Contests: Equilibrium Analysis, Design and Learning

Abheek Ghosh

Department of Computer Science

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2024

Abstract

Contests are games where agents compete by making costly and irreversible investments to win valuable prizes. They model diverse scenarios ranging from crowdsourcing to competition among Bitcoin miners. Using tools from theoretical computer science, we contribute to the understanding of the agents' behavior in contests and make design recommendations to optimize practical objectives. In particular, we (i) analyze learning dynamics in Tullock contests using tools from probabilistic analysis of algorithms and optimization, (ii) design contests that improve diversity in participation, and (iii) study the existence, welfare efficiency, and computational complexity of equilibrium in a class of simultaneous contests.

Acknowledgements

First, I thank my supervisors, Edith Elkind and Paul Goldberg, for their belief, support, and inspiration. They have been immeasurably generous with their time and ideas. From them, I have learned almost everything I know about doing research.

I have benefited greatly by being mentored by some excellent researchers over the years. I especially thank Umang Bhaskar, Dheeraj Nagaraj, Nikolai Gravin, Milind Tambe, Manish Jain, and Ufuk Topcu.

In graduate school, I have been fortunate to work with and talk to many amazing students, researchers, and friends: Jiarui, Nicholas, Mohsen, Sonja, and Martin. Thank you also to my other wonderful friends at Oxford: Tim, Satwik, Julian, Maxime, Luca, Cornelius, and Souvik; and my tennis teammates, including Elliott, Oliver, Teo, Luis, Tom, and Fabian, and our coach, Jarlath.

It is my greatest pleasure to thank my parents, Ashis and Urbi, and my sisters, Uttha and Urna, for their love and for shaping me into who I am. Finally, I thank Sree Sree Thakur, who has been a constant source of inspiration, love, and guidance, and Acharyadev, for being an ever-reliable mentor throughout my life and for encouraging me to do a Ph.D.; they have my eternal love and gratitude.

Contents

List of Figures

Chapter 1

Introduction

There are numerous instances in our society where agents compete for valuable rewards by playing costly actions. Pharmaceutical companies invest in the research and development of drugs and vaccines and compete for market share and profit. Bitcoin miners spend money to acquire and run computational resources with the hope of mining blocks and getting associated rewards. In these examples, the economic agents compete by making irreversible investments before the outcome of the competition is known. Lobbying activities, competitions for promotions, sports, crowdsourcing contests, competition for college seats, to name a few more, all have this property. This type of competition has been widely studied in the economics literature and is generally called a contest [106].

1.1 Models

In the standard formulation, n agents compete by simultaneously producing non-negative outputs. Let $[n] = \{1, 2, \ldots, n\}$. Agent $i \in [n]$ produces an output of $x_i \in \mathbb{R}_{\geq 0}$ and incurs a cost of $c_i(x_i)$, where $c_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is an increasing cost function. Let $\boldsymbol{x} = (x_i)_{i \in [n]}$ denote the output profile. The allocation function $\psi = (\psi_i)_{i \in [n]}$, where $\psi_i : \mathbb{R}^n_{\geq 0} \to \mathbb{R}_{\geq 0}$, maps the output profile $\boldsymbol{x}$ to the reward received by agent i, $\psi_i(\boldsymbol{x})$. The utility of agent i, $u_i(\boldsymbol{x})$, is given by

$$u_i(\boldsymbol{x}) = \psi_i(\boldsymbol{x}) - c_i(x_i). \tag{1.1}$$

The cost function c_i is usually assumed to be a linear or a convex function.

Contests capture a diverse set of real-life applications, and depending upon the use case, several models for the allocation function ψ have been proposed. Below, we introduce a few models that are particularly relevant to our work.

1.1.1 Proportional Allocation—Tullock Contest

In a Tullock contest [104], there is a prize of unit value (normalized) that is allocated to the agents proportionally. In particular, agent i receives a fraction of the prize proportional to x_i in the non-degenerate case where at least one agent produces a strictly positive output, else $1/n$. The utility of agent i is

$$u_i(\boldsymbol{x}) = \frac{x_i}{\sum_j x_j} - c_i(x_i). \tag{1.2}$$

Tullock contests are one of the most widely used and have been applied to study political lobbying, rent-seeking, and crowdsourcing, among other applications (see, e.g., [67]). As an example of a concrete and relatively new application, it almost exactly models the game among Bitcoin miners [31, 72]. Bitcoin, the largest cryptocurrency with a market cap of over \$1 trillion (February 2024), maintains its platform in a decentralized manner with the help of miners who do costly computational tasks to compete for valuable rewards.

1.1.2 Rank-Order Allocation

In a rank-order allocation contest, there are n prizes $\boldsymbol{w} = (w_1, w_2, \ldots, w_n)$, where $w_j \geq w_{j+1} \geq 0$ for $j \in [n-1]$. The prize w_1 is awarded to the agent with the highest output, w_2 to the second highest, and so on; an agent receives one of the prizes based on their rank. Let us assume there are no ties for now. Suppose that $x_{i_1} > x_{i_2} > \ldots > x_{i_n}$. The utility of agent i is

$$u(\boldsymbol{x}_i) = \sum_{j \in [n]} w_j \, \mathbb{1}\{x_i = x_{i_j}\} - c_i(x_i), \tag{1.3}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function.

The all-pay auction is a classic example that can be captured by rank-order allocation contests. It would correspond to a rank-order allocation contest where $w_1 = 1$ and $w_j = 0$ for $j > 1$, i.e., all prizes except for the highest output (bid) is 0, and the cost is $c_i(x_i) = x_i/v_i$, where v_i is agent i's valuation for the item.[1] Other applications include sports, crowdsourcing, and grading [81, 29].

[1] By scaling the utility function of agent i by v_i, which does not change the strategies of the agents, we get the standard all-pay auction formulation.

1.1.3 Equal-Sharing Allocation

In equal-sharing allocation, there is a threshold $\tau > 0$ such that all agents who produce output above τ are considered successful and get an equal share of a total reward that is a function of the number of successful candidates. Let the total reward function be $\rho : [n] \to \mathbb{R}_{\geq 0}$, where $\rho(k)$ is the total reward when there are k successful candidates. Let the number of successful candidates be $\eta_\tau(\boldsymbol{x}) = |\{i \mid x_i \geq \tau\}|$. The utility of agent i is given by

$$u_i(\boldsymbol{x}) = \mathbb{1}\{x_i \geq \tau\} \frac{\rho(\eta_\tau(\boldsymbol{x}))}{\eta_\tau(\boldsymbol{x})} - c_i(x_i). \tag{1.4}$$

These allocation functions have been used to model applications where the reward is associated with reputation and social status [62]. They are generally accompanied by extensions where the agents produce multi-dimensional output, and several contests may occur simultaneously (discussed below) [76].

1.1.4 Other Modeling Choices

Two natural extensions of the models discussed previously are (i) multi-dimensional output by agents and (ii) simultaneous contests.

In the case of multi-dimensional output, instead of producing a non-negative output in $\mathbb{R}_{\geq 0}$, each agent produces an output in $\mathbb{R}^d_{\geq 0}$ for some $d \geq 2$. The d dimensions correspond to the different activities (or the different aspects of a job) an agent can spend their effort on. The agent may have varying costs for producing output along these activities, and the contest may value these activities differently [76].

In the case of simultaneous contests, multiple contests run in parallel, and the agents have to distribute their limited resources across these contests. The agents either have a budget constraint or a non-separable cost function for their output across these contests, which forces them to pick and choose among these contests. Such models are relevant in diverse applications, including crowdsourcing [39] and blockchains [18], and have also been combined with multi-dimensional output to model competitions for market and social influence [76].

1.2 Overview of Our Contributions

We group the chapters in this thesis into three parts, discussed briefly below. Broadly, the research goals in contest theory are to (i) understand the behavior of agents in these contests and (ii) make design recommendations to influence this behavior towards certain objectives. The natural first step toward understanding the behavior of agents is to perform an equilibrium analysis. However, analytical characterization of the equilibrium is not always possible, and even if this is possible, the agents might not always reach this equilibrium. We use tools from theoretical computer science to contribute to this research direction. Specifically, in Part I, we use techniques from probabilistic analysis of algorithms, convex optimization, and dynamical systems to study learning in Tullock contests. In Part II, we shift our focus to contest design. Although, in many applications, the rules of the contests may be intrinsic, in other applications, they may be decided by a designer. Our work in contest design studies two types of objectives motivated by practical applications. Finally, in Part III, we consider a model of simultaneous contests and study the existence, computational complexity, and price of anarchy of equilibria in these contests. We discuss these contributions in more detail below.

1.2.1 Part I: Dynamics in Tullock Contests

The existence of an equilibrium may not always be a good predictor of the agents' behaviors [52]. Assumptions like the availability of information and rationality of agents, required to ensure equilibrium behavior in games, may not hold in practice for many applications of contests. This line of research aims to address this issue by studying natural learning dynamics in contests.

In Chapter 2, we study best-response (BR) dynamics in Tullock contests with convex cost functions. We show that the dynamics converges for homogeneous agents (all agents have the same cost function) but not for non-homogeneous agents. For homogeneous agents, the time taken to converge to an ϵ-approximate equilibrium has a logarithmic dependency on ϵ for three or more agents (and double-logarithmic for two agents), polynomial dependency on the number of agents n, and a double-logarithmic dependency on the initial state. Our bounds are tight w.r.t. (with respect to) ϵ and the initial state and polynomially tight w.r.t. the number of agents n. We get these results by reducing the BR dynamics to a dynamics we call the discounted-sum (DS) dynamics. We perform a suitable mix of worst-case and probabilistic analysis of the DS dynamics and use this to show the convergence of the BR dynamics. This result indicates that we can expect simple, myopic agents with similar convex cost functions to rapidly converge to the equilibrium in Tullock contests.

In Chapter 3, still focusing on BR dynamics in Tullock contests, we take a deeper look into the specific case of linear cost functions, also known as lottery contests. Our analysis for general convex costs in Chapter 2 provides tight convergence-rate bounds w.r.t. ϵ and the initial state, but not w.r.t. the number of agents n, which we address in Chapter 3 for lottery contests using a different analysis. We provide almost tight $\tilde{O}(n)$ bounds using techniques from convex optimization and Markov processes.

Our results in Chapters 2 and 3 show convergence of BR dynamics for homogeneous agents but non-convergence for non-homogeneous agents. Given these results, in Chapter 4, we study other related learning dynamics. We show that *continuous-time* BR dynamics in Tullock contests with convex costs converges to the unique equilibrium using Lyapunov-style arguments.[2] Notice that this convergence result holds for non-homogeneous agents. With mild technical assumptions, we then extend this result to related discrete-time dynamics, e.g., when the agents best respond to the empirical average action of other agents or when the agents play a discrete approximation of the continuous-time BR dynamics (with small but not necessarily infinitesimally small steps). We can also use these results to provide an algorithm for computing an approximate equilibrium. These results indicate that the equilibrium is a reliable predictor of the agents' behavior in these games.

1.2.2 Part II: Contest Design

One of the primary motivations for organizing contests is to elicit effort from participants, e.g., olympiads and hackathons encourage students to put effort and learn about specific subjects and technologies. Here, the designer may want to elicit an adequate output from several agents instead of a very high output from a few agents. For example, in a crowdsourcing task, such as a survey, it may be more valuable to get many contributors to give adequate responses than to get only a few people to submit perfect responses. In other applications, it may be important to elicit higher participation from a target group, e.g., hackathons to get women interested in AI [71] and crowdsourcing to get machine learning data from underrepresented groups. In Chapters 5 and 6, we design contests with rank-order allocation of prizes and *general* prizes to optimize such objectives, in the context of incomplete information.

In Chapter 5, we study contests where the designer's objective is an extension of the widely studied objective of maximizing the total output: The designer gets zero marginal utility from an agent's output if the output of the agent is very low or very high. We consider two variants of this setting, which correspond to two objective functions: *binary threshold*, where the designer's utility is a non-decreasing function

[2]In continuous-time BR dynamics, agents move smoothly from their current action to their BR action, i.e., take infinitesimally small steps towards BR, rather than jumping from their current action to their BR action.

of the number of agents with output above a certain threshold; and *linear threshold*, where an agent's contribution to the designer's utility is linear in her output if the output is between a lower and an upper threshold, and becomes constant below the lower and above the upper threshold. We characterize the contests that maximize the designer's objective and indicate techniques to efficiently compute them.

In Chapter 6, we study how to incentivize agents in a target subpopulation to produce a higher output by means of rank-order allocation contests. We describe a symmetric Bayes–Nash equilibrium for contests that have two types of rank-based prizes: (a) prizes that are accessible only to the agents in the target group; (b) prizes that are accessible to everyone. We also specialize this equilibrium characterization into two important sub-cases: (i) contests that do not discriminate while awarding the prizes, i.e., only have prizes that are accessible to everyone; (ii) contests that have prize quotas for the groups, and each group can compete only for prizes in their share. For these models, we characterize the properties of the contest that maximize the expected target-group total output.

1.2.3 Part III: Simultaneous Contests

An active area of research in contest theory is to understand simultaneous contests, i.e., when multiple contests run in parallel and the agents have to strategically allocate their limited resources across the contests. Applications include blockchains [18] and social media platforms [76].

In Chapter 7, we study a general class of simultaneous contests with equal-sharing allocation of prizes and multi-dimensional output by agents. We consider two variations of the model: (i) agents have costs for producing outputs; (ii) agents do not have costs but have generalized budget constraints. We observe that these games are exact potential games and hence always have a pure-strategy Nash equilibrium. The price of anarchy is 2 for the budget model but can be unbounded for the cost model. Our main results are for the computational complexity of these games. We prove that for general versions of the model exactly or approximately computing a best response is NP-hard. For natural restricted versions where the best response is easy to compute, we show that finding a pure-strategy Nash equilibrium is PLS-complete, and finding a mixed-strategy Nash equilibrium is CLS-complete (CLS=PPAD$\cap$PLS). On the other hand, an approximate pure-strategy Nash equilibrium can be found in pseudo-polynomial time. These games are a strict but natural subclass of explicit congestion games, but they still exhibit the same equilibrium computation hardness.

Part I

Dynamics in Tullock Contests

Chapter 2

Best-Response Dynamics

2.1 Introduction

We study the convergence of best-response (BR) dynamics in Tullock contests with convex costs (or concave utility functions). Convex costs capture situations where the marginal cost of producing output increases relative to the marginal value for doing so. Tullock contests with convex costs have received much more attention in the literature (compared to non-convex costs) partially because of their more widespread applications and partially because of the technical challenges in analyzing the models with non-convex costs [67].

Given an environment with strategic agents, such as a contest, it is desirable to be able to reliably predict the agents' behavior, so as to reason about possible outcomes. Nash equilibrium strategies serve as an initial approximation to this goal, but the existence of an equilibrium is not always a good predictor of the agents' behavior. Indeed, the traditional explanation of equilibrium is that it results from analysis and introspection by the agents in a situation where the rules of the game, the rationality of the agents, and the agents' payoff functions are all common knowledge. Both conceptually and empirically, these assumptions are not always satisfied in real-life scenarios [52]. A model provides a more robust prediction of the outcome of a game if it explains how the outcome can be attained in a decentralized manner, ideally via a process that involves agents responding to incentives provided by their environment. BR dynamics is arguably the simplest and the most intuitive of these models. In BR dynamics, agents sequentially update their current strategy with one that best responds to that of the other agents. It is especially appropriate in settings, such as Tullock contests with convex costs, where pure-strategy Nash equilibria are guaranteed to exist (see, e.g., [106, Chapter 4]).

In our model for BR dynamics, we assume that at each time step exactly one of the agents best responds and updates her current strategy. For two agents, this leads to a deterministic process where the agents play alternate moves.[1] For more than two agents, notice that the dynamics is under-specified because all agents except the agent who made the most recent transition can make a non-trivial move. So, we consider two models for (the agents themselves, intrinsically) selecting the agent who makes the BR move at each time step. First, we consider a randomized model where an agent i makes the BR transition with some probability $p_{t,i}(\boldsymbol{x}_t)$ at time t, given the strategy profile at time t is $\boldsymbol{x}_t$.[2] In the second model, we assume that the agents are selected to ensure fast convergence to equilibrium; this model is primarily used to provide convergence-rate lower bounds. We analyze the rate of convergence of the BR dynamics to an ϵ-approximate equilibrium in these models.

2.1.1 Results

For two homogeneous agents, we show that the BR dynamics converges to an ϵ-approximate pure-strategy Nash equilibrium in $\Theta(\log\log(1/\epsilon) + \log\log(1/\gamma))$ steps, where γ is a function of the initial action profile (generally equal to the smallest positive action in the initial profile). We also show faster convergence for restricted classes of convex cost functions.

For the randomized agent selection model, we show convergence in $O(\alpha \log(n/(\epsilon\delta)) + \beta \log\log(1/\gamma))$ steps with probability $1-\delta$, where α and β depend upon the randomized agent selection process (Theorem 2.9), e.g., $\alpha = n^2\log(n)$ and $\beta = 1$ if agents are selected uniformly at random each time step. We also provide a lower bound of $\Omega(\eta\log(n\delta) + \log(1/\epsilon)/\log(n) + \log\log(1/\gamma))$ with probability $1-\delta$, where η is again related to the randomized agent selection process, e.g., $\eta = n$ for uniform selection. These bounds are tight w.r.t. ϵ and γ, and are polynomially tight w.r.t. n. For the best-case agent selection model, we provide upper and lower bounds of $O(n\log(n/\epsilon)+\log\log(1/\gamma))$ and $\Omega(n+\log(1/\epsilon)/\log(n)+\log\log(1/\gamma))$, respectively.

Our analysis also implies convergence of the BR dynamics for any agent selection process that does not *starve* any agent, i.e., everyone gets to play infinitely often. To the best of our knowledge, a convergence result for the BR dynamics was not previously known for Tullock contests with homogeneous agents with general convex

[1] Two consecutive moves by the same agent is redundant because the agent is already best responding after the first move and the second move does not change the action.

[2] As exactly one agent moves each time step, so $\sum_i p_{t,i}(\boldsymbol{x}_t) = 1$ We make the technical assumption that $p_{t,i}(x_t) > 0$ for all agents except the one who played the most recent move. See Section 2.2.2 for more details.

costs. For non-homogeneous agents, we provide examples where the BR dynamics goes into a cycle. The set of agent types and initial states that lead to non-convergence has a positive measure.

