

Exploring the basic game theory models of contests found in online services.

BY MILAN VOJNOVIĆ

Contest Theory


A VARIETY OF Internet online services are designed based on contests. A canonical example is crowdsourcing services, which solicit solutions to tasks by open calls to online communities. Here the tasks can be of different categories, such as art design, software development, data-science problems, and various challenges such as planetary-scale locating of objects.^{12,28} These services operate under certain contest rules that include specifying a prize allocation mechanism, for example, awarding only a first-place prize or several position prizes. The prizes can be monetary, or in-kind rewards such as in terms of attention, status, or computing resources, for example, CPU, bandwidth, and storage. We refer to a contest as any situation in which agents invest irreversible and costly efforts toward winning a prize, which is allocated based on relative performance. We use the term “contest theory” in a broad sense to refer to a set of theories developed for the better understanding and informed design of contests.

A central question in contest theory is: *How to allocate prizes to maximize a desired objective?* The objective may be to maximize the utility of production to the agent who solicits solutions to a task, or to the whole society. The question of how to allocate prizes was studied as early as 1902 by Galton.¹⁵ A study of how to allocate prizes necessitates to consider the incentives of contestants, who act strategically in investing costly production efforts.^{1,11,44} Game theory models of contests have been studied in auction theory, economic theory, operations research, as well as theoretical biology; for example, Bishop and Smith.³ The use of compensation schemes based on an individual's ordinal rank rather than absolute performance in firms have been studied by economists; for example, Lazear and Rosen.²⁷ Game theory and pertinent computational questions have been studied by computer scientists.^{31,33,36} Several new contributions have been made on optimal allocation of prizes in crowdsourcing contests, equilibrium outcomes in games that model simultaneous contests, and the worst-case efficiency of production in equilibrium outcomes of various games that model contests.

The skill-rating methods that use observations of relative performance comparisons as input, such as ranking outcomes in contests, have been studied extensively in the past. They are now widely used in various applications, such as sport competitions, online gaming, and online labor platforms.

» key insights

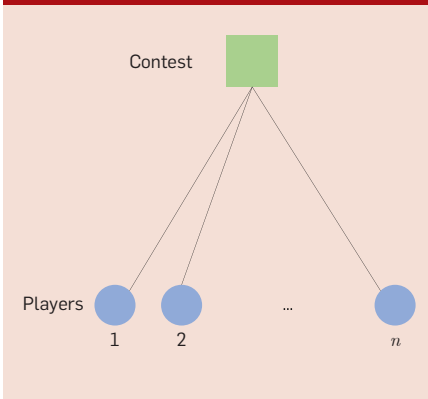
- The operation of various online platforms relies on incentive mechanisms for eliciting user contributions, which take the form of a contest.
- Contest theory refers to a set of theories for the better understanding and informed design of contests.
- The theory provides insights into what user behavior may arise in equilibrium, guidelines of how to allocate prizes, and algorithms for estimating skills of individuals based on observed contest outcomes.



The design of skill-rating methods is based on statistical models of ranking outcomes developed from 1920s onward. More recent developments include skill-rating methods that allow for contests among two or more teams of players, which are common in online gaming and online labour platforms. New results have been recently developed in the area of statistical inference for statistical models of ranking data, including new characterizations of the accuracy of various skill parameter estimators and new iterative methods for skill parameter estimation.

In this article, we survey some main results of contest theory. Specifically, we discuss basic game theory models of contests that are found in online services. We explain the conditions under which to optimally allocate prizes to maximize a given objective, such as the total effort or the maximum individual effort, in a strategic equilibrium. We will focus on games in which players make simultaneous effort investments; the games that involve some aspect of sequential play are only briefly discussed.

Figure 1. Single contest.



We consider both games that model a single contest (see Figure 1) and games that model a system of two or more simultaneous contests (Figure 2). Simultaneous contests are common in the context of online crowdsourcing platforms. We explain basic principles of popular skill rating systems and point out some new results in this area. We conclude with an outlook on future research directions.

This article complements existing surveys on the game-theoretic aspects in contest theory, for example, Corchon,⁹ Konrad,²⁵ and Nitzan.³² We provide an overview of some of the topics covered in the book by Vojnovic,⁴² where the reader may find a more extensive coverage of references.

Strategic Game Models of Contests

The standard game theory framework for studying contests is based on the assumption that agents are rational and strategic players who invest effort with a selfish goal to maximize their individual payoffs. The payoff of a player combines the utility of winning a prize and the cost of production. Specifically, we consider a normal-form game that models a contest, defined by:

- ▶ *Set of two or more players:* $N = \{1, 2, \dots, n\}$;

- ▶ *Payoff functions:* for any given vector of efforts $\mathbf{b} = (b_1, b_2, \dots, b_n)$, the payoff of player i is given by

$$s_i(\mathbf{b}) = v_i x_i(\mathbf{b}) - c(b_i)$$

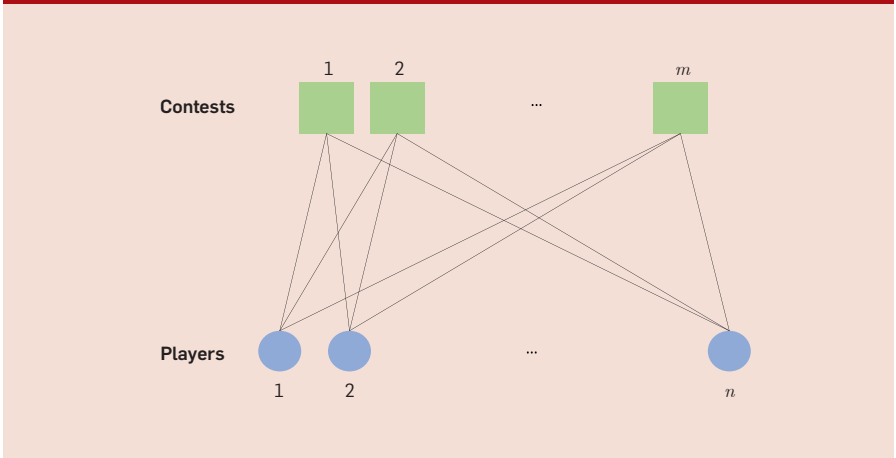
where

- ▶ v_1, v_2, \dots, v_n are positive-valued *skill parameters*,

- ▶ $x(\mathbf{b}) := (x_1(\mathbf{b}), x_2(\mathbf{b}), \dots, x_n(\mathbf{b}))$ is *prize allocation*, and

- ▶ $c(x)$ is a *production cost function*.

