in the Barclays Premier League, according to a contract, the UK broadcast revenue is distributed to the 20 clubs of the league such that 50% of the revenue is split equally among the clubs, 25% of the revenue is split based on merit payments, i.e., depending on where a club finishes in the final league table, and the final 25% is paid in facility fees each time a club's match is broadcast in the United Kingdom. All international broadcast revenue and central commercial revenues are split equally among the 20 clubs. See Figure 1.2 for the values of payments to clubs in the premiership league season 2013/2014 depending on their placement in the final league table. The prizes generally play an important role as an incentive mechanism in sport contests. An interesting historical example of the design of prizes in the context of the game of soccer is the introduction of the system known as the *three points for a win*, where in a match, the winner is awarded three points and the loser zero points, and in the case of a draw outcome each team is awarded one point. This system replaced the original two points for a win system in England in 1981 and was adopted by most European national leagues by the mid-1990s. The new system was introduced under a premise that it would encourage more attacking play because teams would not settle for a draw if they have the prospect of gaining two extra points.

Sport contests have served as a motivation for the development of various contest design methodologies. A prominent example is the design of tournament plans and scheduling of contests. A formal study of the question of how to design a tournament plan using statistics was attempted already in 1883 by Charles Dodgson (better known by his pen name Lewis Carroll) in his article "Lawn Tennis Tournaments: The True Method of Assigning Prizes with a Proof of the Fallacy of the Present Method." A testimony to the richness of the research in this area is an extensive annotated biography of scheduling in sports by Kendall et al. (2010). The design of tournament plans is discussed in Chapter 8. Economic models have also been applied to various sport contests, some of which are covered in Chapters 2, 3, 4, and 7; see the survey by Szymanski (2003). The design of systems ranking and rating skills of players has also been greatly influenced by sport competitions. For example, the methods for rating skills of players

in the game of chess have been in use at least since the late 1930s. The currently most popular chess rating system, namely the Elo rating system (Elo 1978), was first adopted by the United States Chess Federation (UCSF) in 1960 and then by the World Chess Federation (FIDE) in 1970. This rating system has also been applied in several other sports, including American college football, basketball, football, and tennis. Chapters 9 and 10 provide an exposition of the principles of rating and ranking methods.

1.1.2 Rent-Seeking

The term "rent" refers to gaining control of a resource. An example of a rent-seeking activity in our modern economy is spending money on political lobbying for government benefits or subsidies in order to gain a share of wealth that has already been created, or to impose regulations on competitors in order to increase market share. Rent-seeking activities may result in social efficiency loss, reduced wealth creation, government revenue losses, and income inequality. In economics, rent-seeking activities were first studied by Tullock (1967), while the term *rent-seeking* was coined later in the work of Krueger (1974). The model of rent-seeking introduced by Tullock (1980) has elicited general interest as a model of a contest. The key property captured by this model is that a contestant increases his or her chance of winning the prize by increasing his or her effort investment, but it is not guaranteed that the one who invests the most wins the prize. Chapter 4 provides a comprehensive coverage of this and more general families of such models.

1.1.3 Patent Races

Firms invest in research and development (R&D) under both technological and market uncertainties. There is uncertainty about the relationship between a firm's investment in R&D and the time at which the innovation may be introduced by the firm. The market uncertainty occurs because no firm can be sure when any of its rivals' R&D efforts will be successful. Patent races are common in pharmaceutical, software, and other industries and are often referred to as *R&D races* or *research contests*. A classic model of a patent race, which we discuss in Chapter 4, was pioneered by Loury (1979) and subsequently studied by Lee and Wilde (1980) and Dasgupta and Stiglitz (1980). The practice of using patents as a mechanism to promote innovation is under much debate among economists, e.g., Stiglitz (2008). One of the main concerns is that giving exclusive rights to corporations enforces monopoly and social inequality, especially between the developed and the developing world. Some economists, e.g., Stiglitz (2006), have advocated the use of alternative mechanisms such as innovation inducement prizes.

1.1.4 Innovation Inducement Prizes

As a way to promote innovation, governments and private parties award prizes for solutions to predefined scientific or technological problems or for demonstrations of the feasibility of specified unprecedented accomplishments. We refer to these as *innovation inducement prizes*. For concreteness, we next discuss a selection of such prizes from a large number of representative examples.

The Netflix prize challenge is a contest that was run from October 2, 2006, to September 21, 2009; it sought to substantially improve the accuracy of predictions about how much someone is going to enjoy a movie based on his or her movie preferences. This is an example of a contest awarding the main prize subject to a minimum required level of quality. To qualify for the grand prize in the amount of $1,000,000, the accuracy of a submitted prediction was required to be at least 10% better than the accuracy of a benchmark prediction method.