Discounted-Sum Dynamics A novel contribution of this work is the introduction and analysis of the discounted-sum (DS) dynamics (Section 2.4), which is used in the analysis of the BR dynamics. This dynamics proceeds as follows: starting from an initial state $\boldsymbol{x}_0 \in \mathbb{R}^n$, at each time t, the dynamics picks a coordinate $i_t \in [n]$ and updates the value at the i_t-th coordinate with the negative discounted sum of the values at the remaining coordinates, i.e., $x_{t+1,i_t} = -\beta_t \sum_{j \neq i_t} x_{t,j}$. The discount factor $\beta_t \in [0, 1)$ can be picked arbitrarily, possibly adversarially, given t, i_t, and $\boldsymbol{x}_t$. We show that this dynamics rapidly converges to 0 using a potential function.

Intuition for the Proof of Convergence We divide the analysis into three main parts. The first part corresponds to a *warm-up* phase of the dynamics. Intuitively, there is a certain domain $\mathcal{D}$ such that the BR dynamics avoids corner cases, e.g., BR to 0, if the output profile is in $\mathcal{D}$. The warm-up phase corresponds to the time it takes to ensure that the output profile is inside $\mathcal{D}$ (where it stays thereafter) and is proportional to the time it takes for every agent to make at least one move. In the second part of the analysis, we derive a set of sufficient conditions that ensures that the BR dynamics behaves like the DS dynamics described earlier. We bound the time it takes for the BR dynamics to reach a state that satisfies these conditions, and show that the BR dynamics satisfies these conditions thereafter. In the third part of the analysis, we use our results for the DS dynamics to show convergence of the output profile to the equilibrium output profile in the BR dynamics, which implies convergence to an approximate equilibrium by showing that the utility function is Lipschitz continuous.

For readability, we defer certain proofs to Section 2.8.

2.1.2 Related Work

Ewerhart [44] shows that a lottery contest (Tullock contest with linear costs) with homogeneous agents is a BR potential game [107, 69], which is a strict generalization of the better known classes of ordinal and exact potential games [83]. Lottery contests are known not to be exact potential games even with homogeneous agents [44] and not to be ordinal potential games with non-homogeneous agents [45].

Moulin and Vial [84] implicitly show strategic equivalence between contests and zero-sum games. This directly implies convergence of fictitious play dynamics for two agents [47], but no such result has been proven for three or more agents. A Tullock contest corresponds to a Cournot game with isoelastic inverse demand and constant marginal costs. There are convergence results of learning dynamics for specific types of Cournot games, such as Cournot oligopoly with strictly declining best-response functions [38, 102], Cournot game with linear demand [100], aggregative games that allow monotone best-response selections [61, 42, 63], and others [41, 24]. However, all these methods do not apply to the Tullock contest whose best-response function is not monotone [40]. A different line of research has studied the convergence (or chaotic behavior) of learning dynamics in other types of contests (like all-pay auctions) and Cournot games (e.g., [90, 108, 32]), but these techniques and results also do not apply to Tullock contests.

2.2 Preliminaries

Let $\lg(x) = \log_2(x)$ and $\ln(x) = \log_e(x)$. Recall that in Tullock's model [104], n agents participate in a contest with unit prize (normalized). The agents simultaneously produce non-negative outputs; we denote the output of agent i by $x_i \in \mathbb{R}_{\geq 0}$ and the output profile by $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n_{\geq 0}$. Let $\boldsymbol{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$, $s = \sum_{j \in [n]} x_j$, and $s_{-i} = \sum_{j \neq i} x_j$. Agent i incurs a cost of $c_i(x_i)$ for producing the output x_i and receives a fraction of the prize proportional to x_i if at least one agent produces a strictly positive output, else $1/n$.[3] The utility of agent i, $u_i(\boldsymbol{x})$, is

$$u_i(\boldsymbol{x}) = \begin{cases} \frac{x_i}{\sum_j x_j} - c_i(x_i) = \frac{x_i}{x_i + s_{-i}} - c_i(x_i), & \text{if } \sum_j x_j > 0, \\ \frac{1}{n}, & \text{otherwise.} \end{cases} \tag{2.1}$$

Notice that the utility of agent i depends upon her output x_i and the total output of other agents s_{-i} but not upon the distribution of s_{-i} across the $n-1$ agents. To signify this, we shall denote $u_i(\boldsymbol{x})$ as $u_i(x_i, s_{-i})$.

We make the following assumptions on the cost functions: for every agent i, the cost function c_i is (a) twice differentiable; (b) zero cost for non-participation, $c_i(0) = 0$; (c) increasing, $c_i'(z) > 0$ for all $z \geq 0$; (d) weakly convex, $c_i''(z) \geq 0$ for all $z \geq 0$. These assumptions are standard in the literature, and without them, a pure-strategy Nash equilibrium may not exist (see, e.g., [106, Chapter 4], [46]). The

[3] Some papers in the literature, e.g., [34], assume that all agents get a prize of 0 if they all produce a 0 output. Our analysis and results remain the same with this alternate assumption as well.

twice differentiability assumption may be relaxed to arrive at similar results, but we assume this to provide a cleaner presentation. We also make an additional Lipschitz continuity assumption: $|c_i(z) - c_i(\overline{z})| \leq K|z - \overline{z}|$ for all $z, \overline{z}$ in the neighborhood of the equilibrium,[4] where K is the Lipschitz constant; this constant K is used in our convergence-rate bounds to ensure that if an output profile is close to the equilibrium output profile, then the output profile is an approximate equilibrium.

2.2.1 Best-Response Dynamics

Given the total output $s_{-i} = \sum_{j \neq i} x_j$ of all agents except i, the best response (BR) of agent i is an action x_i such that

$$x_i \in \arg\max_{z \geq 0} u_i(z, s_{-i}) = \arg\max_{z \geq 0} \left(\frac{z}{z + s_{-i}} - c_i(z) \right).$$

First, notice that an agent has no BR if the output produced by every other agent is 0.[5] To circumvent this technical issue, we assume that there is a small positive value $a < 1$ such that $x_i = a$ is the action of any agent i if $s_{-i} = 0$. We shall work with this assumption—an agent plays a if everyone else plays 0.[6]

On the other hand, agent i has a unique BR if $s_{-i} > 0$, i.e., if the output produced by at least one other agent j is non-zero. This unique BR can be computed by taking a derivative of $u_i(z, s_{-i})$ with respect to z. If $s_{-i} > 0$, then

$$\frac{\partial u_i(z, s_{-i})}{\partial z} = \frac{s_{-i}}{(z + s_{-i})^2} - c_i'(z), \tag{2.2}$$

$$\frac{\partial^2 u_i(z, s_{-i})}{\partial z^2} = \frac{-2s_{-i}}{(z + s_{-i})^3} - c_i''(z) < 0, \tag{2.3}$$

where the last inequality is true because $z \geq 0$, $c_i''(z) \geq 0$, and $s_{-i} > 0$. So, $u_i(z, s_{-i})$ is strictly concave w.r.t. z, i.e., $\frac{\partial u_i(z, s_{-i})}{\partial z}$ is strictly decreasing w.r.t. z. Let $BR_i(s_{-i})$

[4] For convergence to an ϵ-approximate equilibrium, we will assume Lipschitz continuity in a neighborhood of size proportional to ϵ.

[5] If $s_{-i} = 0$, then by producing an output of $\epsilon > 0$, agent i gets a utility of $u_i(\epsilon, 0) = 1 - c_i(\epsilon)$, which is strictly more than $u_i(0, 0) = 1/n$ for small ϵ (because c_i is continuous and $c_i(0) = 0$, which implies $c_i(\epsilon)$ is close to 0), so $x_i = 0$ cannot be a BR. Further, any $\epsilon > 0$ cannot be a BR because $u_i(\epsilon/2, 0) = 1 - c_i(\epsilon/2) > 1 - c_i(\epsilon) = u_i(\epsilon, 0)$ (and because c_i is increasing).

[6] This issue of not having any BR to 0 can also be resolved by an alternate technical assumption: the prize is given to agent i with probability $\frac{x_i}{b + \sum_j x_j}$ and agent i's expected utility is $\frac{x_i}{b + \sum_j x_j} - c_i(x_i)$, where b is a small positive constant. Notice that, with this alternate assumption, the prize may not get allocated to any agent (or is allocated to a pseudo agent) with a positive probability of $\frac{b}{b + \sum_j x_j}$, unlike our model. We *expect* all our results to hold for this alternate model as well.

denote the BR of agent i given that the total output of other agents is s_{-i}. The first-order conditions for the BR are

$$BR_i(s_{-i}) > 0 \text{ and } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=BR_i(s_{-i})} = 0, \text{ if } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=0} > 0; \tag{2.4}$$

$$BR_i(s_{-i}) = 0 \text{ and } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=BR_i(s_{-i})} \leq 0, \text{ if } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=0} \leq 0. \tag{2.5}$$

If $s_{-i} > 0$, notice that at $z = \max(1, c_i^{-1}(1))$, the agent has a cost of at least 1 but a prize of less than 1, therefore $\left.\frac{\partial u_i(z,s_{-i})}{\partial z}\right|_{z=\max(1,c_i^{-1}(1))} < 0$. So, the BR is less than $\max(1, c_i^{-1}(1))$. For $s_{-i} = 0$, the BR is $a < 1$ by assumption.

Slightly overloading notation, let $\boldsymbol{x}_t = (x_{t,i})_{i\in[n]}$ denote the action profile of the agents at time t in the BR dynamics. Similarly, let $s_t = \sum_j x_{t,j}$ and $s_{t,-i} = \sum_{j\neq i} x_{t,j}$. The BR dynamics starts with an initial profile $\boldsymbol{x}_0 = (x_{0,i})_{i\in[n]}$. At each time step $t \geq 0$, an agent $i_t \in [n]$ makes a BR move. Formally, $x_{t+1,i_t} = BR_{i_t}(s_{t,-i_t})$ and $x_{t+1,j} = x_{t,j}$ for $j \neq i_t$. We study the convergence (or non-convergence) of this BR dynamics.

2.2.2 Agent-Selection Models for Three or More Agents

When there are just two agents, we may assume that the BR dynamics proceeds uniquely: the two agents make alternate BR moves because making consecutive moves is redundant. But, when there are $n \geq 3$ agents, the BR dynamics is not unique. At any time step, the $n - 1 \geq 2$ agents who did not make the most recent BR move can possibly make a non-trivial move. The rate of convergence depends upon which agent makes a move at a given time step. We consider two models for selecting the agent playing the BR move at any given time step.

2.2.2.1 Randomized Agent-Selection Model

In the random selection model, agent i makes the BR move w.p. (with probability) $p_{t,i}(\boldsymbol{x}_t)$ at time t given output profile $\boldsymbol{x}_t$. The probability $p_{t,i}(\boldsymbol{x}_t)$ models the relative activity of agent i at time t given the current state $\boldsymbol{x}_t$. We assume that $p_{t,i}(\boldsymbol{x}_t) \geq L > 0$ for all agents except the agent who made the last transition. Our rate of convergence analysis will be worst-case over all $p_{t,i}(\boldsymbol{x}_t)$ values given the parameter $L > 0$. Even if this requirement on the $p_{t,i}(\boldsymbol{x}_t)$ values is not satisfied, our analysis still implies convergence of the BR dynamics assuming every agent gets to play infinitely often.

An important special case of the model is to assume that $p_{t,i}(\boldsymbol{x}_t) = 1/n$ for all t, i, and $\boldsymbol{x}_t$, i.e., the agent making the move is selected uniformly at random.

2.2.2.2 Best-Case Agent-Selection Model

In the best-case selection model, we assume that the selection process picks the sequence of agents making the BR moves to reach the ϵ-approximate equilibrium as quickly as possible. We provide lower bounds on the time till convergence for this model, which by definition automatically apply to all other selection models, including the randomized model. We also prove almost tight upper bounds for this model.

2.2.3 Equilibrium

A Tullock contest with convex costs always has a pure-strategy Nash equilibrium (which is also the unique equilibrium, including mixed-strategy Nash equilibria; see, e.g., [46]). So, we exclusively focus on pure equilibria.

Definition 2.1 (Pure-Strategy Nash Equilibrium)**.** An action profile $\boldsymbol{x}^* = (x_1^*, \ldots, x_n^*)$ is a pure-strategy Nash equilibrium if it satisfies

$$u_i(x_i^*, \boldsymbol{x}_{-i}^*) \geq u_i(x_i', \boldsymbol{x}_{-i}^*),$$

for every agent i and every action x_i' for agent i.

Definition 2.2 (Approximate Pure-Strategy Nash Equilibrium)**.** An action profile $\boldsymbol{x} = (x_1, \ldots, x_n)$ is an ϵ-approximate pure-strategy Nash equilibrium, for $\epsilon > 0$, if it satisfies

$$u_i(x_i, \boldsymbol{x}_{-i}) \geq (1-\epsilon) u_i(x_i', \boldsymbol{x}_{-i}),$$

for every agent i and every action x_i' for agent i.

In general, BR dynamics in a Tullock contest never exactly reaches the equilibrium, rather it may *converge* to the equilibrium. The dynamics converges to an equilibrium if it reaches an ϵ-approximate equilibrium in finite time for any $\epsilon > 0$.

2.2.4 Homogeneous Agents

The agents are homogeneous if they all have the same cost function. To keep our analysis cleaner, we make the following additional assumption on the (homogeneous) cost function: for every agent i, $c_i(x_i) = \frac{n-1}{n^2} c(x_i)$, where c satisfies $c'(1) = 1$. This assumption is w.l.o.g. (without loss of generality) by suitable change of notation.[7]

[7] Say the (homogeneous) cost of agent i is $c_i(x_i) = \tilde{c}(x_i)$ for some $\tilde{c}$, for every i. First, we can introduce the $\frac{n-1}{n^2}$ factor by writing $c_i(x_i) = \frac{n-1}{n^2}\widehat{c}(x_i)$, where $\widehat{c}(x_i) = \frac{n^2}{n-1}\tilde{c}(x_i)$. Then, let $\gamma > 0$ be the solution to $\widehat{c}'(\gamma) = \frac{1}{\gamma}$. Such a γ always exists because $\widehat{c}'$ is increasing, while $\frac{1}{\gamma}$ is decreasing and converges to ∞ as γ converges to 0. So, setting $c(x_i) = \widehat{c}(\gamma x_i)$, we get $c'(1) = \gamma\widehat{c}'(\gamma) = 1$. Finally, we can change the notation and say that an agent i produces an output of $y_i = \gamma x_i$ instead of x_i.

The utility of agent i can be written as $u_i(x_i, s_{-i}) = \frac{x_i}{x_i+s_{-i}} - \frac{n-1}{n^2}c(x_i)$ for every $i \in [n]$. Notice that the BR of each agent is now identical given the same action profile of the remaining agents; let us, therefore, denote the best response function as $BR = BR_i$ for every $i \in [n]$ by suppressing i. There is a unique equilibrium where each agent i plays the same action (see, e.g., [106, Chapter 4]), and given our notation, this is equal to 1. Indeed, if $x_j = 1$ for all j, then $s_{-i} = n-1$ and $\frac{\partial u_i(z,s_{-i})}{\partial z}\Big|_{z=1} = \frac{s_{-i}}{(1+s_{-i})^2} - \frac{n-1}{n^2}c'(1) = \frac{n-1}{(1+n-1)^2} - \frac{n-1}{n^2} = 0$.

2.3 Two Homogeneous Agents

In this section, we study the rate of convergence of best-response dynamics for two homogeneous agents. W.l.o.g. we assume that agent 1 makes the best-response move when t is odd, i.e., $i_t = 1$ if t odd, and agent 2 makes the best-response move when t is even, i.e., $i_t = 2$ if t even. Formally, for $t = 0, 1, 2, \ldots$, the BR dynamics proceeds as:

$$\begin{aligned} x_{t+1,2} &= BR(x_{t,1}), \text{ and } x_{t+1,1} = x_{t,1}, \qquad \text{if } t = 0, 2, 4, \ldots, \\ x_{t+1,1} &= BR(x_{t,2}), \text{ and } x_{t+1,2} = x_{t,2}, \qquad \text{if } t = 1, 3, 5, \ldots. \end{aligned}$$

Next, we prove a double-logarithmic rate of convergence of the BR dynamics w.r.t. the initial state and the approximation parameter.

Theorem 2.1. *Best-response dynamics in Tullock contests with two homogeneous agents reaches an ϵ-approximate equilibrium within $\lg\lg(\frac{1}{\epsilon}) + \lg\lg(\frac{1}{\gamma}) + O(1)$ steps, where γ is a function of the initial state: $\gamma = x_{0,1}$ if $0 < x_{0,1} < 1$; $\gamma = (c')^{-1}\left(\frac{1}{x_{0,1}}\right)$ if $x_{0,1} > 1$ and $c'(0) < \frac{4}{x_{0,1}}$; and $\gamma = a$ otherwise.*

Proof. Notice that the evolution of the states of the BR dynamics $(\boldsymbol{x}_t)_{t\geq 0}$ can be tracked by a single sequence $(z_t)_{t\geq 0}$ defined as follows:

$$z_t = \begin{cases} x_{t,1} & \text{if } t = 0, 2, 4, \ldots, \\ x_{t,2} & \text{if } t = 1, 3, 5, \ldots. \end{cases}$$

z_t has a one-to-one correspondence with $\boldsymbol{x}_t$ because: (i) Only the action of agent 1 in the initial profile is relevant for the BR dynamics, because, at $t = 0$, agent 2 makes the BR move, which is a function of only $x_{0,1}$, and which overwrites $x_{0,2}$. (ii) For each time step, one of the agents moves and the other stays at their current action, and z_t is equal to the new action at each time point.

We next prove the following properties about the sequence $(z_t)_{t\geq 0}$: (i) $z_t \leq 1$ for all $t \geq 1$; (ii) if $z_t < 1$, then $z_t < z_{t+1} < 1$. Technical proofs omitted for readability, including the proof of Lemma 2.2, are provided in Section 2.8.