Figure 2. A system of simultaneous contests: Edges indicate effort investment opportunities.



The skill parameters reflect the abilities of players: the larger the value of a player’s skill parameter, the more proficient is the player. If the production cost is according to a linear function, we can normalize the payoff functions such that a player’s skill parameter can be interpreted as the reciprocal of his or her production cost per unit effort. The game as defined here allows us to study equilibrium outcomes under different prize allocation mechanisms, such as assigning fixed shares of a prize budget in decreasing order of invested efforts, or splitting a prize budget among players in proportion to their effort investments. The prize allocation can be interpreted as the winning probabilities for an indivisible prize item, or as the shares of an infinitely divisible prize. The game allows us to study equilibrium outcomes for different types of production costs. For example, it is common to consider *linear production costs*, we refer to as *constant marginal production costs*, under which the production cost per unit effort is constant; in particular, we refer to *unit marginal production cost* when the production cost per unit effort is of unit value. We may also consider production costs with either *decreasing* or *increasing marginal costs*. It is noteworthy that the game as defined here formally corresponds to an auction, where efforts, skills, and production costs are in correspondence with bids, valuations, and payments, respectively


We say a game is *with complete information* if the players have perfect information about each other’s skill parameters. A game with complete information can be used as a model of a contest when the players are informed about who is going to participate in the contest and about the skills of the participants. For example, a situation like this can be found in competition-based software development platforms such as TopCoder, where a contest takes place after a registration phase, which reveals identities of participants. A game is said to be *with incomplete information* if the value of each player’s skill parameter is his or her private information. In a game with incomplete information, skill parameters are assumed to be random variables according to a prior distribution, which is a common knowledge.

A game with incomplete information allows us to model uncertainty about skills of competitors in a contest; in the context of online services, such an uncertainty may arise because it may not be a priori known who is going to participate in a contest.


The strategic effort investment by a player can be according to a pure strategy, specifying a value of the effort investment, or according to a mixed strategy, specifying a probability distribution over pure strategies. An investment of efforts by players is a *pure-strategy Nash equilibrium* if no player can increase his or her payoff by a unilateral deviation. Similarly, a set of mixed strategies is a *mixed-strategy Nash equilibrium* if no player can increase his or her expected payoff by a unilateral deviation. A *Bayes-Nash equilibrium* is a mapping of an individual's skill to a value of effort such that no player can increase his or her expected payoff by a unilateral deviation.

The utility of production is typically studied with respect to the following two metrics: the *total effort* and the *maximum individual effort*. The total effort has been studied extensively because it corresponds to the revenue accrued in an all-pay auction, and the total outlay accrued in a rent-seeking contest.^{26,40} The maximum individual effort has been studied motivated by applications in contests, such as in crowdsourcing services, where a contest owner makes use only of the best submitted solution. The utility of production has also been studied from a societal perspective, defined by a *social welfare* function, which is commonly defined as the sum of payoffs of all the parties involved (players and the contest owner). For example, when players incur unit marginal production costs and the payoff to the contest owner is the total effort invested by the players, social welfare corresponds to the total valuation of prizes by those who win them. Social welfare in an equilibrium can be smaller than optimal value; in some instances, optimum social welfare is achieved only if a given prize budget is fully assigned to highest-skill players, while in equilibrium a lower-skill player can have a strictly positive winning probability.

Single contest. We now consider a normal-form game that models a sin-



A game with complete information can be used as a model of a contest when the players are informed about who is going to participate in the contest and about the skills of the participants.



gle contest among two or more players, for different prize allocation mechanisms and production cost functions. A model of a single contest allows us to study situations in which players have no outside options such as investing effort in an alternative contest; we will later discuss games that model simultaneous contests, which provide players with such outside options.

Standard all-pay contest. A classic game that models a contest, we refer to as the standard *all-pay contest*, assumes a prize allocation mechanism that allocates entire prize budget to a highest-effort player with random tie break, and unit marginal production costs. This game corresponds to the well-known game that models an all-pay auction, studied in auction theory. The given prize allocation mechanism is commonly referred to as *perfect discrimination*, because it assumes perfect identification of a highest-effort player, achieved by some flawless mechanism for comparison of individual efforts.

We first discuss Nash equilibrium outcomes in the game with complete information that models the standard all-pay contest. This game does not have a pure-strategy Nash equilibrium. It can be easily verified that for any given effort investments, there is always a player who has a beneficial unilateral deviation. On the other hand, the game always has one or more mixed-strategy Nash equilibria, which were first fully characterized by Baye, Kovenock, and de Vries.²

The game has a unique mixed-strategy Nash equilibrium only in some special cases, such as in a two-player contest, or in a contest with three or more players but where two players have individual skills larger than that of any other player. In general, the game has a continuum of mixed-strategy Nash equilibria. This may be considered a drawback because it implies a lack of predictive power. The mixed-strategy Nash equilibria are payoff equivalent: whenever a game has two or more mixed-strategy Nash equilibria, the expected payoffs in these equilibria are equivalent. In general, the equilibrium outcomes are not equivalent with respect to either the expected total effort or the expected maximum individual effort. It is noteworthy that

Figure 3. Rank order allocation of prizes: Allocation of fixed shares $w_1 \geq w_2 \geq \dots \geq w_n \geq 0$ of a prize budget in decreasing order of effort.



there always exists a mixed-strategy Nash equilibrium in which all but two highest-skill players invest zero effort. The expected total effort in this equilibrium is at least as large as in any other equilibrium.