The DARPA Network Challenge was launched in 2009 with the broad goal of exploring the role of the Internet and social networking in real-time communications, wide-area collaborations, and practical actions required to solve broad-scope, time-critical problems. A cash prize in the amount of $40,000 was offered to the team that could first identify the locations of 10 red balloons moored at different locations across the United States.

Kaggle is a platform for hosting challenges on data science problems, which offer a wide range of monetary prizes. The users are also awarded reputation points based on the outcomes in contests; these points are discounted over time. The number of points earned by a team member for participation in a contest is based on several factors including the placement of the team, the number of competing teams, and the number of team members.

X Prize is a non-profit organization founded in 1995 that designs and manages public competitions intended to bring about major breakthroughs for the benefit of humanity. The X Prize competitions are typically big projects with prizes in the order of millions of dollars.

Challenge.gov is a platform established in 2010 for hosting challenge and prize competitions, all of which are run by various agencies across the U.S. federal government. These include technical, scientific, ideation, and creative competitions where the government seeks innovative solutions from the public, with the goal of bringing the best ideas and talent together to solve mission-centric problems.

1.1.5 Labor Markets

Contest theory provides models for labor compensations in firms that reward employees using prizes such as bonuses, stock shares, and job title promotions, often based on an individual's performance relative to that of his or her co-workers. Typically, there is a limited prize budget, such as a percentage of a firm's annual revenue or a fixed number of stock shares, as well as a limited number of available positions at each level of an organization hierarchy. The question is how to design a promotion system to induce employees to work the hardest, or perhaps to maximize the chance of promoting the best worker. A first formal study of these questions using economic theory was given by Lazear and Rosen (1981) in which they compared compensation schemes that pay according to an individual's ordinal rank of performance to those that pay based on an individual's absolute production output. In Chapters 2 through 8, we cover a wide range of models of contests that award prizes

based on the relative performance of contestants, which are of interest in the context of labor markets.

1.1.6 Scientific Projects

Academic research is mostly realized through scientific projects funded by various non-profit and profit organizations. Scientists form teams to work on projects that result in publications and other deliverables. An illustration of the scale of scientific collaborations is presented in Figure 1.3 showing co-authorship of publications in computer science. Contributors to a project receive shares of scientific credit, whose values depend on the resulting impact and perhaps on the values of individual contributions and other factors. The impact is measured through performance indicators such as the number of citations and other indicators of societal impact. How should scientific credit be allocated to incentivize production that has high social value? Such questions were considered in studies of the allocation of scientific credit and division of cognitive labor, such as Merton (1968, 1973) and Kitcher (1990, 1993). Chapter 6 discusses the social efficiency of a model of a general production system in which individuals invest efforts in one or more available projects under different assumptions on how the efforts invested relate to the value of the utility produced and different utility sharing mechanisms.

1.1.7 Crowdsourcing Services

Crowdsourcing is the process of obtaining needed services, ideas, or content by soliciting contributions from a large group of people, especially from an online community, rather than from traditional employees or supplies. The first definition of the term "crowdsourcing" was posted in a blog by Howe (2006a) on June 2, 2006, which

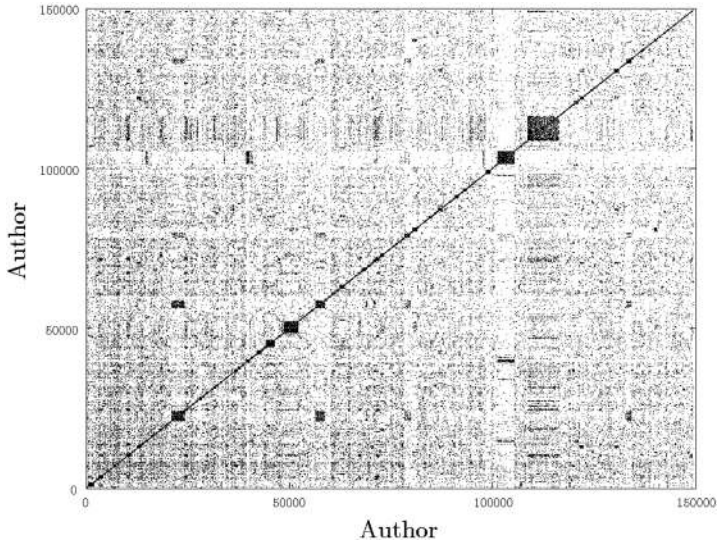

Figure 1.3. Co-authorship of journal articles published in 2010 according to DBLP: each point (x, y) represents co-authorship of a journal article by authors x and y.