Lemma 2.2. *The BR of agent i, $BR(x_{-i})$, given any strategy x_{-i} of the other agent, satisfies the following properties:*

1. *if $x_{-i} < 1$, then $x_{-i} < BR(x_{-i}) < 1$,*
2. *if $x_{-i} = 1$, then $BR(x_{-i}) = 1$, and*
3. *if $x_{-i} > 1$, then $BR(x_{-i}) < 1$.*

Dependency on the Initial State Let us now analyze the effect of the initial state z_0 on the BR dynamics. If $z_0 = 1$, from Lemma 2.2 (2), we know that $z_t = 1$ for all t. This can alternately be deduced from the fact that $(1, 1)$ is the equilibrium profile. For the remaining proof, we assume that $z_0 \neq 1$. As $z_0 \neq 1$, from Lemma 2.2 (1) and (3), we know that $z_1 < 1$. Again using Lemma 2.2 (1), we know that $z_t < z_{t+1} < 1$ for all $t \geq 1$, i.e., $(z_t)_{t\geq 1}$ is a strictly increasing sequence going towards 1. Let us define the variable γ used in the theorem statement as follows:

- If $z_0 = 0$, set $\gamma = a < 1$. Notice that $z_1 = a = \gamma$.
- If $0 < z_0 < 1$, set $\gamma = z_0 < 1$.
- If $z_0 > 1$ and $c'(0) < \frac{4}{z_0}$ set $\gamma = (c')^{-1}\left(\frac{1}{z_0}\right)$. From the first-order condition, equation (2.4), and using $0 \leq z_1 < 1 < z_0$, we have

$$\frac{z_0}{(z_0 + z_1)^2} = \frac{1}{4}c'(z_1) \implies \frac{z_0}{(z_0 + z_0)^2} = \frac{1}{4z_0} \leq \frac{1}{4}c'(z_1) \leq \frac{1}{z_0} = \frac{z_0}{(z_0 + 0)^2}$$
$$\implies z_1 = (c')^{-1}\left(\frac{\kappa}{z_0}\right) \text{ for some } \kappa \in [1, 4]. \quad (2.6)$$

  We will later prove that the dependency on γ is $\lg\lg(\frac{1}{\gamma}) + O(1)$, so we may ignore the constant κ as it can be absorbed in the $O(1)$ term.

- If $z_0 > 1$ and $c'(0) \geq \frac{4}{z_0}$, set $\gamma = a$. In this case, it is easy to check that $z_1 = 0$ and $z_2 = \gamma$.

By the definition of γ, either z_0, z_1, or z_2 is equal to $\gamma < 1$. So, for the remaining portion of the proof, by shifting the time by at most two steps, w.l.o.g., let us assume that $z_0 = \gamma < 1$. We know that the sequence $(z_t)_{t\geq 0}$ is strictly increasing. We next try to find the number of steps required by the sequence to increase from γ to $1-\epsilon$. We break this into two parts: the number of steps required to reach (i) $\frac{1}{2}$ from γ in Lemma 2.3; (ii) $1-\epsilon$ from $\frac{1}{2}$ in Lemma 2.4.

Lemma 2.3. *Given $z_0 = \gamma \in (0, \frac{1}{2}]$, $z_t \geq \frac{1}{2}$ for all $t \geq \lg\lg(\frac{1}{\gamma})$.*

Lemma 2.4. *Given $z_0 \in [\frac{1}{2}, 1]$, $z_t \geq 1-\epsilon$ for all $t \geq \lg\lg(\frac{1}{\epsilon}) + O(1)$.*

Combining the previous two lemmas, we know that after $t-1$ steps, where $t-1 \geq \lg\lg(\frac{1}{\epsilon}) + \lg\lg(\frac{1}{\gamma}) + O(1)$, we get $z_{t-1} \geq 1-\epsilon$ and $z_t \geq 1-\epsilon$. So, $x_{t,i} \geq 1-\epsilon$ for $i \in [2]$. Essentially, we have shown that the output profile $\boldsymbol{x}_t$ is close to the equilibrium output profile. Lemma 2.5 below completes the proof by showing that if the output profile is close to the equilibrium profile then the agents cannot benefit much by deviating.

Lemma 2.5. *If the output profile $\boldsymbol{x} = (x_1, x_2)$ satisfies $1-\epsilon \leq x_i \leq 1$ for $i \in [2]$, then $\boldsymbol{x}$ is a 3ϵ-approximate equilibrium.*

□

We complement the upper bound in Theorem 2.1 with a lower bound of $\lg\lg(\frac{1}{\epsilon}) + \lg\lg(\frac{1}{\gamma}) + \Omega(1)$ for linear cost functions (Theorem 2.6 with parameters $p = q = 0$). So, for the class of all convex cost functions, our bound is tight up to additive constants. But, faster convergence rates are possible for specific classes of convex cost functions. For example, if $c'(z) = z^q$ and $q \to \infty$, then $BR(x_{-i}) \to 1$ for any $x_{-i} \in [0,1]$ and $i \in [2]$. So, within a constant number of steps, we can approximately reach the equilibrium. In Theorem 2.6, we provide tight bounds for cost functions whose derivatives can be bounded as $z^p \leq c'(z) \leq z^q$ for $z \in [0,1]$ and $p \geq q \geq 0$. Notice that the case $p = q = 1/r$ for $r \in (0,1]$ is equivalent to Tullock contests with concave contest success functions $(\frac{x_i^r}{\sum_j x_j^r})$ and linear costs.

Theorem 2.6. *Best-response dynamics in Tullock contests with two homogeneous agents and a convex cost function c that satisfies $z^p \leq c'(z) \leq z^q$ for $z \in [0,1]$ and $p \geq q \geq 0$ reaches an ϵ-approximate equilibrium in $\lg\lg(\frac{1}{\epsilon}) + \frac{1}{\lg(2+r)}\lg\lg(\frac{1}{\gamma}) - \lg\lg(2+r) + \Theta(1)$ steps, where $r \in [q,p]$ and γ is a function of the initial state: $\gamma = x_{0,1}$ if $0 < x_{0,1} < 1$; $\gamma = (\frac{1}{x_{0,1}})^{1/r}$ if $x_{0,1} > 1$ and $c'(0) < \frac{4}{x_{0,1}}$; and $\gamma = a$ otherwise.*

2.4 Discounted-Sum Dynamics

In this section, we focus on a dynamic process we call the discounted-sum (DS) dynamics and study its convergence properties. We will use these results for our subsequent analysis of the BR dynamics for three or more agents. The DS dynamics, its convergence, and the potential function used to prove its convergence may be of independent interest.

Definition 2.3 (Discounted-Sum Dynamics)**.** Let $\boldsymbol{z}_0 = (z_{0,i})_{i\in[n]} \in \mathbb{R}^n$ be the initial state. At time $t \in \mathbb{Z}_{\geq 0}$, an $i_t \in [n]$ is selected, and the next state is computed as follows:

$$z_{t+1,i_t} = -\beta_{t,i}(\boldsymbol{z}_t)\sum_{j\neq i_t} z_{t,j}; \qquad z_{t+i,j} = z_{t,j}, \text{ for } j \neq i_t;$$

where $0 \leq \beta_{t,i}(\boldsymbol{z}_t) \leq B$ for some non-negative constant $B < 1$.

We assume that $\beta_{t,i}(\boldsymbol{z}_t)$ can be chosen arbitrarily (possibly adversarially), but must satisfy $0 \leq \beta_{t,i}(\boldsymbol{z}_t) \leq B$. For convenience, we refer to the n coordinates of the DS dynamics as agents, because these n coordinates will correspond to the n agents when we use this dynamics to analyze the BR dynamics in subsequent sections. We analyze the DS dynamics for two methods of selecting agent i_t—randomized and best-case—similarly to the agent selection processes of the BR dynamics as described in Section 2.2.2.

We study ϵ-approximate convergence of the DS dynamics w.r.t. the ℓ_1-distance, i.e., for any $\epsilon > 0$, for sufficiently large t, we want $||\boldsymbol{z}_t||_1 = \sum_j |z_{t,j}| \leq \epsilon$. It can be observed that the unique stable point of this dynamics is $\mathbf{0}$ (our rate of convergence analysis will also formally imply this claim). Let $\sigma_t = \sum_j z_{t,j}$ and $\sigma_{t,-i} = \sum_{j\neq i} z_{t,j}$.

Lemma 2.7. *For the randomized agent selection model, the DS dynamics ϵ-approximately converges to $\mathbf{0}$ in $O(\frac{\log(n)}{L^2(1-B)}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon\delta}))$ time w.p. $1-\delta$, i.e., $||\boldsymbol{z}_t||_1 \leq \epsilon$ for all $t \geq T$ and $T = O(\frac{\log(n)}{L^2(1-B)}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon\delta}))$ w.p. $1-\delta$.*

Before we prove Lemma 2.7, let us briefly discuss the behavior of the DS dynamics. Say agent $i = i_t$ makes the move at time t, then the dynamics updates the i-th coordinate as $z_{t+1,i} = -\beta_{t,i}(\boldsymbol{z}_t)\sigma_{t,-i}$, where $\beta_{t,i}(\boldsymbol{z}_t)$ may be selected adversarially in $[0, B]$. Let us make the following simplifying assumption: say $\beta_{t,i}(\boldsymbol{z}_t) = B$ always, so $z_{t+1,i} = -B\sigma_{t,-i}$. With this simplifying assumption, we can construct a potential function $f_{simple}(\boldsymbol{z}_t) = \frac{B}{2}\sigma_t^2 + \frac{1-B}{2}||\boldsymbol{z}_t||_2^2$ and prove fast convergence in $O(\frac{1}{L}\log(\frac{n}{\epsilon\delta}))$ time with probability $1-\delta$ using techniques similar to the ones used for analyzing linear cost functions in Chapter 3. In particular, we can show that f_{simple} is $(1-B)$-strongly convex and $(1+(n-1)B)$-smooth and then use techniques used to analyze coordinate descent [110] to get the required convergence rate. But for arbitrary $\beta_{t,i}(\boldsymbol{z}_t) \in [0, B]$ coefficients, our attempts at constructing smooth and well-behaved potential functions failed (e.g., we can formally show that no degree-2 polynomial can be used as a potential function by constructing suitable examples; this observation likely holds for higher degree polynomials as well). Therefore, we introduce a novel non-smooth *weak* potential function and use it to prove the rate of convergence. The potential either decreases or stays the same (therefore, we call it weak) as the dynamics proceeds.

Proof of Lemma 2.7, partial. Given an n-dimensional vector $\boldsymbol{z} \in \mathbb{R}^n$, we define the following potential function

$$f(\boldsymbol{z}) = \max\left(\sum_{j\in[n]} \mathbb{1}(z_j > 0)z_j, -\sum_{j\in[n]} \mathbb{1}(z_j < 0)z_j\right), \tag{2.7}$$

where $\mathbb{1}$ is the indicator function. The potential function f separates the positive and negative parts of $\boldsymbol{z}$ and takes the max of the sums of their absolute values. Notice that $f(\boldsymbol{z}) \geq 0$ for all $\boldsymbol{z} \in \mathbb{R}^n$ and $f(\boldsymbol{z}) = 0 \iff \boldsymbol{z} = \mathbf{0}$. The intuition for this potential function is as follows: In the dynamics, we sum up all the z_j values except along one coordinate, say z_i, and plug-in a negative discounted value of this sum into z_i. In this sum, either the sum of the positive terms (left side of f) or the negative terms (right side of f) will have a higher magnitude, and after replacing z_i with the negative discounted-sum, we shall observe that the new left and right sides of f do not exceed the larger of the two old sides of f.

Let $v_{t,j} = \mathbb{1}(z_{t,j} > 0)z_{t,j}$ and $w_{t,j} = -\mathbb{1}(z_{t,j} \leq 0)z_{t,j}$. Let $V_t = \sum_j v_{t,j}$ and $W_t = \sum_j w_{t,j}$. Let also $\mathcal{V}_t = \{j \in [n] \mid z_{t,j} > 0\}$ and $\mathcal{W}_t = \{j \in [n] \mid z_{t,j} \leq 0\}$. Notice that $\mathcal{V}_t \cap \mathcal{W}_t = \emptyset$ and $\mathcal{V}_t \cap \mathcal{W}_t = [n]$. Also notice that $f(\boldsymbol{z}_t) = \max(V_t, W_t)$.

First, let us prove that the potential never increases. Let us assume $V_t \geq W_t$; this is w.l.o.g. because the dynamics and the potential function are symmetric w.r.t. the positive and the negative parts of $\boldsymbol{z}_t$.[8] Say agent i is picked at time t, we have the following cases:

- $i \in \mathcal{V}_t$. Depending upon the value of $v_{t,i}$, we have the following two sub-cases:

[8]An agent j with $z_{t,j} = 0$ can be put on either the positive or the negative side; as a convention, we associate them to the negative side, which causes slight asymmetry. It can be checked that our proof works irrespective of the side we assign to the agents with $z_{t,j} = 0$.

– $W_t \leq V_t - v_{t,i}$. Then agent i will change sign after this move: $v_{t+1,i} = 0$ and $w_{t+1,i} = \beta_{t,i}(\boldsymbol{z}_t)(V_t - v_{t,i} - W_t) \leq B(V_t - v_{t,i} - W_t)$. So, the updated value of the potential is

$$\begin{aligned} f(\boldsymbol{z}_{t+1}) &= \max(V_t - v_{t,i}, W_t + w_{t+1,i}) \\ &\leq \max(V_t - v_{t,i}, (1-B)W_t + B(V_t - v_{t,i})) \\ &= V_t - v_{t,i} \leq V_t = \max(V_t, W_t) = f(\boldsymbol{z}_t). \end{aligned}$$

– $W_t > V_t - v_{t,i}$. Then agent i will keep the same sign after the move: $v_{t+1,i} \leq B(W_t - (V_t - v_{t,i}))$. So, the updated value of the potential is

$$\begin{aligned} f(\boldsymbol{z}_{t+1}) &= \max(V_t - v_{t,i} + v_{t+1,i}, W_t) \\ &\leq \max((1-B)(V_t - v_{t,i}) + BW_t, W_t) \\ &= W_t \leq \max(V_t, W_t) = f(\boldsymbol{z}_t). \end{aligned}$$

- $i \in \mathcal{W}_t$. As $W_t \leq V_t$, the agent i will keep the same sign after the move: $w_{t+1,i} \leq B(V_t - (W_t - w_{t,i}))$. So, the updated value of the potential is

$$\begin{aligned} f(\boldsymbol{z}_{t+1}) &= \max(V_t, W_t - w_{t,i} + w_{t+1,i}) \\ &\leq \max(V_t, BV_t + (1-B)(W_t - w_{t,i})) \\ &= V_t = f(\boldsymbol{z}_t). \end{aligned}$$

We provide the rate of convergence analysis in Section 2.8. □

Lemma 2.7 proves a high probability bound of $O(\frac{\log(n)}{L^2})$ on the time till convergence (w.r.t. n and L). We think that this bound can be improved to $O(\frac{\log(n)}{L})$ using alternate techniques. However, this improvement is unlikely using the potential function f given in equation (2.7) and by measuring progress w.r.t. f only for small time periods. This is illustrated in Example 2.4.

Example 2.4. Let us assume that $\beta_{t,i}(\boldsymbol{z}_t) = \frac{1}{2}$ for all t, i, and $\boldsymbol{z}_t$.[9] Let $\boldsymbol{z}_t = (-1, 1, 0, \ldots)$. Note that $f(\boldsymbol{z}_t) = 1$. At time t, if we select an agent $i_t \geq 3$, then i_t has $z_{t,i_t} = 0$ by definition and $z_{t+1,i_t} = 0$ because $\sum_{j \neq i_t} z_{t,j} = 1 - 1 = 0$. So, $\boldsymbol{z}_{t+1} = \boldsymbol{z}_t$, and we make no progress w.r.t. the potential f. To make progress, we need to pick $i_t \in \{1, 2\}$, which, in the worst case, occurs with a probability of $2L$. So, it easy to check that we will need about $\frac{1}{2L}$ time steps before we pick an agent in $\{1, 2\}$.

[9] This ensures a $O(\frac{\log(n)}{L})$ rate of convergence using the smooth potential function f_{simple} discussed earlier in the section.

Now, after we pick $i_\tau \in \{1, 2\}$, say $i_\tau = 1$ w.l.o.g. at time $\tau \geq t$, we have $\boldsymbol{z}_{\tau+1} = (-\frac{1}{2}, 1, 0, \ldots, 0)$. The potential is still $f(\boldsymbol{z}_{\tau+1}) = 1$. At time $\tau + 1$, if we pick $i_{\tau+1} = 2$, which happens with probability L in the worst case, then $f(\boldsymbol{z}_{\tau+2}) = \frac{1}{2}$. But if we pick $i_{\tau+1} \geq 3$, say $i_{\tau+1} = 3$ w.l.o.g., then $\boldsymbol{z}_{\tau+2} = (-\frac{1}{2}, 1, -\frac{1}{4}, 0, \ldots, 0)$ and $f(\boldsymbol{z}_{\tau+2}) = 1$. And then if we pick $i_{\tau+2} \in \{1, 2\}$, then $f(\boldsymbol{z}_{\tau+3}) \geq \frac{3}{4}$. Extending this argument, if we pick an agent in $\{1, 2\}$ the first time after time τ at time $\tau + 3$, then $f(\boldsymbol{z}_{\tau+4}) \geq \frac{7}{8}$. Essentially, the longer the delay in picking an agent in $\{1, 2\}$, the less progress we make. By simple calculations, we can check that at time $\tau' > \tau$ when we eventually pick an agent in $\{1, 2\}$ the second time, the expected value of the potential is $\mathbb{E}[f(\boldsymbol{z}_{\tau'+1})] \geq 1 - 2L$.

To summarize, after about $\frac{2}{L}$ steps, we decreased the potential by a multiplicative factor of about $1 - 2L$. So, if we measure only progress w.r.t. f for small contiguous time periods, and we do not track any other properties of the state $\boldsymbol{z}_t$, then it is unlikely that we can get a convergence time better than $O(\frac{1}{L^2})$.

Notice that Example 2.4 applies to the special case of uniform distribution, i.e., $L = 1/n$. This indicates that our method is unlikely to provide a better bound even for the case of a fixed distribution that does not vary over time. Techniques from linear dynamical systems may help get a better bound. The discounted-sum dynamics is an instance of a time-varying linear dynamical system, with a specific structure for the matrices defining the linear transformation in each time step. One can study the product of the time-varying matrices of the linear dynamical system. The most natural idea would be to bound the spectral norm of these matrices (or some other suitable matrix norm). But the spectral norm for the individual matrices can be strictly larger than 1 (the same result holds for similar norms). So, we will need other ideas to show the convergence using this method.