We now discuss some properties that hold in any mixed-strategy Nash equilibrium. Without loss of generality, assume that players' identities are in decreasing order of their skill parameters. The expected total effort is of value between $v_2/2$ and v_2 . Interestingly, the expected maximum individual effort is always at least half of the expected total effort. This provides a theoretical support for the efficiency of competition-based crowdsourcing services in which a contest owner solicits solutions from multiple workers, but makes use only of the best submitted solution. Intuitively, one would expect that such a production system is bound to be highly inefficient because much of the invested work ends up being wasted. However, by this result, inefficiency can only be to a limited extent in any mixed-strategy Nash equilibrium. With regard to social welfare, there can be some efficiency loss in equilibrium, because a player whose skill is not the highest may have a strictly positive winning probability. However, this can only be up to a limited extent in any mixed-strategy Nash equilibrium: the expected social welfare is always at least 4/5 of the optimum social welfare.

Another noteworthy property is the so-called *exclusion principle*, which refers to the existence of game instances for which the expected total effort in equilibrium can be increased by excluding some players from the competition. In particular, for some game instances, it can be beneficial to exclude the highest-skill player. Intuitively, such exclusion may result in a more intense competition among players with more balanced skills, and, as a result, yield a higher expected total effort.

We now move on to discuss the game with incomplete information that models the standard all-pay contest. We restrict our discussion to prior distributions according to which skills of players are independent and identically distributed random variables. The game has a unique symmetric Bayes-Nash equilibrium, in which players play identical strategies. The expected total effort in this equilibrium is equal to the expected value of the second-highest skill of a player. Interestingly, the expected maximum individual effort is at least half of the expected total effort in any symmetric Bayes-Nash equilibrium, which was established by Chawla, Hartline, and Sivan.⁷ This is exactly the same relation we previously noted to hold between the expected total effort and the expected maximum individual effort in any mixed-strategy Nash equilibrium of the game with complete information.

Rank-order allocation of prizes. We now consider a more general situation where a prize budget can be arbitrarily split among two or more position prizes, which are assigned to players in decreasing order of effort, subject to the constraint that any position prize is at least as large as any lower position prize (see Figure 3). For example, a prize budget may be split between two position prizes such that 2/3 of the prize budget is allocated to first place prize and the remaining part is allocated to second place prize; a common way of splitting a prize budget in Top-Coder contests.

We consider the question of how should a prize budget be split among position prizes to maximize a given objective in equilibrium. Clearly, the answer depends on the choice of the objective, equilibrium concept, heterogeneity of skills, and production costs. Suppose the objective is to maximize the expected total effort in equilibrium of the game with incomplete information, where players have identical prior distributions of skills and unit marginal production costs. Under these assumptions, it is optimal to allocate the entire prize budget to first place prize, which was shown by Moldovanu and Sela.²⁸ Under the same assumptions, allocating the entire prize budget to the first-place prize is also optimal for the objective of maximizing the expected maximum individual effort, which was shown by Chawla, Hartline, and Sivan.⁷ These results hold even more generally for any production cost function with decreasing marginal costs. In contrast, for production cost functions with increasing marginal costs, it may be optimal to split a prize budget among two or more position prizes. The optimality of allocating entire prize budget to first place prize holds also for the game with complete information, under the assumption that players have identical skills and decreasing marginal production costs, as shown by Glazer and Hassin¹⁷ and Ghosh and McAfee.¹⁶

The assumption that in the game with incomplete information the skills of players have identical prior distributions is critical for the optimality of allocating entire prize budget to first place prize. Similarly, the assumption

that in the game with complete information the skills of players are identical is critical for the optimality of allocating entire prize budget to first place prize. If the skills of players have non-identical prior distributions, then there exist game instances such that it is profitable to split the prize budget over two or more position prizes, which is shown by the following example.

Three players, two prizes example. Consider a game where a unit prize budget is split between two position prizes such that $\frac{1}{2} \leq \alpha \leq 1$ is allocated to the first place prize and the remaining part is allocated to the second place prize. Assume there are three players: a high-skill player with the skill parameter of value $v > 1$ and two low-skill players whose skill parameters are of value 1. Assume that each player incurs unit marginal production cost. This game has a mixed-strategy Nash equilibrium such that the two low-skill players play symmetric strategies. This equilibrium is such that in the limit of asymptotically large skill of the high-skill player, the mixed strategy of the high-skill player converges to a uniform distribution on $[1 - \alpha, \alpha]$, and that of low-skill players converges to a uniform distribution on $[0, 1 - \alpha]$. In this limit, the expected effort of the high-skill player is $\frac{1}{2}$, and that of each low-skill player is $(1 - \alpha)/2$. This adds up to the expected total effort of value $\frac{3}{2} - \alpha$. Therefore, we observe the more balanced the split of the prize budget between the two position prizes, the larger the expected total effort.