was soon followed by a popular book by Howe (2006b) on this subject. The term *competition-based* crowdsourcing refers to soliciting solutions to various tasks through organizing contests. The sponsor (a private person or a company) identifies a specific problem, offers a cash prize, and broadcasts an invitation to submit solutions. Contests work well in situations when it is not clear what combination of skills or technical approach will lead to the best solution for a problem at hand, or the sponsor lacks information about which specific workers are most skillful for the given problem. The term *paid-labor* crowdsourcing refers to online labor marketplaces for matching workers and jobs where the compensation for labor is typically based an a priori agreed contract between an employer and a worker or a team of workers. An important feature of a crowdsourcing platform is that it typically hosts a number of open contests at any given time, providing workers with alternative options for effort investments.

Competition-Based Crowdsourcing Services

We discuss here two examples of competition-based crowdsourcing platforms, which help illustrate the type of tasks, the structure and value of prizes, and participation in contests.

TopCoder is a competition-based software development crowdsourcing platform using monetary and reputation prizes. Contests are held for the design, development, specification, and architecture of software development tasks. Table 1.1 provides some summary statistics. Such contests typically offer two prizes, splitting a prize purse between the first and the second place prize in the ratio 2:1. See Figure 1.4 for a

Table 1.1. Statistics of participation in some crowdsourcing contests[a]

Category	Tasks	Users	Mean User/Task	Median User/Task
TopCoder				
Design	2,041	535	2.35	2
Development	2,147	1,450	3.47	2
Specification	199	71	2.33	2
Architecture	378	94	1.83	2
Taskcn				
Website	13,673	1,946	11.70	7
Design	11,306	3,388	17.51	13
Coding	896	4,758	7.15	4
Writing	3,398	6,328	50.75	15
Multimedia	155	2,305	15.98	9
Other	1,824	3,768	17.14	11

[a] (TopCoder) data covering approximately a 10-year period from early 2003 until early 2013, (Taskcn) data covering approximately a 7-year period from mid-2006 until early 2013.

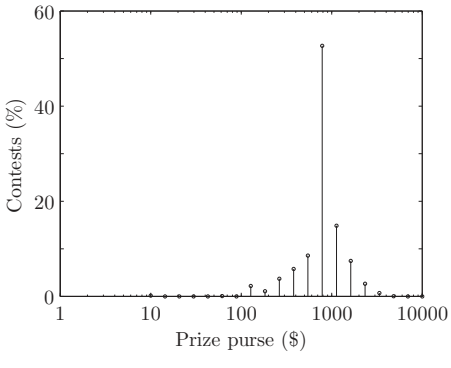

Figure 1.4. TopCoder histogram of prize values.

histogram of the values of prizes. The algorithmic contests award rating scores using TopCoder's own rating system, which is described in Chapter 9. An annual tournament, the TopCoder Open Tournament, is held for algorithmic competitions, which is an illustrative example of a tournament design with multiple elimination stages. See Figure 1.6 for a graphical representation of this tournament design.

Taskcn is a platform for soliciting solutions to various tasks such as website design, design, coding, writing, multimedia, and other types of tasks. Tasks typically attract in the order of ten submissions and some, such as writing a slogan, attract a considerably larger number of submissions. Table 1.1 shows statistics of participation in contests, and Figure 1.5 shows participation of users in contests over month-long periods in a year.

The key feature of competition-based crowdsourcing is the all-pay nature of contests: online workers invest irreversible efforts in producing solutions to tasks, and only those whose solutions are selected as the winning solutions for a given task are awarded a prize. Chapters 2, 3, and 4 discuss models of all-pay contests under different prize allocation mechanisms that specify the relation between the values of effort investments and the winning probabilities. Another key feature of competition-based crowdsourcing is that, typically, at any given time there are a multitude of open contests in which a user may participate. Chapter 5 covers models of simultaneous contests that are of interest in such scenarios.

Paid-Labor Crowdsourcing Services

We now discuss three representative examples of paid-labor crowdsourcing platforms that are described next.

Amazon Mechanical Turk is a crowdsourcing Internet marketplace where requesters post tasks known as HITs (Human Intelligence Tasks). HITs are typically simple tasks such as choosing a preferred photo for a storefront, writing a product description, and categorizing restaurants. Workers (sometimes referred to as Turkers) select and complete tasks for a given monetary payment.