We next prove almost tight upper and lower bounds for the convergence time for the best-case selection model. By the definition of the best-case model, the lower bound applies to the randomized model as well, which implies that our bound is tight w.r.t. ϵ and polynomially tight w.r.t. n.

Lemma 2.8. *For the best-case selection model, the DS dynamics ϵ-approximately converges in $O(\frac{n}{1-B}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon}))$ time, i.e., $||\boldsymbol{z}_t||_1 \leq \epsilon$ for all $t \geq T = O(\frac{n}{1-B}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon}))$. Further, for the best-case selection model, the dynamics may take $\Omega(n + \frac{1}{1-B}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon}))$ time to ϵ-approximately converge, i.e., $||\boldsymbol{z}_t||_1 > \epsilon$ for all $t < T = \Omega(n + \frac{1}{1-B}\log(\frac{f(\boldsymbol{z}_0)}{\epsilon}))$ for some starting points $\boldsymbol{z}_0$.*

2.5 Three or More Homogeneous Agents

We now study BR dynamics for $n \geq 3$ agents. As introduced in Section 2.2.2, we consider a randomized model and a best-case deterministic model for selecting the agent who makes the transition at any given time point t.

For the randomized model, agent i makes the transition at time t w.p. $p_{t,i}(\boldsymbol{x}_t)$ given that the action profile at t is $\boldsymbol{x}_t$. We assume $p_{t,i}(\boldsymbol{x}_t) \geq L > 0$ for all agents i except the one who played at time $t-1$. An important special case of our model is to assume that $p_{t,i}(\boldsymbol{x}_t) = 1/n$ for all t, i, and $\boldsymbol{x}_t$, i.e., the agent making the move is selected uniformly at random. We do a worst-case analysis over all $p_{t,i}$ given the parameter $L > 0$.

For the best-case deterministic model, we assume that the agents are selected to ensure the fastest possible convergence. We provide a lower bound on convergence time for this model, which automatically carries over to the randomized model.

Theorem 2.9. *BR dynamics in Tullock contests with $n \geq 3$ homogeneous agents with convex costs and randomized selection of agents reaches an ϵ-approximate equilibrium w.p. $1-\delta$ in $O(\frac{1}{nL}\log\log(\frac{1}{\gamma}) + \frac{\log(n)}{L^2}\log(\frac{nK}{\epsilon\delta}))$ steps for every $\epsilon, \delta \in (0,1)$, where γ is a function of the initial output profile as given below (and discussed in Lemma 2.16), and assuming that the cost function c is K-Lipschitz continuous in the interval $[1-\epsilon, 1+\epsilon]$.*

$\gamma = \min(\gamma_{lb}, \frac{1}{n-1})$, where $\gamma_{lb} = \min(\{a\} \cup \mathcal{A} \cup \mathcal{B})$, where $\mathcal{A} = \{x_{0,j} \mid x_{0,j} > 0, j \in [n]\}$ and $\mathcal{B}$ is either (1) $\mathcal{B} = \{BR(x_{0,j}+1) \mid j \in [n]\}$ if $c'(0) = 0$ or (2) $\mathcal{B} = \{\min(\frac{\kappa - x_{0,j}}{4}, BR(\frac{\kappa + x_{0,j}}{4})) \mid x_{0,j} < \kappa, j \in [n]\}$ if $c'(0) > 0$, where $\kappa = \frac{n^2}{c'(0)(n-1)}$.

Essentially, the value of γ above corresponds to the smallest positive output at the first time point when no agent produces an excessively large output and at least one agent produces a positive output.

Notice that the above theorem implies that when agents are selected uniformly at random at each time step (i.e., $L = 1/n$), the time till convergence is bounded above by $O(\log\log(\frac{1}{\gamma}) + n^2\log(n)\log(\frac{nK}{\epsilon\delta}))$ w.p. $1-\delta$.

Proof of Theorem 2.9. We shall use an extension of the well-known coupon collector problem in the first part of our analysis. In the coupon collector problem, we select a coupon out of n coupons uniformly at random at each time step, and we want to bound the time it takes to collect all coupons. A tight $\Theta(n \log n)$ high probability bound is known for this problem (see, e.g., [80]), which can be used to derive the following bounds on the time it takes for each agent to play at least once in the BR dynamics.

Lemma 2.10. *Let T be the time it takes for all agents to play at least once in the BR dynamics, we have the following high probability bounds: (i) upper bound, $\mathbb{P}[T \leq \frac{1}{L}\ln(\frac{n}{\delta})] \geq 1-\delta$; (ii) lower bound, $\mathbb{P}[T \geq \frac{1}{L}\log(n\delta)] \geq 1-\delta$.*

Keeping the same notation as our previous discussions, let $x_{t,i}$ be the output of agent $i \in [n]$ at time $t \geq 0$. Note that $x_{t,i}$ is a random variable and $(x_t)_{t\geq 0}$ is a stochastic process because the agent i_t making the move at time t is selected randomly. As before, let $s_t = \sum_j x_{t,j}$ and $s_{t,-i} = \sum_{j\neq i} x_{t,j}$; s_t and $s_{t,-i}$ are also random variables.

We call an initial time period of the BR dynamics the *warm-up* phase (Definition 2.5). When the warm-up phase ends, we ensure that the output profile is, and stays thereafter, in a certain well-behaved region where we can avoid the corner case of BR to 0.

Definition 2.5 (Warm-Up Phase)**.** The time period $\{0, 1, \ldots, T_{warm}-1\}$ denotes the warm-up phase, where T_{warm} is the smallest time such that for every $t \geq T_{warm}$:

1. the total output is strictly less than $\frac{n^2}{(n-1)c'(0)}$: $s_t < \frac{n^2}{(n-1)c'(0)}$;
2. all agents produce output of at most $\frac{n^2}{4(n-1)}$: $0 \leq x_{t,i} \leq \frac{n^2}{4(n-1)}$ for every i;
3. there are at least two agents i and $j \neq i$ with positive output: $x_{t,i} > 0$ and $x_{t,j} > 0$.

In the first part of the analysis, we bound the time it takes for the warm-up phase to finish with high probability.

Analysis Part 1 (Warm-Up Phase) We shall use the following properties (Lemma 2.11) of BR dynamics in our subsequent analysis.

Lemma 2.11. *The BR dynamics satisfies the following properties:*

1. *$s_t > 0$ for every $t \geq 1$.*
2. *$s_t < \frac{n^2}{(n-1)c'(0)} \implies s_{t+1} < \frac{n^2}{(n-1)c'(0)}$ for every $t \geq 0$.*
3. *For $i, j \in [n]$ such that $i \neq j$, if $x_{t,i} > 0$, $x_{t,j} > 0$, and $s_t < \frac{n^2}{(n-1)c'(0)}$, then $x_{t+1,i_t} > 0$ for every $t \geq 0$.*

The next lemma gives a high probability bound on the time taken by the warm-up phase.

Lemma 2.12. *Time till completion of the warm-up phase $T_{warm} = O(\frac{1}{L}\log(\frac{n}{\delta}))$ w.p. $1-\delta$.*

Analysis Part 2 (Reduction to Discounted-Sum Dynamics) We will now reduce the BR dynamics to the discounted-sum dynamics discussed in Section 2.4. We assume that all conditions for the completion of the warm-up phase (Definition 2.5) are satisfied.

Let $z_{t,i} = x_{t,i} - 1$. We will show that under certain conditions, $z_{t,i}$ follows the discounted-sum dynamics with suitably selected parameter $B \in [0, 1)$ (used in Definition 2.3). In particular, we need to show that $z_{t+1,i} = -\beta_{t,i}(\boldsymbol{z}_t) \sum_{j \neq i} z_{t,j}$ for some suitable $\beta_{t,i}(\boldsymbol{z}_t) \in [0, B]$ for all t, i, $\boldsymbol{z}_t$. Let $\sigma_t = \sum_j z_{t,j} = s_t - n$ and $\sigma_{t,-i} = \sum_{j \neq i} z_{t,j} = s_{t,-i} - (n-1)$.

Lemma 2.13. *If $s_{t,-i_t} \geq \frac{1}{n-1}$, then $z_{t+1,i_t} = x_{t+1,i_t} - 1 = -\beta_{t,i}(\boldsymbol{x}_t)(s_{t,-i_t} - (n-1)) = -\beta_{t,i}(\boldsymbol{x}_t)\sigma_{t,-i_t}$, where $\beta_{t,i}(\boldsymbol{x}_t) \in [0, B]$ and $B \leq \frac{1}{2}$.*

Lemma 2.13 proves that the BR dynamics behaves like the discounted-sum dynamics if $s_{t,-i_t} \geq \frac{1}{n-1}$. In the next few lemmas, we try to bound the time it takes to satisfy this condition with high probability.

Lemma 2.14. *(Assuming the completion of the warm-up phase.) If at time t, the output profile $\boldsymbol{x}_t$ satisfies the following condition, then $\boldsymbol{x}_\tau$ also satisfies it for all $\tau \geq t$. Condition: There are at least two agents i and $j \neq i$ such that $x_{t,i} \geq \frac{1}{n-1}$ and $x_{t,j} \geq \frac{1}{n-1}$.*

Notice that, if the condition in Lemma 2.14 is satisfied by $\boldsymbol{x}_t$, then there are two agents i and j with $x_{t,i} \geq \frac{1}{n-1}$ and $x_{t,j} \geq \frac{1}{n-1}$, which implies that for every $k \in [n]$, $s_{t,-k} \geq \min(x_{t,i}, x_{t,j}) \geq \frac{1}{n-1}$, and Lemma 2.13 ensures that the BR dynamics resembles a discounted-sum dynamics. We next bound the time it takes to satisfy the condition in Lemma 2.14.

Lemma 2.15. *For every $t \geq T_{warm} + T_{sum}$, there are at least two agents i and $j \neq i$ such that $x_{t,i} \geq \frac{1}{n-1}$ and $x_{t,j} \geq \frac{1}{n-1}$, where T_{sum} is defined as:*

- $T_{sum} = O(\frac{1}{nL} \ln(\frac{1}{\delta}))$ *w.p.* $1 - \delta$ *if* $\gamma = \min_{j \in [n]} s_{T_{warm},-j} \geq \frac{1}{n-1}$;
- $T_{sum} = O(\frac{1}{n^2 L^2} \ln(\frac{1}{\delta}) + \frac{1}{nL} \log\log(\frac{1}{\gamma}))$ *w.p.* $1 - \delta$ *if* $\gamma = \min_{j \in [n]} s_{T_{warm},-j} < \frac{1}{n-1}$.

Notice that the time duration T_{sum} given in Lemma 2.15 depends upon $\gamma = \min_{j \in [n]} s_{T_{warm},-j}$ when $\gamma < \frac{1}{n-1}$. In Lemma 2.15, γ is defined as a function of the output profile $\boldsymbol{x}_{T_{warm}}$ at time T_{warm}, which may be considered unsatisfactory, so we next try to lower bound γ as a function of the initial output profile $\boldsymbol{x}_0$.

Lemma 2.16. *Let* $\gamma = \min_{j\in[n]} s_{T_{warm},-j}$*. If* $\gamma < \frac{1}{n-1}$*, then it is lower bounded as* $\gamma \geq \gamma_{lb} = \min(\{a\} \cup \mathcal{A} \cup \mathcal{B})$*, where* $\mathcal{A} = \{x_{0,j} \mid x_{0,j} > 0, j \in [n]\}$ *and* $\mathcal{B}$ *is either (1)* $\mathcal{B} = \{BR(x_{0,j} + 1) \mid j \in [n]\}$ *if* $c'(0) = 0$ *or (2)* $\mathcal{B} = \{\min(\frac{\kappa - x_{0,j}}{4}, BR(\frac{\kappa + x_{0,j}}{4})) \mid x_{0,j} < \kappa, j \in [n]\}$ *if* $c'(0) > 0$*, where* $\kappa = \frac{n^2}{c'(0)(n-1)}$*.*

Lemma 2.16 intuitively says that the value of γ at the end of the warm-up phase is lower bounded by either a (the response of an agent to 0), or the smallest positive output by a player at time $t = 0$, or is close to the smallest positive BR to the largest output (that has a positive BR) by any player at time $t = 0$.

Let us summarize our results till now. In Lemma 2.12, we showed that the time taken by the warm-up phase is $O(\frac{1}{L}\log(\frac{n}{\delta'}))$ w.p. $1 - \delta'$. Then, assuming the completion of the warm-up phase, we have a $1 - \delta'$ high-probability bound of $O(\frac{1}{n^2L^2}\ln(\frac{1}{\delta'}) + \frac{1}{nL}\log\log(\frac{1}{\gamma}))$ on the time taken to reach a phase of the BR dynamics that corresponds to the discounted-sum dynamics (Lemma 2.14).

Next, we use our results in Section 2.4 to get a high-probability bound on the time it takes for the output profile to reach close to the equilibrium profile. In particular, as $z_{t,i} = x_{t,i} - 1 \geq -1$ and $z_{t,i} = x_{t,i} - 1 \leq \frac{n^2}{4(n-1)}$ for all $t \geq T_{warm}$, so $\sum_i |z_{t,i}| \leq \frac{n^3}{4(n-1)} \leq n^2$. Using Lemma 2.7, we get $||\boldsymbol{z}_{\tau+T_{warm}+T_{sum}}||_1 \leq \epsilon$ for all $\tau \geq T = O(\frac{\log(n)}{L^2}\log(\frac{n}{\epsilon\delta'}))$ w.p. $1-\delta'$. Finally, using Lemma 2.17 below, we show that small ℓ_1-distance implies approximate equilibria. Setting $\delta' = \delta/3$ and using union bound on the total probability of failure completes our analysis.

Lemma 2.17. *Given an output profile* $\boldsymbol{x} = (x_i)_{i\in[n]}$ *and the equilibrium profile* $\boldsymbol{x}^* = (1, \ldots, 1)$*, if* $||\boldsymbol{x} - \boldsymbol{x}^*||_1 \leq \epsilon$*,* $|BR(s_{-i}) - 1| \leq \epsilon$ *for all* $i \in [n]$*, and the (homogeneous) cost function* c *is* K*-Lipschitz continuous in the interval* $[1-\epsilon, 1+\epsilon]$*, then* $\boldsymbol{x}$ *is an* $4Kn\epsilon$*-approximate equilibrium for* $\epsilon \leq \frac{1}{4Kn}$*.*

□

We next prove almost tight upper and lower bounds for the time required for the convergence of the best-case selection model.

Theorem 2.18. *BR dynamics in Tullock contests with* $n \geq 3$ *homogeneous agents with convex costs and best-case selection of agents reaches an* ϵ*-approximate equilibrium in* $O(\log\log(\frac{1}{\gamma}) + n\log(\frac{nK}{\epsilon}))$ *steps for every* $\epsilon \in (0, 1)$*, where* γ *is a function of the initial output profile as given in Lemma 2.16, and assuming that the cost function* c *is* K*-Lipschitz continuous in the interval* $[1-\epsilon, 1+\epsilon]$*. As a lower bound, we show that the convergence takes at least* $\Omega(n + \log\log(\frac{1}{\gamma}) + \frac{1}{\log(n)}\log(\frac{1}{\epsilon}))$ *steps.*

The lower bound for the best-case selection model automatically applies to the randomized selection model; the dependency of $\Omega(n)$ given in Theorem 2.18 can be slightly improved to $\Omega(\frac{1}{L}\log(n\delta))$ w.p. $1-\delta$ using Lemma 2.10. Note that these bounds are in the worst case over all convex cost functions. In contrast, for every ϵ and γ, there exists a convex cost function that reaches an approximate equilibrium after just one BR transition by each agent. E.g., if $c'(z) = z^r$ for all $z \geq 0$ and $r \to \infty$, then $BR(s_{-i}) \to 1$ for any i and $s_{-i} < \infty$. The lower bounds w.r.t. γ and ϵ in Theorem 2.18 use the linear cost function: $c'(z) = 1$ for all $z \geq 0$. We defer tighter analyses of lower and upper bounds for specific classes of convex cost functions, as done for the two-agent case (Theorem 2.6), for future work.

2.6 Non-Homogeneous Agents

In this section, we provide examples with two non-homogeneous agents to demonstrate that the BR dynamics may not converge.[10] These examples also imply non-convergence for three or more non-homogeneous agents. In the examples below, let the cost of agent 1 be 1 (normalized) and agent 2 be $c \leq 1$ (w.l.o.g.). Notice that agent 2 has a lower cost and is *stronger* than agent 1.

Example 2.6. Let $c = 1/10$. Let $a = 10^{-5}$. The agents start from the initial profile $x_0 = (0, a)$, which leads to the cycle given in Table 2.1.

$x_{t,1}$	0.00000	0.00315	0.00315	0.24321	0.24321	0.00000	0.00000	...
$x_{t,2}$	0.00001	0.00001	0.17439	0.17439	1.31631	1.31631	0.00001	...

Table 2.1: Non-convergence of BR dynamics for non-homogeneous agents (Example 2.6).

The intuition for the cycle in Example 2.6 is that the stronger (lower cost $c = 1/10$) agent's best response may overshoot the upper limit on the output of the weaker agent (cost 1), which makes the weaker agent best respond with 0. Then, the stronger agent plays a, then the weaker agent plays $\sqrt{a} - a$, and so on, and the output keeps on increasing until the stronger agent overshoots the upper limit of the weaker agent again. Such an issue never occurs if the agents are homogeneous.

We know that when $c = 1$, i.e., the two agents are homogeneous, we have convergence to the equilibrium. On the other hand, Example 2.6 gave a cycle for $c = 1/10$. So, a natural question is: How non-homogeneous should the agents be to get a cycle?

[10]Link to code to generate these examples: `https://github.com/abheekg/br-tullock-convex`

Example 2.7 below shows that if $c \leq 4/25 = 1/6.25$, then we can get a cycle. For $c \in (4/25, 1)$, our simulations show convergence. We leave the formal analysis of the dynamics with almost homogeneous agents as an open problem.