An interesting question to ask is how should a prize budget be allocated to maximize a given objective in equilibrium, without making a commitment to allocate the entire prize budget to players, no matter what effort investments they make. This question has been resolved for the game with incomplete information and the objective of maximizing the expected total effort by the celebrated work of Myerson.²⁸ In particular, if the skill parameters are independent and identically distributed according to a prior distribution that satisfies a certain regularity condition, it is optimal to award the entire prize budget to a highest-effort player subject to his or her effort being larger than or equal to a minimum required effort,

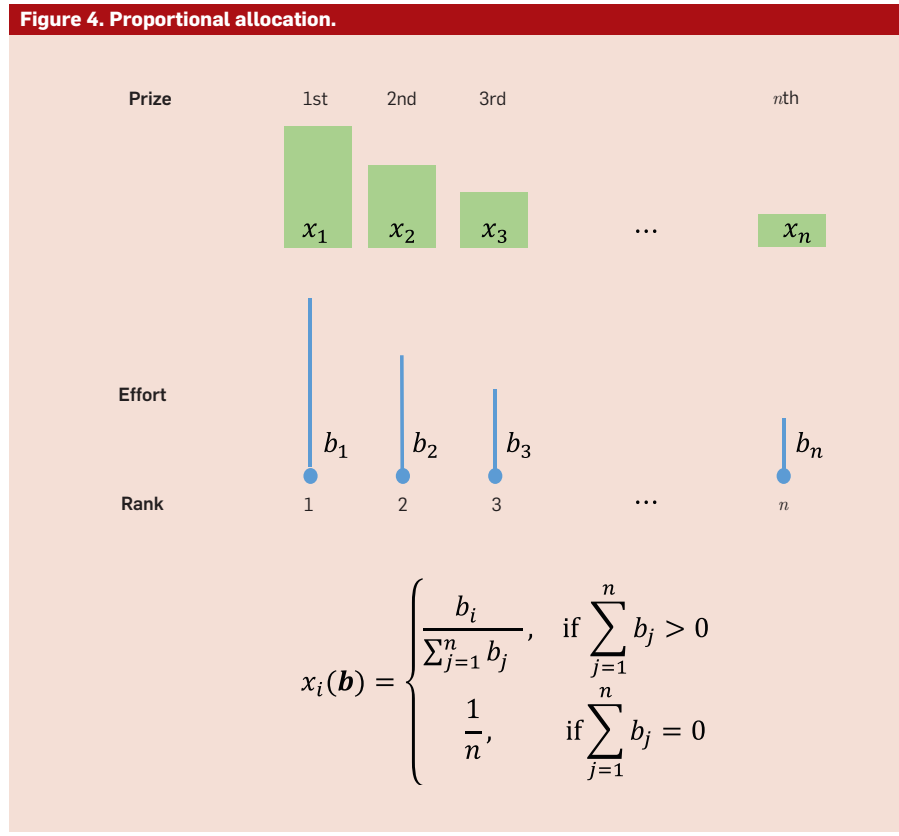
and withhold the prize by the contest owner, otherwise. Chawla, Hartline, and Sivan⁷ have recently established similar characterization of the optimum prize allocation for the objective of maximizing the expected maximum individual effort.

Smooth allocation of prizes. Now consider prize allocation mechanisms that have a positive bias to awarding players who invest high effort, but do not guarantee that the prize is allocated to a highest-effort player. Such prize allocation mechanisms can arise due to various factors. One factor is the stochasticity of production, where individual production outputs are random variables, positively correlated with invested efforts. Another factor is allocation of prizes based on a ranking of players derived from noisy observations of individual production outputs. Such prize allocation mechanisms are referred to be with *imperfect discrimination*. The stochasticity of production may result in prize allocation according to a smooth function of invested efforts, for all vectors of efforts except for some corner cases such as when all players invest zero efforts.

An example of a smooth allocation of prizes is *proportional allocation*

that splits a prize budget among players in proportion to invested efforts, conditional on at least one player investing a strictly positive effort; otherwise, the prize is evenly split among players (Figure 4). A smooth prize allocation may be enforced by the design of a resource allocation mechanism. For example, proportional allocation has been used for allocation of computing resources³⁷ and network bandwidth.²² Such resources typically consist of a large number of small units and, thus, for any practical purposes, can be regarded as infinitely divisible resources.

A more general class of smooth allocations is defined by allocating in proportion to an increasing positive-valued function of invested effort, referred to as a *general logit allocation*. A special case is allocation in proportion to a power function of invested effort, with a positive exponent parameter r . This is commonly referred to as *Tullock allocation*, which has been studied extensively in the literature on rent-seeking contests.⁴⁰ Proportional allocation is a special case of a Tullock allocation for the value of parameter r equal to 1. The larger the value of parameter r , the larger the share of



the prize allocated to a highest-effort player. For more details about smooth allocations, see, for example, Corchon and Dahm¹⁰ and Vojnović.⁴²

One may ask how do equilibrium outcomes in the game that models the standard all-pay contest compare with those in the game with a smooth prize allocation, say, according to proportional allocation. A first notable difference is that unlike the game that models the standard all-pay contest, the game with proportional allocation has a pure-strategy Nash equilibrium, which is unique.

The total effort in any pure-strategy Nash equilibrium is guaranteed to be at least $v_2/2$. The total effort increases in the highest-skill parameter v_1 and it can be larger than v_2 . This is in contrast to the game that models the standard all-pay contest, where the total effort in any mixed-strategy Nash equilibrium is at most v_2 . One may ask whether there exists a smooth allocation of prizes that guarantees the total effort to be within a constant factor of v_1 in any pure-strategy Nash equilibrium. The answer is negative.⁴¹ This gives us a useful insight that randomized prize allocations can achieve a larger total effort, but there are fundamental limits that cannot be surpassed.

The maximum individual effort can be an arbitrarily small fraction of the total effort; for example, this is so for the simple game instance with equally skilled players by taking the number of players to be sufficiently large. This is in contrast to the game that models the standard all-pay contest where we noted that in any mixed-strategy Nash equilibrium, the expected maximum individual effort is at least 1/2 of the expected total effort.

The social welfare in any pure-strategy Nash equilibrium of the game with proportional allocation is always at least 3/4 of the optimum value, a result by Johari and Tsitsiklis.²¹ It has been shown the game with proportional allocation is a smooth game (for example, see Roughgarden³⁴), which implies that the expected social welfare is at least 1/2 of the optimum value in any mixed-strategy Nash equilibrium.

Unlike the game that models the standard all-pay contest, the exclusion principle does not hold for the game that models the contest with propor-



An interesting question to ask is how should a prize budget be allocated to maximize a given objective in equilibrium, without making a commitment to allocate the entire prize budget to players, no matter what effort investments they make.



tional allocation.¹⁴ For the game that models the contest with proportional allocation, the total effort in the pure-strategy Nash equilibrium cannot be increased by excluding some of the players from competition.