Example 2.7. Let $c = 4/25$. Notice that, if $a = 1/4$, we get the following cycle:

$$(0, 1/4) \longrightarrow (1/4, 1/4) \longrightarrow (1/4, 1) \longrightarrow (0, 1) \longrightarrow (0, 1/4) \longrightarrow \dots.$$

We next show that such cycles can also be constructed for certain arbitrarily small values of a. Notice that we can write a *reverse* BR dynamics as follows: If agent 1 makes the move at time t, then we have $x_{t+1,1} = \sqrt{x_{t,2}} - x_{t,2}$, which can be reversed as $x_{t,2} = \frac{(1 \pm \sqrt{1-4x_{t+1,1}})^2}{4}$ (notice that there can be two possible values of $x_{t,2}$, indicated by the $\pm$ sign, that lead to the same $x_{t+1,1}$; we will focus on the smaller $x_{t,2}$ in this example). Similarly, if agent 2 made the move at time t, then we have $x_{t+1,2} = \sqrt{\frac{x_{t,1}}{c}} - x_{t,1}$, which can be reversed as $x_{t,1} = \frac{(1 \pm \sqrt{1-4cx_{t+1,2}})^2}{4c}$. So, starting from $(1/4, 1/4)$ and going backward, we can have the following reversed sequence going to 0:

$$\left(\frac{1}{4}, \frac{1}{4}\right) \longleftarrow \left(\approx \frac{1.1}{100}, \frac{1}{4}\right) \longleftarrow \left(\approx \frac{1.1}{100}, \approx \frac{1.2}{10^4}\right) \longleftarrow \left(\approx \frac{2.3}{10^9}, \approx \frac{1.2}{10^4}\right) \longleftarrow \dots.$$

By selecting one of the $x_{t,2}$ values as a in the (reversed) sequence above, we can construct a cycle.

If $c < 4/25$, we can extend the examples above to have a set of values of a that has a strictly positive measure in $[0, b]$, for any $b > 0$, and that leads to a cycle. See Example 2.8. Note that the set of values of a that cause a cycle can never cover the entire domain because there is always the equilibrium point $\left(\frac{c}{(1+c)^2}, \frac{1}{(1+c)^2}\right)$ (and, likely, a neighborhood of this point of positive measure).

Example 2.8. Let $c = 1/100$. Starting from $x_{t,2} \geq 1$, we track the reverse BR dynamics (as done in Example 2.7) as follows: $x_{t-1,1} \in [0.0098, 1)$, $x_{t-2,2} \in [9.80/10^5, 0.9803]$, $x_{t-3,1} \in [9.61/10^{11}, 0.0094]$, and so on. If we set a equal to a point in the intervals corresponding to $x_{\tau,2}$, for some τ, we will get a cycle.

All the examples above use linear cost functions and have cycles that pass through a step where an agent has to best respond to 0. As the BR to 0 is not well defined, we may wonder whether there is an example that does not pass through a BR to 0. Also, we may ask whether cycles are possible for strictly convex cost functions. Example 2.9 is such an example that uses strictly convex cost functions and does not pass through a BR to 0.

Example 2.9. Let $c_1(z) = z^{1.1}$ and $c_2(z) = z^{1.1}/20$ for $z \geq 0$. BR dynamics leads to the cycle given in Table 2.2.

$x_{t,1}$	0.00102	0.24588	0.24588	0.00102	0.00102	0.24588	...
$x_{t,2}$	0.14930	0.14930	1.80686	1.80686	0.14930	0.14930	...

Table 2.2: Non-convergence of BR dynamics for non-homogeneous agents with strictly convex costs (Example 2.9).

In the cycle presented in Example 2.9 has a flavor similar to Example 2.6: when the weaker agent (high cost $c_1(z) = z^{1.1}$) produces low output, then the stronger agent also produces relatively low output. But then the weaker agent tries to compete by producing higher output, then the stronger agent responds with a relatively higher output, and then the weaker agent backs off and produces low output. And this cycle continues.

We note that Example 2.9 satisfies the relevant properties discussed earlier using linear costs. In particular, the initial states that lead to the cycle can have positive measure. The initial profile can also be close to 0, e.g., starting from $(\cdot, 10^{-5})$ and agent 1 making the first move leads to the cycle presented above. In our numerical simulations, we observe that *stronger* convexity leads to faster and more robust convergence results, i.e., when the cost functions are *more* convex, non-convergence becomes unlikely and the rate of convergence improves. Our intuition for these observations is the fact that convexity of the cost function causes an agent to play closer to the equilibrium output, i.e., given the same total output by other agents, s_{-i}, the BR of agent i, $BR(s_{-i})$, is closer to the equilibrium output if the cost function is more convex. (A version of this behavior can be derived using the first-order condition, equation (2.4), and has been used in our convergence analysis in Section 2.5.) We leave the formal analysis of these aspects of the BR dynamics to future work.

2.7 Conclusion

We showed fast convergence of BR dynamics in Tullock contests with homogeneous agents and convex cost functions. There are several open problems:[11] (1) Our upper bound for the discounted-sum dynamics has a $\tilde{O}(1/L^2)$ dependency on L (an $\tilde{O}(n^2)$ dependency for uniform sampling of agents), but our lower bound is a $\tilde{\Omega}(1/L)$ ($\tilde{\Omega}(n)$

[11]Relevant simulations for some of these directions is available on `https://github.com/abheekg/br-tullock-convex`.

for uniform sampling). A direct open problem is to close this gap. We think that the slack is in the upper bound, which can be improved to $\tilde{O}(1/L)$. As highlighted by Example 2.4, our potential function and analysis might not help improve the upper bound, and this example applies even for fixed distributions that do not vary over time, particularly uniform sampling. Techniques from linear dynamical systems may help get a better bound. The discounted-sum dynamics is an instance of a time-varying linear dynamical system, with a specific structure for the matrices defining the linear transformation in each time step. One can study the product of the time-varying matrices of this linear dynamical system. The most natural idea would be to bound the spectral norm of these matrices (or some other suitable matrix norm). But the spectral norm for the individual matrices can be strictly larger than 1 for this dynamics (the same result holds for similar norms). So, we will likely need additional ideas. (2) Analyze the convergence of BR dynamics for *almost* homogeneous agents.[12] (3) Study the effect of the convexity of the cost functions on the convergence rate.

2.8 Omitted Proofs

2.8.1 From Section 2.3

Proof of Lemma 2.2. For cleaner exposition, let $v = x_{-i} = s_{-i}$ and $y = BR(x_{-i}) = BR(v)$. If $v = 0$, then $y = BR(v) = a < 1$ by definition. Let us assume $v > 0$. From the first order condition on the utility function, equation (2.4), we have

$$\left.\frac{\partial u_i(z,v)}{\partial z}\right|_{z=y=BR(v)} = 0 \implies \frac{v}{(y+v)^2} - \frac{1}{4}c'(y) = 0. \tag{2.8}$$

We know that $u_i(z,v)$ is a strictly concave function of z, or, in other words, $\frac{\partial u_i(z,v)}{\partial z}$ is a strictly decreasing function of z. Substituting z by v in $\frac{\partial u_i(z,v)}{\partial z}$, we get

$$\left.\frac{\partial u_i(z,v)}{\partial z}\right|_{z=v} = \frac{v}{4v^2} - \frac{c'(v)}{4} = \frac{1}{4v}\left(1 - vc'(v)\right).$$

If $v < 1$, then $\left.\frac{\partial u_i(z,v)}{\partial z}\right|_{z=v} = 1 - vc'(v) > 0$ because $c'(v) < c'(1) = 1$. And, as $\frac{\partial u_i(z,v)}{\partial z}$ is strictly decreasing in z, we have $y > v$. With a similar argument, if $v = 1$, then $y = 1$, and if $v > 1$, then $y < v$.

[12]The ratio of the slope of the cost functions of the agents is bounded by a small constant, e.g., ≤ 4.

Notice that equation (2.8) implicitly defines y as a function of v. Let us now compute the derivative of y w.r.t. v by differentiating equation (2.8) w.r.t. v,

$$\begin{aligned}
&\frac{d}{dv}\left(\frac{v}{(y+v)^2}-\frac{1}{4}c'(y)\right)=0\\
&\qquad\implies \frac{1}{(y+v)^2}-\frac{2v}{(y+v)^3}+\frac{dy}{dv}\left(\frac{-2v}{(y+v)^3}-\frac{1}{4}c''(y)\right)=0\\
&\qquad\implies \frac{(y+v)-2v}{(y+v)^3}=\frac{dy}{dv}\frac{1}{(y+v)^3}\left(2v+\frac{(y+v)^3}{4}c''(y)\right)\\
&\qquad\implies \frac{dy}{dv}=\frac{y-v}{2v+\frac{(y+v)^3}{4}c''(y)}.
\end{aligned}$$

As $c''(y)\geq 0$, if $y\geq v$ then $\frac{dy}{dv}\geq 0$. We showed earlier that $v\leq 1 \iff y\geq v$, so $v\leq 1 \implies \frac{dy}{dv}\geq 0$. With a similar argument, $v\geq 1 \iff y\leq v \implies \frac{dy}{dv}\leq 0$. Therefore, the maximum value of y occurs at $v=1$, where y is also 1. Same steps also show that, if $v<1$, then $y>v$, then $\frac{dy}{dv}>0$, which implies $y<1$. $\square$

Proof of Lemma 2.3. We know from the first-order condition that $\frac{z_t}{(z_t+z_{t+1})^2}=\frac{1}{4}c'(z_{t+1})$. As $z_{t+1}<1 \implies c'(z_{t+1})\leq 1$ and $z_{t+1}>z_t$ for every t, we have

$$\frac{z_t}{(z_t+z_{t+1})^2}=\frac{c'(z_{t+1})}{4} \implies \frac{z_t}{(2z_{t+1})^2}\leq\frac{1}{4} \implies z_{t+1}\geq\sqrt{z_t} \implies z_t\geq z_0^{\frac{1}{2^t}}.$$

Now, we want $z_t\geq\frac{1}{2}$. So, if we ensure that $z_0^{\frac{1}{2^t}}=\gamma^{\frac{1}{2^t}}\geq\frac{1}{2}$, then we are done.

$$\gamma^{\frac{1}{2^t}}\geq\frac{1}{2}\iff\frac{1}{2^t}\lg(\gamma)\geq -1\iff\lg\left(\frac{1}{\gamma}\right)\leq 2^t\iff t\geq\lg\lg\left(\frac{1}{\gamma}\right).$$

$\square$

Proof of Lemma 2.4. As $z_{t+1}<1$ for every t, which implies $c'(z_{t+1})\leq 1$, we have

$$\frac{z_t}{(z_t+z_{t+1})^2}=\frac{c'(z_{t+1})}{4} \implies 4z_t\leq(z_t+z_{t+1})^2 \implies z_{t+1}\geq\sqrt{z_t}(2-\sqrt{z_t}).$$

Let $\zeta_t=1-\sqrt{z_t}\iff z_t=(1-\zeta_t)^2$. From the above inequality, we have

$$\begin{aligned}
z_{t+1}\geq\sqrt{z_t}(2-\sqrt{z_t}) &\iff (1-\zeta_{t+1})^2\geq(1-\zeta_t)(1+\zeta_t)=1-\zeta_t^2\\
&\iff 1+\zeta_{t+1}^2-2\zeta_{t+1}\geq 1-\zeta_t^2\iff\zeta_{t+1}\leq(\zeta_{t+1}^2+\zeta_t^2)/2.
\end{aligned}$$

As $z_t<z_{t+1}$, so $\zeta_t>\zeta_{t+1}$, which implies

$$\zeta_{t+1}\leq(\zeta_{t+1}^2+\zeta_t^2)/2 \implies \zeta_{t+1}\leq\zeta_t^2 \implies \zeta_t\leq\zeta_0^{2^t}.$$

From the initial condition, we know that $z_0 \geq \frac{1}{2} \implies \zeta_0 \leq 1 - \frac{1}{\sqrt{2}} \leq \frac{1}{2}$. We want $z_t \geq 1-\epsilon \impliedby \sqrt{z_t} \geq \sqrt{1-\epsilon}$. Using standard inequalities, $\sqrt{1-\epsilon} \leq \sqrt{e^{-\epsilon}} = e^{-\epsilon/2} \leq 1 - \epsilon/4$ for every $\epsilon \leq 1$. So, $\sqrt{z_t} \geq \sqrt{1-\epsilon}$ is implied by $\sqrt{z_t} \geq 1 - \epsilon/4$. Putting everything together,

$$
\begin{aligned}
z_t \geq 1-\epsilon \impliedby \sqrt{z_t} \geq 1 - \frac{\epsilon}{4} \iff \zeta_t \leq \frac{\epsilon}{4} \impliedby \zeta_0^{2^t} \leq \frac{\epsilon}{4} \iff 2^t \lg(\zeta_0) \leq \lg(\epsilon) - 2 \\
\iff 2^t \lg\left(\frac{1}{\zeta_0}\right) \geq 2 + \lg\left(\frac{1}{\epsilon}\right) \impliedby 2^t \geq 3 \lg\left(\frac{1}{\epsilon}\right) \impliedby t \geq 2 + \lg\lg\left(\frac{1}{\epsilon}\right).
\end{aligned}
$$

□

Proof Lemma 2.5. Fix an arbitrary agent i. Let $u_- = u_i(x_i, x_{-i})$ and $u_+ = u_i(BR(x_{-i}), x_{-i})$. We want to prove that

$$u_- \geq (1-3\epsilon)u_+ \iff \frac{u_-}{u_+} \geq 1 - 3\epsilon.$$

We are given that $1-\epsilon \leq x_i, x_{-i} \leq 1$, and from Lemma 2.2, we also know that $1 - \epsilon \leq BR(x_{-i}) \leq 1$. Further, as c is continuous and $c'(z) \leq 1$ for $z \leq 1$, we have $c(z) \geq c(1) - \epsilon$ for $z \geq 1 - \epsilon$. So, we can lower bound u_- as

$$
\begin{aligned}
u_- = u_i(x_i, x_{-i}) = \frac{x_i}{x_i + x_{-i}} - \frac{1}{4}c(x_i) \geq \min_{y,z \in [1-\epsilon,1]} \left(\frac{y}{y+z} - \frac{1}{4}c(y)\right) \\
\geq \min_{y,z \in [1-\epsilon,1]} \frac{y}{y+z} - \max_{y \in [1-\epsilon,1]} \frac{1}{4}c(y) \geq \min_{y \in [1-\epsilon,1]} \frac{y}{y+1} - \frac{1}{4}c(1) \geq \frac{1-\epsilon}{2-\epsilon} - \frac{1}{4}c(1).
\end{aligned}
$$

Similarly, we can upper bound u_+ as

$$
\begin{aligned}
u_+ = u_i(BR(x_{-i}), x_{-i}) &= \frac{BR(x_{-i})}{BR(x_{-i}) + x_{-i}} - \frac{1}{4}c(BR(x_{-i})) \\
&\leq \max_{y,z \in [1-\epsilon,1]} \left(\frac{y}{y+z} - \frac{1}{4}c(y)\right) \leq \max_{y,z \in [1-\epsilon,1]} \frac{y}{y+z} - \min_{y \in [1-\epsilon,1]} \frac{1}{4}c(y) \\
&\leq \max_{y \in [1-\epsilon,1]} \frac{y}{y+1-\epsilon} - \frac{1}{4}c(1-\epsilon) \leq \frac{1}{2-\epsilon} - \frac{1}{4}(c(1)-\epsilon)
\end{aligned}
$$

Putting these two bounds together, we get

$$
\begin{aligned}
\frac{u_-}{u_+} &\geq \frac{\frac{1-\epsilon}{2-\epsilon} - \frac{1}{4}c(1)}{\frac{1}{2-\epsilon} - \frac{1}{4}(c(1)-\epsilon)} = \frac{1 - \frac{2-\epsilon}{4}c(1) - \epsilon}{1 - \frac{2-\epsilon}{4}c(1) + \frac{2-\epsilon}{4}\epsilon} \geq \frac{1 - \frac{1}{2}c(1) - \epsilon}{1 - \frac{1}{2}c(1) + \frac{1}{2}\epsilon} \\
&\geq \frac{1-2\epsilon}{1+\epsilon}, \text{ because } c(1) \leq 1 \text{ as } c(0) = 0 \text{ and } c'(z) \leq 1 \text{ for } z \leq 1, \\
&\geq (1-2\epsilon)(1-\epsilon) \geq 1 - 3\epsilon, \text{ because } \frac{1}{1+z} \geq 1 - z \text{ for } 0 \leq z \leq 1.
\end{aligned}
$$

□

Proof of Theorem 2.6. The proof is similar to the proof of Theorem 2.1. We define the sequence $(z_t)_{t\geq 0}$ the same way as Theorem 2.1. We can use the analysis for Theorem 2.1 until Lemma 2.3; then we prove stronger versions of Lemmas 2.3 and 2.4 assuming $c'(x_i) \in [x_i^p, x_i^q]$, given in Lemmas 2.19 and 2.20, respectively. We can also simplify the value of γ: if $z_0 = x_{0,1} > 1$ and $c'(0) < \frac{4}{z_0}$, then $\gamma = (c')^{-1}\left(\frac{1}{z_0}\right) = \left(\frac{1}{z_0}\right)^{1/r}$ for some $r \in [q, p]$.