Simultaneous contests. In the context of online services, a contest is often run simultaneously with other contests. For example, in competition-based crowdsourcing services, there are typically many open contests at any given time. Similarly, in online labor marketplaces, there are usually many open jobs at any given time. Multiple open contests provide players with alternative options to invest efforts, which can have a significant effect on the effort invested in any given contest. A player can invest effort only in a limited number of contests over a period of time, or he or she has a limited effort budget to invest over available contests. A worker may only be able to produce a high-quality work by focusing to a small number of projects at any given time, or he or she may only be able to devote a limited number of work hours per week. Game theory provides us with a framework to study the relation between the values of prizes offered by different contests and the effort investments across different contests in a strategic equilibrium.

We consider games that model *simultaneous standard all-pay contests* that offer prizes of arbitrary values. Such games have been studied for different types of production costs. We first consider the case where production costs are such it is feasible for each player to participate in at most one contest, in which he or she incurs a unit marginal production cost. In such a game, strategic decision making of a player consists of two components: choosing in which contest to invest effort, and deciding how much effort to invest in the chosen contest. This strategic decision making is informed by the available information, which consists of the values of prizes offered by different contests and the prior information about the skills of players. We consider the game with incomplete information, where the skill parameters of players are independent and identically distributed according to a prior distribution. This game has a symmetric Bayes-

Nash equilibrium, which admits an explicit characterization, established in DiPalantino and Vojnovic'.¹¹ In this equilibrium, there is a *segregation* of players in different skill levels, such that the players of the same skill level choose contests according to identical mixed strategies. A player of a higher skill level chooses a contest to participate from a smaller set of contests that offer highest prizes. A higher expected participation is attracted by contests that offer high prizes, according to a relation that exhibits diminishing returns with respect to the values of the prizes.

Another type of production costs is when each player is endowed with an effort budget that he or she can split arbitrarily over available contests. This game is closely related to so-called *Colonel Blotto game*: there are two colonels and two or more battlefields; each colonel is endowed with a number of troops that are simultaneously deployed over the battlefields; a battlefield is won by the colonel who places a larger number of troops on this battlefield, and the game is won by the colonel who wins more battlefields. A *continuous Colonel Blotto game* assumes that each colonel is endowed with an infinitely divisible amount of army force.

The game with players endowed with effort budgets has a rich set of equilibrium properties. There are game instances with a continuum of mixed-strategy Nash equilibria. For example, this is the case for the game with two players that have non-identical effort budgets and two or more standard all-pay contests that offer identical prizes. When players have identical effort budgets, the game has both pure and mixed-strategy Nash equilibria in which each player invests all his or her effort in one contest, provided that the number of players is sufficiently large. In the limit of many players, the equilibrium participation of players across different contests is proportional to the values of prizes.

Games that model simultaneous contests with players endowed with effort budgets have also been studied for other prize allocation mechanisms, including proportional allocation and equal-share allocation. The game that models simultaneous

contests with proportional allocation and players endowed with effort budgets is not guaranteed to have a pure-strategy Nash equilibrium. A sufficient condition for the existence of a pure-strategy Nash equilibrium is that each contest has at least two players with strictly positive skill parameters. The social efficiency in a pure-strategy Nash equilibrium can be arbitrarily low in a worst case.

Sharing of the utility of production.

There have been various studies of production systems where agents invest effort in one or more activities, which results in a utility of production that is shared among contributors according to a utility sharing mechanism. Some online services rely on user contributions and award credits to incentivize contributions. For example, some online services rely on user-generated content, such as questions and answers in online Q&A services, and award credits in terms of attention or reputation points, which are commensurate to user contributions. Sharing the utility of production has been also studied in the context of cognitive labor and allocation of scientific credit, for example, Kitcher²³ and Kleinberg and Oren.²⁴

A central question here is about the social efficiency of production in strategic equilibrium outcomes. Several factors can contribute to social inefficiency of production, including the choice of the utility sharing mechanism, the nature of the utility of production functions, and the nature of production cost functions. Special attention has been paid to *local utility sharing mechanisms*, which specify the shares of the utility of production associated with a project exclusively based on the effort investments in this project, and not on the effort investments in other projects. It is of interest to understand social efficiency of *simple* local utility sharing mechanisms, for example, allocating a priori fixed shares of the utility of production in decreasing order of individual contributions or allocating in proportion to individual contributions.

The nature of the utility of production is a critical factor for the social efficiency of equilibrium outcomes. If the utility of production is allowed to be according to a non-monotonic func-

tion of effort investments, then there are game instances for which the utility of production in a pure-strategy Nash equilibrium is an arbitrarily small fraction of the optimum; for example, this can be for a single project game with proportional allocation. This is an instance of a general phenomenon known as the *tragedy of the commons*,¹⁹ referring to an inefficient use of congestible resources that arises from non-cooperative behavior of selfish agents. The nature of the production cost functions is also a critical factor. If, in a single project game with proportional allocation, the utility of production is a monotone function, but players incur unit marginal production costs, then a similar inefficiency of production can arise.

Are there conditions for the games under consideration under which equilibrium is guaranteed to exist and all equilibria are approximately socially efficient? Here we may settle for the utility of production to be at least a constant factor of the optimum value. Such conditions have been identified by Vetta⁴¹ for the class of games referred to as *monotone valid utility games*. A game is said to be a monotone valid utility game if the players' payoffs are according to utility functions whose sum is less than or equal to the value of a social utility function, and the following two conditions hold. The game is required to satisfy a *monotonicity condition*, which restricts to social utility functions whose value cannot increase by some player opting out from participation. The game is also required to satisfy a *marginal contribution condition*, which restricts each player's utility to be at least as large as his or her marginal contribution to the social utility. In the context of games that model simultaneous projects, whether or not the marginal contribution condition holds depends on the nature of the utility of production functions and the utility sharing mechanism. For example, the marginal contribution condition holds if the project utility functions are increasing functions with diminishing returns in the total effort invested in a project, and the utility sharing is according to proportional allocation. For monotone valid utility games, the utility of production in any pure-strategy Nash

equilibrium is guaranteed to be at least $1/2$ of the optimum value.

The approximate social efficiency of the utility of production in any pure-strategy Nash equilibrium has been established under the assumption that project utility functions have diminishing returns. The diminishing returns of the utility of production are representative of production systems in which individual contributions are substitutes. If, on the other hand, individual contributions are complements (that is, the utility of production has increasing returns), then the social efficiency in an equilibrium outcome can be arbitrarily low. In such cases, the utility of production in a pure-strategy Nash equilibrium cannot be guaranteed to be a constant-factor of the optimum value, but it is always at least $1/k$ of the optimum value, where k is the maximum number of players participating in a project.

Sequential contests and tournaments. So far we discussed games that model contests where players simultaneously invest effort. A variety of games have been studied that model contests with some elements of sequential play. A coverage of these games and related work is available.⁴² Here we only mention some of these games: a single contest with sequential effort investments; a multi-round two-player contest where the winner is the player who first wins a given number of rounds more than the opponent, referred to as *tug-of-war*; a contest in which each player continuously invests effort until dropping out and the contest ends as soon as the number of players that are still in the competition is equal to the number of available prizes, referred to as *war-of-attrition*;⁵ a multi-round contest that ends as soon as the utility of cumulative effort exceeds a threshold whose value is private information of the contest owner;³⁵ and, a contest where prizes are allocated over multiple rounds and each player competes until he or she wins a prize.⁸

Common contest architecture has the form of a *single-elimination tournament*, defined by a directed *tree* and a *seeding* of players. Each contest of the tournament has one winner and all players who lose in a contest are

eliminated from further competition. The winner of the tournament is the player who wins all contests in which he or she participates. A typical single-elimination tournament consists of two-player contests and is defined by a binary tree and a seeding of players. Seeding procedures have been studied with respect to various criteria, such as the winning probability of the highest-skill player. These studies have been pursued under two different assumptions: contest outcomes are assumed to be independent random events according to given winning probabilities; and in each round of the tournament, the players who participate in this round make strategic effort investments accounting for their prospective payoffs in subsequent rounds of the tournament.

Skill-Rating Methods

An important component of some online services is a skill-rating system that uses as input observed contest outcomes. For example, a contest outcome may be a full ranking, that is, an ordered list of participants in the contest in decreasing order of individual performance, or a partial ranking such as a top-1 list that contains information about who participated in a contest and who was the winner in this contest. The skill ratings are used for various purposes, such as for creation of league tables, leaderboards, seeding of tournaments, and match-making in online labor platforms. Popular skill-rating systems include TrueSkill, used in online gaming,²⁰ TopCoder skill-rating system, and skill-rating systems used in various sport competitions.¹³

A common requirement for skill-rating systems is to allow for prediction of contest outcomes. For example, such predictions are used in online games for the purpose of matching equally skilled players, which results in interesting matches with uncertain outcomes. The design of skill-rating systems is often required to be based on simple and easy to understand principles, which are often made public information. The skill-rating systems often use only a few parameters to represent an individual's skill; for example, using a scalar parameter for a point estimate

and one extra parameter for the uncertainty of the estimate.

Statistical models of ranking outcomes. The design of skill-rating systems is based on statistical models of ranking outcomes introduced by statisticians as early as in 1920s. A commonly used statistical model of ranking outcomes was introduced by Thurstone.³⁹ Under this model, each comparison of a given set of individuals results in a ranking of these individuals generated as follows. The individuals are associated with latent performance random variables that are assumed to be independent across different individuals and comparisons. The ranking outcome of a comparison is assumed to be in decreasing order of individual performance. Each individual performance is equal to a deterministic skill parameter plus a zero mean noise random variable. The value of the skill parameter is unknown and has to be inferred from the observed ranking outcomes. The noise random variables are assumed to be independent and identically distributed over different individuals and different comparisons.

Specifically, for a comparison of a set S of individuals, each individual $i \in S$ is associated with performance $b_i = v_i + \epsilon_i$, where v_i is a real-valued skill parameter and ϵ_i is a zero mean noise random variable. A ranking outcome is derived from admitting that i is ranked higher than j whenever their respective performances satisfy $b_i > b_j$.

A common assumption is that noise random variables are according to a Gaussian distribution, with zero mean and known variance β^2 . This assumption was made in the original work by Thurstone for pair comparisons, and has been admitted by many popular skill-rating systems, including TrueSkill, TopCoder skill-rating system, and skill-rating systems used in various sport competitions. The probability that individual i is ranked higher than individual j , in a comparison that involves these two individuals, is given by

$$p_{i,j} = \mathbf{P}[b_i > b_j] = \Phi\left(\frac{1}{\sqrt{2}\beta}(v_i - v_j)\right)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution.

Another well-known model is the Bradley-Terry model, first introduced by Zermelo in 1920s⁴⁵ and later popularized in 1950s by the work of Bradley and Terry⁴ and others. Under the Bradley-Terry model, the probability that individual i is ranked higher than individual j , in a comparison that involves these two individuals, is given by

$$p_{i,j} = \frac{\theta_i}{\theta_i + \theta_j}$$

where θ_i and θ_j are positive-valued skill parameters. According to the Bradley-Terry model, the winning probability of an individual in a pair comparison with another individual is proportional to his or her skill parameter. The natural generalization to comparison sets of two or more individuals, where the winning probabilities are proportional to the skill parameters, is known as the Luce's choice model. Another generalization is a model of full ranking outcomes for comparison sets of two or more individuals, defined by sampling individuals from a given comparison set without replacement with probabilities proportional to their skill parameters; this is known as the Plackett-Luce model. The Luce's choice model is a special instance of a Thurstone model with noise random variables according to a double-exponential distribution with zero mean and variance β^2 . In this case, we have

$$p_{i,j} = \mathbf{P}[b_i > b_j] = \frac{1}{1 + e^{-\frac{\pi}{\sqrt{6}\beta}(v_i - v_j)}}$$

that corresponds to the Bradley-Terry model by using the change of parameters $v_i = \sqrt{6}\beta/\pi \log(\theta_i)$.