Lemma 2.19. *Given $z_0 = \gamma \in (0, \frac{1}{2}]$, $z_t \geq \frac{1}{2}$ for all $t \geq \frac{1}{\lg(2+q)} \lg\lg\left(\frac{1}{\gamma}\right)$ and $z_t < \frac{1}{2}$ for all $t < \frac{1}{\lg(2+p)} \lg\lg\left(\frac{1}{\gamma}\right) - O(1)$.*

Proof. We know from the first-order condition that $\frac{z_t}{(z_t+z_{t+1})^2} = \frac{1}{4}c'(z_{t+1})$. As $c'(z_{t+1}) \leq z_{t+1}^q$ and $z_{t+1} > z_t$ for every t,

$$\frac{z_t}{(z_t+z_{t+1})^2} = \frac{c'(z_{t+1})}{4} \implies \frac{z_t}{(2z_{t+1})^2} \leq \frac{z_{t+1}^q}{4} \implies z_{t+1} \geq z_t^{\frac{1}{2+q}}$$
$$\implies z_t \geq z_0^{\frac{1}{(2+q)^t}} = \gamma^{\frac{1}{(2+q)^t}}.$$

Now, we want $z_t \geq \frac{1}{2}$. So, if we ensure that $\gamma^{\frac{1}{(2+q)^t}} \geq \frac{1}{2}$, then we are done.

$$\gamma^{\frac{1}{(2+q)^t}} \geq \frac{1}{2} \iff \frac{1}{(2+q)^t}\lg(\gamma) \geq -1 \iff \lg\left(\frac{1}{\gamma}\right) \leq (2+q)^t$$
$$\iff t \geq \frac{1}{\lg(2+q)} \lg\lg\left(\frac{1}{\gamma}\right).$$

Let us now look at the lower bound. As $c'(z_{t+1}) \geq z_{t+1}^p$ and $z_t \geq 0$ for all t,

$$\frac{z_t}{(z_t+z_{t+1})^2} = \frac{c'(z_{t+1})}{4} \implies \frac{z_t}{(z_{t+1})^2} \geq \frac{z_{t+1}^p}{4} \implies z_{t+1} \leq 4z_t^{\frac{1}{2+p}}$$
$$\implies z_t \leq 4z_{t-1}^{\frac{1}{2+p}} \leq 4^{1+\frac{1}{2+p}} z_{t-2}^{\frac{1}{(2+p)^2}} \leq 4^{1+\frac{1}{2+p}+\frac{1}{(2+p)^2}} z_{t-3}^{\frac{1}{(2+p)^3}} \leq \ldots \leq 16 z_0^{\frac{1}{(2+p)^t}}.$$

Now, we want $z_t < \frac{1}{2}$. So, if we ensure that $16z_0^{\frac{1}{(2+p)^t}} = 16\gamma^{\frac{1}{(2+p)^t}} < \frac{1}{2}$, then we are done.

$$16\gamma^{\frac{1}{(2+p)^t}} < \frac{1}{2} \iff \frac{1}{(2+p)^t}\lg(\gamma) < -5 \iff \lg\left(\frac{1}{\gamma}\right) > 5(2+p)^t$$
$$\iff t < \frac{1}{\lg(2+p)} \lg\lg\left(\frac{1}{\gamma}\right) - \lg(5).$$

□

Lemma 2.20. *Given $z_0 \geq \frac{1}{2}$, $z_t \geq 1-\epsilon$ for all $t \geq \lg\lg(\frac{1}{\epsilon}) - \lg\lg(2+q) + O(1)$. On the other hand, given $z_0 \leq \frac{1}{2}$, $z_t < 1-\epsilon$ for all $t < \lg\lg(\frac{1}{\epsilon}) - \lg\lg(2+p) - O(1)$.*

Proof. As $z_{t+1} < 1$ for every t, which implies $c'(z_{t+1}) \leq 1$, we have

$$\frac{z_t}{(z_t + z_{t+1})^2} = \frac{c'(z_{t+1})}{4} \implies 4z_t \leq (z_t + z_{t+1})^2 \implies z_{t+1} \geq \sqrt{z_t}(2 - \sqrt{z_t}).$$

Let $\zeta_t = 1 - \sqrt{z_t} \iff z_t = (1 - \zeta_t)^2$. From the above inequality, we have

$$\begin{aligned} z_{t+1} \geq \sqrt{z_t}(2 - \sqrt{z_t}) &\iff (1 - \zeta_{t+1})^2 \geq (1 - \zeta_t)(1 + \zeta_t) = 1 - \zeta_t^2 \\ &\iff 1 + \zeta_{t+1}^2 - 2\zeta_{t+1} \geq 1 - \zeta_t^2 \iff 2\zeta_{t+1} \leq \zeta_{t+1}^2 + \zeta_t^2. \end{aligned} \tag{2.9}$$

We will use inequality (2.9), $2\zeta_{t+1} \leq \zeta_{t+1}^2 + \zeta_t^2$, in our subsequent analysis.

We have been given that $z_t^p \leq c'(z_t) \leq z_t^q$ for all t. Making use of the upper bound, we get

$$\begin{aligned} &\frac{z_t}{(z_t + z_{t+1})^2} = \frac{c'(z_{t+1})}{4} \leq \frac{z_{t+1}^q}{4} \implies \sqrt{z_t} \leq z_{t+1}^{\frac{q}{2}} \left(\frac{z_t + z_{t+1}}{2} \right) \\ &\iff 1 - \zeta_t \leq (1 - \zeta_{t+1})^{\frac{q}{2}} \left(\frac{(1 - \zeta_t)^2 + (1 - \zeta_{t+1})^2}{2} \right) \\ &\iff 1 - \zeta_t \leq (1 - \zeta_{t+1})^{\frac{q}{2}} \left(1 - \zeta_t + \frac{\zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1}}{2} \right) \\ &\iff 1 \leq (1 - \zeta_{t+1})^{\frac{q}{2}} \left(1 + \frac{\zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1}}{2(1 - \zeta_t)} \right). \end{aligned}$$

From inequality (2.9), $2\zeta_{t+1} \leq \zeta_{t+1}^2 + \zeta_t^2 \implies \frac{\zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1}}{2(1-\zeta_t)} \leq \zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1}$ for $\zeta_t \leq \frac{1}{2}$. Also, $(1 - \zeta_{t+1})^{\frac{q}{2}} \leq e^{-\frac{\zeta_{t+1} q}{2}} \leq 1 - \frac{\zeta_{t+1} q}{4}$ for $\zeta_{t+1} \leq \frac{1}{8q}$. (If $\zeta_{t+1} > \frac{1}{8q}$, then we can use an argument similar to Lemma 2.19 to reach $\zeta_{t+1} \leq \frac{1}{8q}$ in a constant number of steps from $\zeta_0 \leq \frac{1}{2}$.) We get

$$\begin{aligned} &1 \leq \left(1 - \frac{q\zeta_{t+1}}{4} \right) \left(1 + \zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1} \right) \\ &\implies \zeta_{t+1} \left(2 + \frac{q}{4} \right) \leq (\zeta_t^2 + \zeta_{t+1}^2) - \frac{q\zeta_{t+1}}{4} (\zeta_t^2 + \zeta_{t+1}^2 - 2\zeta_{t+1}) \\ &\implies \zeta_{t+1} \leq \frac{8}{8 + q} \zeta_t^2 \implies \zeta_t \leq \left(\frac{8}{8 + q} \right)^{2^t - 1} \zeta_0^{2^t}. \end{aligned}$$

From the initial condition, we have $z_0 \geq \frac{1}{2} \implies \zeta_0 \leq 1 - \frac{1}{\sqrt{2}} \leq \frac{1}{2}$. We want $z_t \geq 1 - \epsilon \impliedby \sqrt{z_t} \geq \sqrt{1 - \epsilon}$. Using standard inequalities, $\sqrt{1 - \epsilon} \leq \sqrt{e^{-\epsilon}} = e^{-\epsilon/2} \leq 1 - \epsilon/4$ for every $\epsilon \leq 1$. So, $\sqrt{z_t} \geq \sqrt{1 - \epsilon}$ is implied by $\sqrt{z_t} \geq 1 - \epsilon/4$.

Putting everything together,

$$\begin{aligned}
z_t \geq 1-\epsilon \impliedby \sqrt{z_t} \geq 1-\frac{\epsilon}{4} \iff \zeta_t \leq \frac{\epsilon}{4} \impliedby \left(\frac{8}{8+q}\right)^{2^t-1} \zeta_0^{2^t} \leq \frac{\epsilon}{4} \\
\impliedby \left(\frac{8\zeta_0}{8+q}\right)^{2^{t-1}} \leq \frac{\epsilon}{4} \iff 2^{t-1}\lg\left(\frac{8\zeta_0}{8+q}\right) \leq \lg(\epsilon)-2 \\
\iff 2^{t-1}\lg\left(\frac{1+q/8}{\zeta_0}\right) \geq 2+\lg\left(\frac{1}{\epsilon}\right) \impliedby t \geq \lg\lg\left(\frac{1}{\epsilon}\right)-\lg\lg(2+q)+O(1).
\end{aligned}$$

Let us now follow similar steps to derive the lower bound. Using $c'(z_t) \geq z_t^p$, we have

$$\begin{aligned}
\frac{z_t}{(z_t+z_{t+1})^2} = \frac{c'(z_{t+1})}{4} \geq \frac{z_{t+1}^p}{4} &\implies \sqrt{z_t} \geq z_{t+1}^{\frac{p}{2}}\left(\frac{z_t+z_{t+1}}{2}\right) \\
&\iff 1-\zeta_t \geq (1-\zeta_{t+1})^{\frac{p}{2}}\left(\frac{(1-\zeta_t)^2+(1-\zeta_{t+1})^2}{2}\right) \\
&\iff 1-\zeta_t \geq (1-\zeta_{t+1})^{\frac{p}{2}}\left(1-\zeta_t+\frac{\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}}{2}\right) \\
&\iff 1 \geq (1-\zeta_{t+1})^{\frac{p}{2}}\left(1+\frac{\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}}{2(1-\zeta_t)}\right).
\end{aligned}$$

From inequality (2.9), $2\zeta_{t+1} \leq \zeta_{t+1}^2+\zeta_t^2 \implies \frac{\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}}{2(1-\zeta_t)} \geq \frac{\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}}{2}$ for $\zeta_t \geq 0$. Also, $(1-\zeta_{t+1})^{\frac{p}{2}} \geq e^{-\zeta_{t+1}p} \geq 1-\zeta_{t+1}p$ for $\zeta_{t+1} \leq \frac{1}{2}$. We get

$$\begin{aligned}
&1 \geq (1-p\zeta_{t+1})\left(1+\frac{\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}}{2}\right) \\
&\implies \zeta_{t+1}(p+1) \geq \frac{\zeta_t^2+\zeta_{t+1}^2}{2}-\frac{p\zeta_{t+1}}{2}(\zeta_t^2+\zeta_{t+1}^2-2\zeta_{t+1}) \\
&\implies \zeta_{t+1}\left(\frac{3p}{2}+1\right) \geq \frac{\zeta_t^2+\zeta_{t+1}^2}{2} \implies \zeta_{t+1} \geq \frac{1}{2+3p}\zeta_t^2 \implies \zeta_t \geq \left(\frac{1}{2+3p}\right)^{2^t}\zeta_0^{2^t}.
\end{aligned}$$

From the initial condition, we have $z_0 \leq \frac{1}{2} \implies \zeta_0 \geq 1-\frac{1}{\sqrt{2}} \geq \frac{1}{4}$. We want $z_t < 1-\epsilon \impliedby \sqrt{z_t} < \sqrt{1-\epsilon}$. As $\sqrt{1-\epsilon} \geq 1-\epsilon$ for every $\epsilon \leq 1$, so $\sqrt{z_t} < \sqrt{1-\epsilon}$ is implied by $\sqrt{z_t} < 1-\epsilon$.

$$\begin{aligned}
z_t < 1-\epsilon \impliedby \sqrt{z_t} < 1-\epsilon \iff \zeta_t > \epsilon \impliedby \left(\frac{1}{2+3p}\right)^{2^t}\zeta_0^{2^t} > \epsilon \\
\iff 2^t\lg\left(\frac{\zeta_0}{2+3p}\right) > \lg(\epsilon) \iff 2^t\lg\left(\frac{2+3p}{4}\right) < \lg\left(\frac{1}{\epsilon}\right) \\
\impliedby t < \lg\lg\left(\frac{1}{\epsilon}\right)-\lg\lg(2+p)-O(1).
\end{aligned}$$

□

Lemmas 2.19 and 2.20 give tight lower and upper bounds on the time required for z_t to reach $1-\epsilon$ starting from γ. Applying Lemma 2.5 completes the proof for the upper bound. For the lower bound, notice that if $z_t \in [\frac{1}{2}, 1-\epsilon]$ and agent i makes the move at time t, then $BR(x_{t,-i}) - x_{t,i} = z_{t+1} - z_{t-1} \geq z_{t+1} - z_t \geq \Omega(\epsilon)$ using arguments given in Lemma 2.4 (if $z_t = 1-\zeta \leq 1-\epsilon$, then $z_{t+1} \geq 1-\zeta^2$, which implies $z_{t+1} - z_t \geq \zeta - \zeta^2 = \Omega(\zeta) = \Omega(\epsilon)$ for $\zeta \leq 1/2$); applying Lemma 2.23 completes the proof. □

2.8.2 From Section 2.4

Proof of Lemma 2.7, remaining portion. We now bound the rate of convergence. We will show that for every pair of consecutive time steps, the expected value of the potential decreases by at least a multiplicative factor of $(1 - \frac{\kappa(1-B)L^2}{\log(n)})$ for some constant $\kappa > 0$.

Let us look at two consecutive time steps t and $t+1$. Let w.l.o.g. $f(\boldsymbol{z}_t) > 0$ and $V_t \geq W_t$. So, $f(\boldsymbol{z}_t) = V_t > 0$. Using Lemma 2.21 below, we know that that there exists $k, \ell \in [n]$ such that there are at least k agents $i \in \mathcal{V}_t$ with $v_{t,i} \geq \frac{V_t}{4k\lg(n)}$ and at least ℓ agents $j \in \mathcal{W}_t$ with $w_{t,j} \geq \frac{W_t}{4\ell\lg(n)}$.

Lemma 2.21. *For any $(p_i)_{i\in[n]}$ with $p_i \geq 0$ for all $i \in [n]$ and $\sum_{i\in[n]} p_i = 1$, there exists $k \in [n]$ such that $|\{i \in [n] \mid p_i \geq \frac{1}{4k\lg(n)}\}| \geq k$.*

Say at time t, we pick one of the k agents in $\mathcal{V}_t$, say agent i, with $v_{t,i} \geq \frac{V_t}{4k\lg(n)}$, and at time $t+1$, we pick one of the ℓ agents in $\mathcal{W}_t$, say agent j, with $w_{t,j} \geq \frac{W_t}{4\ell\lg(n)}$. We have the following possible scenarios:

1. $W_t \leq V_t - v_{t,i}$. Then agent i will move from $\mathcal{V}_t$ to $\mathcal{W}_{t+1}$ after the transition and will have an updated value of $v_{t+1,i} = 0$ and $w_{t+1,i} \leq B(V_t - v_{t,i} - W_t)$. So, the potential at time $t+1$ is

$$\begin{aligned} f(\boldsymbol{z}_{t+1}) &\leq \max(V_t - v_{t,i}, W_t + B(V_t - v_{t,i} - W_t)) \\ &= \max(V_t - v_{t,i}, (1-B)W_t + B(V_t - v_{t,i})) \\ &= V_t - v_{t,i} \leq \left(1 - \frac{1}{4k\lg(n)}\right) V_t = \left(1 - \frac{1}{4k\lg(n)}\right) f(\boldsymbol{z}_t). \end{aligned}$$

The probability that one of these k agents with $v_{t,i} \geq \frac{V_t}{4k\lg(n)}$ is picked is at least kL, so the expected value of the potential at time $t+1$ is $\mathbb{E}[f(\boldsymbol{z}_{t+1})|\boldsymbol{z}_t] \leq \left(1 - \frac{L}{4\lg(n)}\right) f(\boldsymbol{z}_t) \implies \mathbb{E}[f(\boldsymbol{z}_{t+2})|\boldsymbol{z}_t] \leq \mathbb{E}[f(\boldsymbol{z}_{t+1})|\boldsymbol{z}_t] \leq \left(1 - \frac{L}{4\lg(n)}\right) f(\boldsymbol{z}_t)$.

2. $W_t > V_t - v_{t,i}$. As W_t can be very close to V_t, we will need two transitions to guarantee sufficient progress. As $v_{t,i} > V_t - W_t$, agent i will stay in $\mathcal{V}_{t+1}$ after the transition and will have an updated value of $v_{t+1,i} \le B(W_t - (V_t - v_{t,i}))$. We can bound V_{t+1} as

$$\begin{aligned} V_{t+1} &\le V_t - v_{t,i} + B(W_t - (V_t - v_{t,i})) = (1-B)V_t + BW_t - (1-B)v_{t,i} \\ &\le V_t - (1-B)v_{t,i} \le \left(1 - \frac{1-B}{4k\lg(n)}\right) V_t. \end{aligned}$$

Now, at time $t+1$, agent j with $w_{t+1,j} = w_{t,j} \ge \frac{W_t}{4\ell\lg(n)} = \frac{W_{t+1}}{4\ell\lg(n)}$ makes a move. We can have the following sub-cases depending upon the values of $w_{t+1,j}$, W_{t+1}, and V_{t+1}:

(a) $V_{t+1} < W_{t+1} - w_{t+1,j}$. Then agent j will move from $\mathcal{W}_{t+1}$ to $\mathcal{V}_{t+2}$ after the transition and will have an updated value of $w_{t+2,j} = 0$ and $v_{t+2,j} \le B(W_{t+1} - w_{t+1,j} - V_{t+1})$. So, the potential at time $t+2$ is

$$\begin{aligned} f(\boldsymbol{z}_{t+2}) &\le \max(V_{t+1} + B(W_{t+1} - w_{t+1,j} - V_{t+1}), W_{t+1} - w_{t+1,j}) \\ &= \max((1-B)V_{t+1} + B(W_{t+1} - w_{t+1,j}), W_{t+1} - w_{t+1,j}) \\ &= W_{t+1} - w_{t+1,j} \le \left(1 - \frac{1}{4\ell\lg(n)}\right) W_{t+1} \\ &= \left(1 - \frac{1}{4\ell\lg(n)}\right) f(\boldsymbol{z}_{t+1}) \le \left(1 - \frac{1}{4\ell\lg(n)}\right) f(\boldsymbol{z}_t). \end{aligned}$$

(b) $V_{t+1} \ge W_{t+1} - w_{t+1,j}$. Then agent j will stay in $\mathcal{W}_{t+2}$ after the transition and will have an updated value of $w_{t+2,j} \le B(V_{t+1} - (W_{t+1} - w_{t+1,j}))$. So, the potential at time $t+2$ is

$$\begin{aligned} f(\boldsymbol{z}_{t+2}) &\le \max(V_{t+1}, W_{t+1} - w_{t+1,j} + B(V_{t+1} - (W_{t+1} - w_{t+1,j}))) \\ &= \max(V_{t+1}, (1-B)(W_{t+1} - w_{t+1,j}) + BV_{t+1}) = V_{t+1} \\ &\le \left(1 - \frac{1-B}{4k\lg(n)}\right) V_t = \left(1 - \frac{1-B}{4k\lg(n)}\right) f(\boldsymbol{z}_t). \end{aligned}$$

Now, combining the above two cases, we have

$$\begin{aligned} f(\boldsymbol{z}_{t+2}) &\le \max\left(\left(1 - \frac{1}{4\ell\lg(n)}\right) f(\boldsymbol{z}_t), \left(1 - \frac{1-B}{4k\lg(n)}\right) f(\boldsymbol{z}_t)\right) \\ &= f(\boldsymbol{z}_t) - f(\boldsymbol{z}_t)\frac{1}{4k\ell\lg(n)} \min(k, (1-B)\ell) \le \left(1 - \frac{1-B}{4k\ell\lg(n)}\right) f(\boldsymbol{z}_t). \end{aligned}$$

Now, the probability that one of the k agents with $v_{t,i} \ge \frac{V_t}{4k\lg(n)}$ is picked at time t is at least kL and that one of the ℓ agents with $w_{t,i} \ge \frac{W_t}{4\ell\lg(n)}$ is picked at time

$t+1$ is at least ℓL, so the probability for the pair of transitions is at least $k\ell L^2$. So, the expected value of the potential is $\mathbb{E}[f(\boldsymbol{z}_{t+2})|\boldsymbol{z}_t] \leq \left(1 - \frac{(1-B)L^2}{4\lg(n)}\right) f(\boldsymbol{z}_t)$.