The statistical models of ranking outcomes discussed so far have been extended to accommodate various requirements of modern applications. For example, they have been extended to allow for skill rating based on observed outcomes of team competitions, which arises in online gaming applications. This extension is based on a model that assumes a team performance to be according to a given function of individual performances. For instance, a team performance may be assumed to be a linear function of individual performances, such as in the TrueSkill rating system. An area in which advances have been made is on

An important component of some online services is a skill ratings system that uses as input observed contest outcomes.

statistical inference methods, which we briefly review as follows.

Statistical inference methods.

Having admitted a statistical model of ranking outcomes, it remains to choose a statistical inference method for estimation of skill parameters based on observed ranking outcomes. Two approaches are in common use: a frequentist approach and a Bayesian approach. The frequentist approach considers skill parameters as unknown parameters and estimates them by minimizing a given loss function, for example, the negative log-likelihood in the case of the maximum likelihood estimation. The Bayesian approach considers skill parameters as random variables with a given prior distribution, and amounts to computing the posterior distribution of these random variables conditional on the observed ranking outcomes.

Frequentist inference. Statistical models of pair comparisons, such as the Thurstone model with either Gaussian or double-exponential distribution of noise, have a unique maximum likelihood estimator (up to an additive constant) provided that the adjacency matrix, specifying how many times different pairs of individuals are compared in the input data, is irreducible. An *adjacency matrix* is said to be irreducible if the corresponding graph, we refer to as a comparison graph, is connected. It was recently shown that the accuracy of the maximum likelihood parameter estimator critically depends on how well the comparison graph is connected, for example, Hajek, Ox and Xu¹⁸ and Vojnovic and Yun.⁴³ Specifically, a key parameter is the *algebraic connectivity* of the comparison graph, defined as the second smallest eigenvalue of the Laplacian matrix of the comparison graph. Another line of recent research is on various iterative methods for skill parameter estimation, including gradient-descent based methods for minimizing the negative log-likelihood function, as well as alternative methods based on spectral properties of matrices and random walks, for example, Neghaban, Oh, and Shah.³⁰

Bayesian inference. For statistical models of ranking outcomes according to a Thurstone model, the posterior distribution of an individual's

skill is a marginal distribution of the posterior joint distribution of a multivariate variable that consists of individual skills and individual performances. This posterior joint distribution consists of several factors that are conveniently represented by a graphical model, a way to represent the information about which factors depend on which variables. The marginal posterior distributions of skills can be computed using standard message-passing methods for inference in graphical models, such as the sum-product algorithm. It is common to approximate a marginal distribution of a skill variable with a distribution from an assumed family of distributions; for example, assuming the family of Gaussian distributions, as done in the TrueSkill rating system. The approximate Bayesian inference amounts to approximating marginal posterior distributions of skills by distributions from the given family of distributions, assuming that marginal prior distributions belong to this family of distributions.

Future Directions

Strategic game models of contests provide plenty of interesting hypotheses about what strategic user behavior may arise in different contest situations. Future work must be devoted to narrowing the gap between theoretical results and empirical validations. The availability of online services whose design is based on contests and the collected data provides us with an opportunity to test the existing theories and guide the development of new contributions to contest theory. Another research direction is to study statistical inference methods for various contest designs, such as in the recent study of A/B testing for auctions.⁶

While the skill-rating methods have been studied extensively over many years, some interesting questions still remain open. Most skill-rating methods represent an individual's skill by a scalar parameter. In many situations, however, it is of interest to consider an individual's skill over multiple dimensions; for example, an online worker may have different types of skills such as analytical problem solving, strategic business planning, and software programming

skills. Another interesting direction is to study statistical inference methods for statistical models of ranking outcomes that allow for a larger set of unknown parameters. For example, in an online labor platform, a ranking of job applicants would depend not only on the idiosyncratic skills of the applicants, but also on the specific job requirements, both of which may have uncertainties. Another direction is to develop solid theoretical foundations for individual skill rating based on observed team performance outputs. Current statistical inference methods used in practice assume simple models of team performance, such as that a team performance is the sum of individual performances, which may not always be valid in practice. □