Putting everything together, we have $\mathbb{E}[f(\boldsymbol{z}_{t+2})|\boldsymbol{z}_t] \leq \left(1 - \frac{(1-B)L^2}{4\lg(n)}\right) f(\boldsymbol{z}_t) \implies \mathbb{E}[f(\boldsymbol{z}_{t+2})] = \mathbb{E}[\mathbb{E}[f(\boldsymbol{z}_{t+2})|\boldsymbol{z}_t]] \leq \left(1 - \frac{(1-B)L^2}{4\lg(n)}\right) \mathbb{E}[f(\boldsymbol{z}_t)]$. Therefore, after $T = \frac{8\lg(n)}{(1-B)L^2} \ln\left(\frac{2f(\boldsymbol{z}_0)}{\epsilon\delta}\right)$ time steps, we have

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{z}_T)] &\leq \left(1 - \frac{(1-B)L^2}{4\lg(n)}\right)^{T/2} f(\boldsymbol{z}_0) \leq e^{-\frac{(1-B)L^2}{4\lg(n)}\frac{T}{2}} f(\boldsymbol{z}_0) \\ &= e^{-\ln\left(\frac{2f(\boldsymbol{z}_0)}{\epsilon\delta}\right)} f(\boldsymbol{z}_0) \leq \frac{\epsilon\delta}{2}.\end{aligned}$$

Finally, using Markov inequality we have $\mathbb{P}[f(\boldsymbol{z}_T) \geq \frac{\epsilon}{2}] \leq \frac{2\,\mathbb{E}[f(\boldsymbol{z}_T)]}{\epsilon} = \delta$. And it is easy to check that $||\boldsymbol{z}_T||_1 \leq 2f(\boldsymbol{z}_T) < 2\frac{\epsilon}{2} = \epsilon$. $\square$

Proof of Lemma 2.21. We prove the theorem by contradiction. Let us assume that for every $k \in [n]$ there are strictly less than k coordinates $i \in [n]$ such that $p_i \geq \frac{1}{4k\lg(n)}$. Then we have

$$\begin{aligned}\sum_i p_i = 1 = \sum_i \mathbb{1}\left(\frac{1}{4\lg(n)} \leq v_i \leq 1\right) v_i \\ + \sum_i \sum_{j=1}^{\lceil\lg(n)\rceil-1} \mathbb{1}\left(\frac{1}{2^{j+2}\lg(n)} \leq v_i < \frac{1}{2^{j+1}\lg(n)}\right) v_i \\ + \sum_i \mathbb{1}\left(0 \leq v_i < \frac{1}{2^{\lceil\lg(n)\rceil+1}\lg(n)}\right) v_i\end{aligned}$$

$$\begin{aligned}&\leq \sum_i \mathbb{1}\left(\frac{1}{4\lg(n)} \leq v_i \leq 1\right) \\ &\quad + \sum_i \sum_{j=1}^{\lceil\lg(n)\rceil-1} \mathbb{1}\left(\frac{1}{2^{j+2}\lg(n)} \leq v_i < \frac{1}{2^{j+1}\lg(n)}\right) \frac{1}{2^{j+1}\lg(n)} \\ &\quad + \sum_i \mathbb{1}\left(0 \leq v_i < \frac{1}{2^{\lceil\lg(n)\rceil+1}\lg(n)}\right) \frac{1}{2^{\lceil\lg(n)\rceil+2}\lg(n)} \\ &< \sum_{j=1}^{\lceil\lg(n)\rceil-1} \frac{2^j}{2^{j+1}\lg(n)} + \frac{n}{2^{\lceil\lg(n)\rceil+1}\lg(n)}, \\ &\qquad \text{as there are} < k \text{ values} \geq \frac{1}{4k\lg(n)} \text{ and at most } n \text{ values total,} \\ &\leq \frac{\lceil\lg(n)\rceil}{2\lg(n)} < 1.\end{aligned}$$

$\square$

Proof of Lemma 2.8. The proof for the upper bound is a simplified version of the argument used in the proof of Lemma 2.7. As we do not have any randomness, we can deterministically pick the agent i_t at time t. In particular, we pick the largest element in the larger side of $f(\boldsymbol{z}_t)$ (where the two sides correspond to the sum of the positive and the negative values of $\boldsymbol{z}_t$).

Formally, let $V_t = \sum_j \mathbb{1}(z_{t,j} > 0)z_{t,j}$ and $W_t = \sum_j \mathbb{1}(z_{t,j} < 0)z_{t,j}$. W.l.o.g. let $V_t \geq W_t$. So, $f(\boldsymbol{z}_t) = V_t$. We pick $i_t = \arg\max_{j|z_{t,j}>0} z_{t,j}$, where $z_{t,i_t} \geq \frac{V_t}{n}$ by definition, and (if required) pick $i_{t+1} = \arg\max_{j|z_{t,j}<0}(-z_{t,j})$, where $z_{t,i_{t+1}} \geq \frac{W_t}{n}$ by definition. We have the following cases based on the value of z_{t,i_t}, V_t, and W_t:

- $W_t \leq V_t - z_{t,i_t}$. Following the same steps as the proof of Lemma 2.7 case (1), we have $f(\boldsymbol{z}_{t+1}) \leq V_t - z_{t,i_t} \leq (1 - \frac{1}{n})V_t = (1 - \frac{1}{n})f(\boldsymbol{z}_t)$.
- $W_t > V_t - z_{t,i_t}$. Following the same steps as the proof of Lemma 2.7 case (2), we have $V_{t+1} \leq (1 - \frac{1-B}{n})V_t$. Then either by following Lemma 2.7 case (2a) we have $f(\boldsymbol{z}_{t+2}) \leq (1 - \frac{1}{n})f(\boldsymbol{z}_{t+1}) \leq (1 - \frac{1}{n})f(\boldsymbol{z}_t)$ or by following Lemma 2.7 case (2b) we have $f(\boldsymbol{z}_{t+2}) \leq V_{t+1} \leq (1 - \frac{1-B}{n})f(\boldsymbol{z}_t)$.

Overall, we have the bound $f(\boldsymbol{z}_{t+2}) \leq (1 - \frac{1-B}{n})f(\boldsymbol{z}_t)$. So, after $T = \frac{2n}{1-B}\ln(\frac{2f(\boldsymbol{z}_0)}{\epsilon})$ steps, we have

$$f(\boldsymbol{z}_T) \leq \left(1 - \frac{1-B}{n}\right)^{T/2} f(\boldsymbol{z}_0) = e^{-\frac{(1-B)T}{2n}} f(\boldsymbol{z}_0)$$
$$= e^{-\ln\left(\frac{2f(\boldsymbol{z}_0)}{\epsilon}\right)} f(\boldsymbol{z}_0) = \frac{\epsilon}{2} \implies ||\boldsymbol{z}_1|| \leq \epsilon.$$

Let us now prove the lower bound using two simple examples. In the first example, let $z_{0,i} = 1$ for all $i \in [n]$. $f(\boldsymbol{z}_0) = n$. For $\epsilon < 1$, we will need every agent to play at least once to get $f(\boldsymbol{z}_t) \leq \epsilon < 1$. So, we need at least n steps.

In the second example, let $z_{0,1} = z_{0,2} = \kappa$ for arbitrary $\kappa > 0$. $f(\boldsymbol{z}_0) = \kappa$. Also let $\beta_{t,i}(\boldsymbol{z}_t) = B \geq \frac{1}{2}$ for all t, i, and $\boldsymbol{z}_t$. It can be easily checked that the fastest way to decrease the potential is to pick agents 1 and 2 alternately. And by doing that we have $f(\boldsymbol{z}_{t+1}) = Bf(\boldsymbol{z}_t)$ for all $t \geq 1$, which implies $f(\boldsymbol{z}_t) \geq B^{t-1}\kappa$, so

$$f(\boldsymbol{z}_t) > \epsilon \impliedby B^{t-1}\kappa > \epsilon \iff t - 1 < \frac{1}{\ln(1/B)} \ln\left(\frac{\kappa}{\epsilon}\right)$$
$$\impliedby t - 1 < \frac{1}{2(1-B)} \ln\left(\frac{\kappa}{\epsilon}\right),$$

where the last inequality holds for $B \geq \frac{1}{2}$ because $2(1-B) = \ln(e^{2(1-B)}) \geq \ln(1 + 2(1-B)) \geq \ln(1 + \frac{1-B}{B}) = \ln(\frac{1}{B})$. $\square$

2.8.3 From Section 2.5

Proof of Lemma 2.10. Let τ be the time it takes to collect all n coupons in the coupon collector problem where each coupon is selected w.p. $1/n$, the following result is well-known.

Lemma 2.22. *([80]) For any constant $c > 0$, we have the following high probability bounds: (i) upper bound, $\mathbb{P}[\tau > n \ln n + cn] < e^{-c}$; (ii) lower bound, $\mathbb{P}[\tau < n \ln n - cn] < e^{-c}$.*

The analysis underlying Lemma 2.22 goes as follows: The time to collect all coupons τ can be decomposed as $\tau = \sum_i \tau_i$, where τ_i is the time it takes to collect the i-th coupon. After $(i-1)$ coupons have been collected, the probability that we get a coupon that has not yet been collected in the next time step is equal to $p_i = \frac{n-i+1}{n}$. So, τ_i corresponds to the time till the first head of a geometric random variable with parameter p_i. In particular, $p_1 = 1$, $p_2 = (n-1)/n$, $p_3 = (n-2)/n$, and so on.

For the best-response dynamics, let $T = \sum_i T_i$ denote the time it takes for every agent to play at least once, where T_i is the time between the $(i-1)$-th and the i-th unique agent. Remember that we are doing a worst-case analysis over the random selection process parameterized by L. In the first time step, we get the first unique agent w.p. 1. In the second time step, the probability of selecting a new agent is at least $(n-1)L = \frac{(n-1)Ln}{n} = Lnp_2$. Similarly, we can show that for all i, after exactly $(i-1)$ agents have played at least once, the probability that we select an agent who has not yet played in the next time step is at least Lnp_i. This implies that $\mathbb{P}[T_i \le k] \le Ln\,\mathbb{P}[\tau_i \le k]$ for all $i \in [n]$ and $k \ge 1$. Using the upper bound in Lemma 2.22, we get $\mathbb{P}[T > \frac{1}{L}(\ln n + c)] < e^{-c}$, setting $c = \ln(1/\delta)$, we get $\mathbb{P}[T \le \frac{1}{L}(\ln n + \ln(\frac{1}{\delta}))] \ge 1-\delta$ as required.

The idea for the lower bound is similar. As discussed earlier, after exactly $(i-1)$ unique agents have played, the total probability that the selection process can assign to the agents who have already played, in worst-case, is equal to $1-(n-(i-1))L = 1-nLp_i$. So, using the lower bound of Lemma 2.22, we get $\mathbb{P}[T \ge \frac{1}{L}\log(n\delta)] \ge 1-\delta$. $\square$

Proof of Lemma 2.11. Let's first show that $s_t > 0$ for $t \ge 1$. For contradiction, say $s_t = 0$ for some $t \ge 1$. Let $i = i_{t-1}$ denote the agent that moved at time $t-1$. Notice that $s_t = 0 \implies s_{t-1,-i} = s_{t,-i} \le s_t = 0$. If $s_{t-1,-i} = 0$, then as a best response to this, agent i must have played $x_{t,i} = a > 0$. So, we must have $s_t > 0$. Contradiction.

Let us now prove that for any t, if $s_t < \frac{n^2}{(n-1)c'(0)}$, then $s_{t+1} < \frac{n^2}{(n-1)c'(0)}$. If $c'(0) = 0$, e.g., when $c'(y) = y^r$ for some $r > 0$, then the inequality is satisfied trivially because $\frac{n^2}{(n-1)c'(0)} = \infty$. Assuming $c'(0) > 0$, we have the following two cases

- if $s_{t,-i_t} = 0$, then $s_{t+1} = x_{t+1,i_t} = a < 1 < \frac{n^2}{(n-1)c'(0)}$;
- if $0 < s_{t,-i_t} \le s_t < \frac{n^2}{(n-1)c'(0)}$, then from the first-order condition we have

$$\begin{aligned}
&\frac{s_{t,-i_t}}{(x_{t+1,i_t} + s_{t,-i_t})^2} - \frac{n-1}{n^2} c'(x_{t+1,i_t}) = 0 \\
&\implies \frac{s_{t,-i_t}}{(x_{t+1,i_t} + s_{t,-i_t})^2} = \frac{s_{t,-i_t}}{s_{t+1}^2} = \frac{n-1}{n^2} c'(x_{t+1,i_t}) \ge \frac{n-1}{n^2} c'(0) \\
&\implies s_{t+1} \le \sqrt{s_{t,-i_t} \frac{n^2}{(n-1)c'(0)}} < \frac{n^2}{(n-1)c'(0)}.
\end{aligned}$$

Let us now assume that at time t, $x_{t,i} > 0$, $x_{t,j} > 0$, and $s_t < \frac{n^2}{(n-1)c'(0)}$. Agent i_t makes the move at time t. We know that $s_{t,-i_t} \ge \min(x_{t,i}, x_{t,j}) > 0$ and $s_{t,-i_t} < s_t < \frac{n^2}{(n-1)c'(0)}$. Let us look at the first-order condition for agent i_t,

$$\frac{s_{t,-i_t}}{(x_{t+1,i_t} + s_{t,-i_t})^2} - \frac{n-1}{n^2} c'(x_{t+1,i_t}) = 0.$$

If $c'(0) = 0$, then

$$\frac{s_{t,-i_t}}{(0 + s_{t,-i_t})^2} - \frac{n-1}{n^2} c'(0) = \frac{1}{s_{t,-i_t}} > 0,$$

so the BR must be strictly positive, i.e., $x_{t+1,i_t} > 0$, to satisfy the first-order condition (as the utility function is concave and its derivative decreasing). On the other hand, if $c'(0) > 0$, as $s_{t,-i_t} < \frac{n^2}{(n-1)c'(0)}$, we have

$$s_{t,-i_t} < \frac{n^2}{(n-1)c'(0)} \iff \frac{s_{t,-i_t}}{(0 + s_{t,-i_t})^2} - \frac{n-1}{n^2} c'(0) > 0.$$

The same argument as before implies that $x_{t+1,i_t} > 0$ to ensure that the first-order condition is satisfied with equality. □

Proof of Lemma 2.12. Let T_1 be a random variable that denotes the time it takes for every agent to make at least one move. We next prove that conditions (1) and (2) of the warm-up phase (Definition 2.5) are satisfied for $t \ge T_1$.

Condition (1). Notice that Lemma 2.11 (1) implies $s_t > 0$ for $t \ge T_1 \ge n \ge 1$. For conciseness, let $\kappa = \frac{n^2}{(n-1)c'(0)}$. Let us now prove that $s_t < \kappa$ for $t \ge T_1$. Note that it is enough to prove $s_{T_1} < \kappa$ because $s_{T_1} < \kappa$ implies $s_t < \kappa$ for all $t \ge T_1$ using Lemma 2.11 (2). For contradiction, let us assume that $s_{T_1} \ge \kappa$. Let $t = T_1 - 1$.

- If $s_{t,-i_t} = 0$, then $s_{T_1} = s_{t+1} = x_{t+1,i_t} = a < \kappa$, which contradicts $s_{T_1} \geq \kappa$.

- If $0 < s_{t,-i_t} < \kappa$, then applying the first-order condition for agent i_t, we have $\frac{s_{t,-i_t}}{s_{t+1}^2} - \frac{n-1}{n^2} c'(x_{t+1,i_t}) = 0 \implies s_{t+1} < \sqrt{\kappa \frac{n^2}{(n-1)c'(0)}} = \kappa$. So, $s_{T_1} = s_{t+1} < \kappa$, which also leads to contradiction.

- If $s_{t,-i_t} \geq \kappa$, then $x_{t+1,i_t} = 0$ and $s_t \geq s_{t,-i_t} \geq \kappa$. Doing induction in the backward direction, starting from $\tau = t = T_1 - 1$ and going to $\tau = 0$, $x_{\tau+1,i_\tau} = 0$ for every $\tau < T_1$. As every agent i makes at least one move before T_1, so every agent $i \in [n]$ has $x_{T_1,i} = 0$. So, $s_{T_1} = 0$. Contradiction.

This completes the proof for condition (1) of the warm-up phase.

Condition (2). Let us now prove that $x_{t,i} \leq \frac{n^2}{4(n-1)}$ for all $t \geq T_1$ and $i \in [n]$. Fix a $t \geq T_1$ and an $i \in [n]$. As every agent has played at least once before T_1, so i must have made a move before t. Let $\tau < t$ denote the most recent time before t when agent i made the best-response move. By definition of τ, i did not play between $\tau + 1$ and t, so $x_{t,i} = x_{\tau+1,i}$. Finally, $x_{\tau+1,i} \leq \frac{n^2}{4(n-1)}$ because:

- If $s_{\tau,-i} = 0$, then $x_{\tau+1,i} = a \leq \frac{n^2}{4(n-1)}$.