References

1. Archak, N. Money, glory and cheap talk: analysing strategic behavior of contestants in simultaneous crowdsourcing contests on TopCoder.com. In *Proceedings of WWW '10* (Raleigh, N.C., 2010), 21–30.
2. Baye, M.R., Kovenock, D. and de Vries, C.G. The all-pay auction with complete information. *Econ. Theory* 8, 2 (1996), 291–305.
3. Bishop, D.T. et al. The war of attrition with random rewards. *J. Theoretical Bio* 70, 1 (1978), 85–124.
4. Bradley, R.A. and Terry, M.E. Rank analysis of incomplete block designs: I. Method of paired comparisons. *Biometrika* 20, 3–4 (1952), 334–345.
5. Bulow, J. et al. The generalized war of attrition. *Amer. Econ. Rev.* 39, 3–4 (1997), 324–345.
6. Chawla, S., Hartline, J.D. and Nekipelov, D. A/B testing of auctions. In *Proceedings of ACM EC '16* (Maastricht, Netherlands, 2016), 856–868.
7. Chawla, S., Hartline, J.D. and Sivan, B. Optimal crowdsourcing contests. In *Proceedings of SODA '12*, (Kyoto, Japan, 2012), 856–868.
8. Clark, D.J. and Riis, C. Competition over more than one prize. *Amer. Econ. Rev.* 88, 1 (1998), 276–289.
9. Corchon, L.C. The theory of contests: A survey. *Rev. Econ. Design* 11 (2007), 69–100.
10. Corchon, L. and Dahm, M. Foundations for contest success functions. *Econ. Theory* 88, 1 (2010), 81–98.
11. DiPalantino, D. and Vojnovic, M. Crowdsourcing and all-pay auctions. In *Proceedings of ACM EC '09* (Stanford, CA, 2009), 119–128.
12. Doan, A., Ramakrishnan, R. and Halevy, A.Y. Crowdsourcing systems on the World-Wide Web. *Commun. ACM* 54, 4 (Apr. 2011), 85–96.
13. Elo, A.E. *The rating of chessplayers*. Ishi Press International, 1978.
14. Franke, J., Kanzow, C., Leininger, W. and Schwartz, A. Lottery versus all-pay auction contests: A revenue dominance theorem. *Games and Economic Behavior*, 13 (2014), 116–126.
15. Galton, F. The most suitable proportion between the value of first and second prizes. *Biometrika* 1, 4 (1902), 385–399.
16. Ghosh, A. and McAfee, R.P. Crowdsourcing with endogenous entry. In *Proceedings of WWW '12* (Lyon, France, 2012), 999–1008.
17. Glazer, A. and Hassin, R. Optimal contests. *Economic Inquiry* 26, 1 (1988), 133–143.
18. Hajek, B., Oh, S. and Xu, J. Minimax-optimal inference from partial rankings. In *Proceedings of NIPS '14*, (Montreal, Quebec, 2014), 1475–1483.
19. Hardin, G. The tragedy of the commons. *Science* 162, 3859 (1968), 1243–1248.
20. Herbrich, R., Minka, T. and Graepel, T. TrueSkill: A Bayesian skill rating system. In *Proceedings of NIPS '06*, (Vancouver, B.C., 2006), 569–576.
21. Johari, R. and Tsitsiklis, J.N., Efficiency loss in a network resource allocation game. *Math. Operations Res* 29, 3 (2004), 402–435.
22. Kelly, F. Charging and rate control for elastic traffic. *European Trans. Telecommun.* 8, 1 (1997), 33–37.
23. Kitcher, P. The division of cognitive labor. *J. Philosophy* 87, 1 (1990), 5–22.
24. Kleinberg, J. and Oren, S. Mechanisms for (mis) allocating scientific credit. In *Proceedings of STOC '11* (San Jose, CA, 2011), 529–538.
25. Konrad, K.A. Strategy in Contest—An Introduction. *WZB-Markets and Politics Working Paper N. SP II 2007-01*: 2007 (<http://ssrn.com/abstract=960458>).
26. Krueger, A.O. The political economy of the rent-seeking society. *Amer. Econ. Rev.* 64, (1974), 291–303.
27. Lazear, E.P. and Rosen, S. Rank-order tournaments as optimum labor contracts. *J. Pol. Econ.* 89, 5 (1981), 841–864.
28. Moldovanu, B. and Sela, A. The optimal allocation of prizes in contests. *American Econ. Rev.* 91, 3 (2001), 542–558.
29. Myerson, R.B. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
30. Neghaban, S., Oh, S., and Shah, D. Iterative ranking from pair-wise comparisons. In *Proceedings of NIPS '12*, (Lake Tahoe, NV, 2012), 2483–2491.
31. Nisan, N., Roughgarden, T., Tardos, E. & Vazirani, V. V., 2007. *Algorithmic Game Theory*. Cambridge University Press.
32. Nitzan, S., 1994. Modelling rent-seeking contests. *Eur. J. Polit. Econ.* 10 (1994), pp. 41–60.
33. Roughgarden, T. Algorithmic game theory. *Commun. ACM* 53, 7 (2010), 78–86.
34. Roughgarden, T. Intrinsic robustness of the prize of anarchy. *Commun. ACM* 55, 7 (2012), 116–123.
35. Shaili, J., Yiling, C., Parkes, D.C. Designing incentives for online question and answer forums. In *Proceedings of ACM EC '09* (Stanford, CA, 2009), 129–138.
36. Shoham, Y. Computer science and game theory. *Commun. ACM* 51, 8 (2008), 75–79.
37. Stoica, I. et al. A proportional share resource allocation algorithm for real-time, time-shared systems. In *Proceedings of the 17th Real-Time Systems Symposium*. (Washington, D.C., 1996), 288–299.
38. Tang, J.C. et al. Reflecting on the DARPA Red Balloon Challenge. *Commun.* 54, 4 (2011), 78–85.
39. Thurstone, L.L. A law of comparative judgment. *Psychological Review* 34, 2 (1927), 273–286.
40. Tullock, G. Efficient rent seeking. In *Theory of the Rent-Seeking Society*. A&M University Press (1980), 131–146.
41. Vetta, A., *Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions*. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science* (2002), 416–425.
42. Vojnović, M. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press, 2016.
43. Vojnović, M. and Yun, S.-Y., Parameter estimation for generalized Thurstone choice models. In *Proceedings of ICML '16* (New York City, NY, 2016).
44. Yang, J., Adamic, L., and Ackerman, M. Crowdsourcing and knowledge sharing: Strategic user behavior on Taskcn. In *Proceedings of the ACM EC '09* (Chicago, IL, 2008).
45. Zermelo, E. Die Berechnung der Turnier-Ergebnisse at sein Maximumproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift* 29 (1929), 436–460.

Milan Vojnović (m.vojnovic@lse.ac.uk) is a professor of data science in the Department of Statistics, London School of Economics, U.K., where he serves as program director of MSc in Data Science.

©2017 ACM 0001-0782/17/05 \$15.00.



Watch the author discuss his work in this exclusive *Communications* video. <http://cacm.acm.org/videos/contest-theory>