- If $0 < s_{\tau,-i} < \frac{n^2}{(n-1)c'(0)}$, then we know that $0 < x_{\tau+1,i} \leq s_{\tau+1} < \frac{n^2}{(n-1)c'(0)}$ from Lemma 2.11. If $x_{\tau+1,i} \leq 1$, then $x_{\tau+1,i} \leq \frac{n^2}{4(n-1)}$ because $n \geq 3$. On the other hand, if $x_{\tau+1,i} \geq 1$, then $c'(x_{\tau+1,i}) \geq 1$, and using the first-order optimality condition we have

$$\frac{s_{\tau,-i}}{(x_{\tau+1,i} + s_{\tau,-i})^2} = \frac{n-1}{n^2} c'(x_{\tau+1,i}) \geq \frac{n-1}{n^2} c'(1) = \frac{n-1}{n^2}$$
$$\implies x_{\tau+1,i} \leq \frac{n}{\sqrt{n-1}} \sqrt{s_{\tau,-i}} - s_{\tau,-i} \leq \max_{z \geq 0} \left(\frac{n}{\sqrt{n-1}} z - z^2 \right).$$

  As $\frac{n}{\sqrt{n-1}} z - z^2$ is maximized at $z = \frac{n}{2\sqrt{n-1}}$, so $x_{\tau+1,i} \leq \frac{n^2}{4(n-1)}$ as required.

This completes the proof for condition (2) of the warm-up phase.

We have shown that $s_t > 0$ for $t \geq T_1$, so there must be at least one agent with positive output for all $t \geq T_1$. Let i be an agent that has positive output $x_{T_1,i} > 0$ at time T_1. We also know from condition (1), which we proved earlier, that $x_{T_1,i} \leq s_{T_1} < \frac{n^2}{(n-1)c'(0)}$. Although $x_{T_1,i} > 0$, it is possible that i is the only agent with positive output and every other agent $j \neq i$ has $x_{T_1,j} = 0$. We next resolve this scenario.

Condition (3). Let T_2 denote the additional steps after T_1, if any, required to get at least two agents with positive output. Notice that T_2 can be upper bounded by the time it takes to get the first head (select an agent $j \neq i$) of a geometric random variable with parameter $p \geq (n-1)L$ because each $j \neq i$ is assigned a probability of at least L at each time step. When a $j \neq i$ is selected at time $t = T_2 - 1$, then $x_{T_2,j} = x_{t+1,j} \in (0, \frac{n^2}{(n-1)c'(0)})$, as required, using Lemma 2.11 (3). Further, Lemma 2.11 (3) also implies that there will always be at least two agents with positive output this time onward. So, for $t \geq T_1 + T_2$, the action profile x_t satisfies condition (3) of the warm-up phase.

To summarize, for $t \geq T_1 + T_2$, the action profile x_t satisfies all conditions required for the completion of the warm-up phase (Definition 2.5). So, $T_{warm} \leq T_1 + T_2$. Let us now prove a high probability upper bound on $T_1 + T_2$, which holds for T_{warm} as well. From Lemma 2.10, we know that $\mathbb{P}[T_1 > \frac{1}{L}\ln(\frac{n}{\delta'})] < \delta'$ for any $\delta' \in (0,1)$. Set $\delta' = \frac{\delta}{2}$, we get $\mathbb{P}[T_1 > \frac{1}{L}\ln(\frac{2n}{\delta})] < \frac{\delta}{2}$. As T_2 underestimates the time till the first head for a geometric random variable with parameter $(n-1)L$, we have $\mathbb{P}[T_2 > k] \leq (1-(n-1)L))^k < e^{-(n-1)Lk}$. Setting $k = \frac{1}{(n-1)L}\ln(\frac{2}{\delta})$, we get $\mathbb{P}[T_2 > \frac{1}{(n-1)L}\ln(\frac{2}{\delta})] < \frac{\delta}{2}$. Using union bound, we get $T_{warm} \leq T_1 + T_2 \leq \frac{1}{L}\ln(\frac{2n}{\delta}) + \frac{1}{(n-1)L}\ln(\frac{2}{\delta}) = O(\frac{1}{L}\log(\frac{n}{\delta}))$ w.p. $1-\delta$ as required. $\square$

Proof of Lemma 2.13. Let $i = i_t$. Notice that $z_{t+1,i} = 0 \iff x_{t+1,i} = 1$. If $x_{t+1,i} = 1$, the first order condition for agent i is

$$\frac{s_{t,-i}}{(x_{t+1,i} + s_{t,-i})^2} = \frac{n-1}{n^2}c'(x_{t+1,i}) \iff \frac{s_{t,-i}}{(1+s_{t,-i})^2} = \frac{n-1}{n^2}$$
$$\iff (n-1)(1+s_{t,-i})^2 - n^2 s_{t,-i} = 0.$$

The two solutions for the quadratic equation above are

$$\begin{aligned}
s_{t,-i} = n-1 &\implies (n-1)(1+s_{t,-i})^2 - n^2 s_{t,-i} \\
&= (n-1)(1+n-1)^2 - n^2(n-1) = 0, \\
s_{t,-i} = \frac{1}{n-1} &\implies (n-1)(1+s_{t,-i})^2 - n^2 s_{t,-i} \\
&= (n-1)\left(1+\frac{1}{n-1}\right)^2 + n^2\frac{1}{n-1} = (n-1)\left(\frac{n}{n-1}\right)^2 + n^2\frac{1}{n-1} = 0.
\end{aligned}$$

Further, when $s_{t,-i} \in (\frac{1}{n-1}, n-1)$, then $\frac{s_{t,-i}}{(1+s_{t,-i})^2} > \frac{n-1}{n^2}$, which implies that the BR to $s_{t,-i}$ is strictly more than 1. Similarly, we can verify that if $s_{t,-i} \in (0, \frac{1}{n-1})$ or $s_{t,-i} > n-1$, then $x_{t+1,i} < 1$. Writing this in terms of z and σ, we have

- $s_{t,-i} \in [\frac{1}{n-1}, n-1] \iff \sigma_{t,-i} \in [\frac{1}{n-1} - (n-1), 0]$, then $x_{t+1,i} \geq 1 \iff z_{t+1,i} \geq 0$, and

- $s_{t,-i} \geq n-1 \iff \sigma_{t,-i} \geq 0$, then $x_{t+1,i} \leq 1 \iff z_{t+1,i} \leq 0$.

So, we have shown that if $s_{t,-i} \geq \frac{1}{n-1}$, then $z_{t+1,i}$ has the opposite sign as $\sigma_{t,-i}$, as required. We now upper bound the ratio $\frac{-z_{t+1,i}}{\sigma_{t,-i}}$.

Let $r = \frac{-z_{t+1,i}}{\sigma_{t,-i}}$, we want to upper bound r. Let us write the first order condition for agent i using z and σ

$$\frac{s_{t,-i}}{(x_{t+1,i}+s_{t,-i})^2} = \frac{n-1}{n^2}c'(x_{t+1,i}) \iff \frac{n-1+\sigma_{t,-i}}{(n+z_{t+1,i}+\sigma_{t,-i})^2} = \frac{n-1}{n^2}c'(x_{t+1,i}).$$

If $x_{t+1,i} > 1 \iff z_{t+1,i} > 0$, then $c'(x_{t+1,i}) \geq c'(1) = 1$. Also, $z_{t+1,i} > 0 \implies \sigma_{t,-i} < 0$. Using these, we get

$$\begin{aligned}
&\frac{n-1+\sigma_{t,-i}}{(n+z_{t+1,i}+\sigma_{t,-i})^2} = \frac{n-1+\sigma_{t,-i}}{(n+(1-r)\sigma_{t,-i})^2} \geq \frac{n-1}{n^2}\\
&\implies 1+\frac{\sigma_{t,-i}}{n-1} \geq \left(1+\frac{(1-r)\sigma_{t,-i}}{n}\right)^2 \implies \frac{(1-r)\sigma_{t,-i}}{n} \leq \sqrt{1+\frac{\sigma_{t,-i}}{n-1}}-1\\
&\implies 1-r \geq \frac{n}{\sigma_{t,-i}}\left(\sqrt{1+\frac{\sigma_{t,-i}}{n-1}}-1\right) \text{ as } \sigma_{t,-i} < 0.
\end{aligned}$$

Let $g(y) = \frac{\sqrt{1+y}-1}{y}$. Notice that $1-r \geq \frac{n}{n-1}g(\frac{\sigma_{t,-i}}{n-1})$. As $\sigma_{t,-i} < 0$, so $\frac{\sigma_{t,-i}}{n-1} < 0$. Further, as $\sigma_{t,-i} \geq -(n-1)+\frac{1}{n-1}$, so $\frac{\sigma_{t,-i}}{n-1} > -1$. We will later lower bound $g(y)$ in $y \in (-1, 0]$, but before that let us look at the case $x_{t+1,i} < 1$.

If $x_{t+1,i} < 1 \iff z_{t+1,i} < 0$, then $c'(x_{t+1,i}) \leq c'(1) = 1$. Also, as $z_{t+1,i} < 0 \implies \sigma_{t,-i} > 0$. Using these, we get

$$\begin{aligned}
&\frac{n-1+\sigma_{t,-i}}{(n+(1-r)\sigma_{t,-i})^2} \leq \frac{n-1}{n^2} \implies 1+\frac{\sigma_{t,-i}}{n-1} \leq \left(1+\frac{(1-r)\sigma_{t,-i}}{n}\right)^2\\
&\implies \frac{(1-r)\sigma_{t,-i}}{n} \geq \sqrt{1+\frac{\sigma_{t,-i}}{n-1}}-1 \implies 1-r \geq \frac{n}{\sigma_{t,-i}}\left(\sqrt{1+\frac{\sigma_{t,-i}}{n-1}}-1\right),
\end{aligned}$$

where the last inequality holds because $\sigma_{t,-i} > 0$. Again, for $g(y) = \frac{\sqrt{1+y}-1}{y}$, we have the same inequality as before $1-r \geq \frac{n}{n-1}g(\frac{\sigma_{t,-i}}{n-1})$. Now, if $\sigma_{t,-i} \geq 2$, we trivially have $r = \frac{-z_{t+1,i}}{\sigma_{t,-i}} \leq \frac{1}{2}$ because $z_{t+1,i} \geq -1$ always. So, we can assume $\sigma_{t,-i} \leq 2 \implies \frac{\sigma_{t,-i}}{n-1} \leq \frac{2}{n-1}$. Earlier, for $z_{t+1,i} > 0$, we needed to lower bound $g(y)$ in domain $y \in [-1, 0]$, and now for $z_{t+1,i} < 0$, we need to lower bound $g(y)$ in domain $[0, \frac{2}{n-1}]$.

Let us differentiate $g(y) = \frac{\sqrt{1+y}-1}{y}$ for $y > -1$, we get

$$\begin{aligned}
g'(y) &= \frac{1}{2y\sqrt{1+y}} - \frac{\sqrt{1+y}-1}{y^2} = \frac{y-2\sqrt{1+y}(\sqrt{1+y}-1)}{2y^2\sqrt{1+y}}\\
&= \frac{y-2(1+y)+2\sqrt{1+y}}{2y^2\sqrt{1+y}} = \frac{-1-(1+y)+2\sqrt{1+y}}{2y^2\sqrt{1+y}} = \frac{-(1-\sqrt{1+y})^2}{2y^2\sqrt{1+y}} < 0.
\end{aligned}$$

So, $g(y)$ is a decreasing function for all $y > -1$. Therefore, the minimum value of $g(y)$ in domain $[0, \frac{2}{n-1}]$ occurs at $y = \frac{2}{n-1}$. Plugging it in, we get

$$1 - r \geq \frac{n}{n-1} g\left(\frac{\sigma_{t,-i}}{n-1}\right) \geq \frac{n}{n-1} g\left(\frac{2}{n-1}\right) = \frac{n}{2}\left(\sqrt{1 + \frac{2}{n-1}} - 1\right).$$

Let us now lower bound $\frac{n}{2}\left(\sqrt{1 + \frac{2}{n-1}} - 1\right)$. Let $v = 1 + \frac{2}{n-1} \iff n = 1 + \frac{2}{v-1} = \frac{v+1}{v-1}$. Notice that $v \geq 1$ as $n \geq 2$, and as $n \to \infty$, $v \to 1$. Let $h(v) = \frac{n}{2}\left(\sqrt{1 + \frac{2}{n-1}} - 1\right) = \frac{v+1}{2(v-1)}(\sqrt{v} - 1) = \frac{v+1}{2(\sqrt{v}+1)}$. We need to lower bound $h(v)$ for $v > 1$. Differentiating $h(v)$ w.r.t. v, we have

$$\begin{aligned} h'(v) = \frac{1}{2(\sqrt{v}+1)} &- \frac{v+1}{2} \frac{1}{(\sqrt{v}+1)^2} \frac{1}{2\sqrt{v}} \\ &= \frac{2\sqrt{v}(\sqrt{v}+1) - (v+1)}{4\sqrt{v}(\sqrt{v}+1)^2} = \frac{v + 2\sqrt{v} - 1}{4\sqrt{v}(\sqrt{v}+1)^2} > 0, \text{ as } v > 1. \end{aligned}$$

So, the minimum value of $h(v)$ in domain $v \in (1, \infty)$ occurs at $v = 1$, and we have $h(v) \geq h(1) = \frac{1}{2}$. So, we have proven that $1 - r \geq \frac{1}{2} \implies r \leq \frac{1}{2}$. □

Proof of Lemma 2.14. Note that we are assuming the completion of the warm-up phase. Let $\boldsymbol{x}_t$ have at least two agents i and $j \neq i$ with $x_{t,i} \geq \frac{1}{n-1}$ and $x_{t,j} \geq \frac{1}{n-1}$. We shall prove that $\boldsymbol{x}_{t+1}$ also has two agents with output at least $\frac{1}{n-1}$. By induction, this implies the lemma.

If an agent $i_t \notin \{i, j\}$ makes the move at time t. Then $x_{t+1,i} = x_{t,i} \geq \frac{1}{n-1}$ and $x_{t+1,j} = x_{t,j} \geq \frac{1}{n-1}$, and we are done trivially. So, let us assume that one of i or j, w.l.o.g. say i, makes the move at time t. If there is an agent $k \notin \{i, j\}$ with $x_{t,i} \geq \frac{1}{n-1}$, then again we are done trivially because j and k will have output at least $\frac{1}{n-1}$ at time $t+1$. So, let us assume that $x_{t,k} < \frac{1}{n-1}$ for all $k \notin \{i, j\}$.

Now, as the warm-up phase has completed, so $x_{t,j} \leq \frac{n^2}{4(n-1)}$. Therefore, $s_{t,-i} = x_{t,j} + \sum_{k \notin \{i,j\}} x_{t,k} \leq \frac{n^2}{4(n-1)} + \frac{n-2}{n-1} \leq \frac{n^2}{2(n-1)}$. Also, $s_{t,-i} \geq x_{t,j} \geq \frac{1}{n-1}$. Let us now look at the first order condition for agent i. If $x_{t+1,i} \geq 1$, then we are done trivially again. So, assume $x_{t+1,i} \leq 1$, which implies $c'(x_{t+1,i}) \leq c'(1) = 1$. We have

$$\frac{s_{t,-i}}{(x_{t+1,i} + s_{t,-i})^2} = \frac{n-1}{n^2} c'(x_{t+1,i}) \leq \frac{n-1}{n^2} \implies x_{t+1,i} \geq \sqrt{\frac{n^2 s_{t,-i}}{n-1}} - s_{t,-i}.$$

Notice that the function $g(y) = \sqrt{\frac{n^2 y}{n-1}} - y$ is concave in y, so the minimum values occur at extreme points. As $s_{t,-i} \in [\frac{1}{n-1}, \frac{n^2}{2(n-1)}]$, plugging in the extreme points we get:

$$g\left(\frac{1}{n-1}\right) \geq \sqrt{\frac{n^2}{(n-1)^2}} - \frac{1}{n-1} = \frac{n}{n-1} - \frac{1}{n-1} = 1 \geq \frac{1}{n-1},$$
$$g\left(\frac{n^2}{2(n-1)}\right) \geq \sqrt{\frac{n^2 n^2}{(n-1)2(n-1)}} - \frac{n^2}{2(n-1)} = \frac{n^2}{n-1}\frac{\sqrt{2}-1}{2} \geq \frac{1}{n-1},$$

where the last inequality holds for all $n \geq 3$. So, we get $x_{t+1,i} \geq \frac{1}{n-1}$ as required. □

Proof of Lemma 2.15. Let us look at the sequence of moves after the warm-up phase. Let $t \geq T_{warm}$ and $i = i_t$, and let's assume that there are less than two agents with output at least $\frac{1}{n-1}$ at time t. Then, we have the following cases based on the value of $s_{t,-i}$.

Let $s_{t,-i} \geq \frac{1}{n-1}$. Then following exactly the same argument as the proof of Lemma 2.14, we have $x_{t+1,i} \geq \frac{1}{n-1}$. Now, at time $t+1$, if we have two agents with output $\geq \frac{1}{n-1}$, we are done. Else, let $j \neq i$ be the agent who plays the next BR move, say at time $\tau \geq t+1$ (ignoring the redundant consecutive BR moves by i that do not change the output profile). Again, following exactly the same argument as the proof of Lemma 2.14, we have $x_{\tau+1,j} \geq \frac{1}{n-1}$. Finally, the time taken to get this non-redundant move by an agent $j \neq i$ is at most $m = \frac{1}{(n-1)L}\ln(\frac{1}{\delta})$ w.p. $1-\delta$ because the probability that agent i makes m consecutive moves is $\leq (1-(n-1)L)^m \leq e^{-(n-1)Lm} \leq \delta$.

Let $s_{t,-i} < \frac{1}{n-1}$. Then the two possible scenarios are either $x_{t+1,i} \geq \frac{1}{n-1}$ or $x_{t+1,i} < \frac{1}{n-1} \leq 1$. If $x_{t+1,i} \leq 1$, then $c'(x_{t+1,i}) \leq c'(1) = 1$, and the first-order condition implies

$$\frac{s_{t,-i}}{(x_{t+1,i}+s_{t,-i})^2} = \frac{n-1}{n^2}c'(x_{t+1,i}) \leq \frac{n-1}{n^2} \implies x_{t+1,i} \geq \sqrt{\frac{n^2 s_{t,-i}}{n-1}} - s_{t,-i}$$
$$\implies x_{t+1,i} \geq \sqrt{s_{t,-i}}\left(\frac{n}{\sqrt{n-1}} - \sqrt{s_{t,-i}}\right) \geq \sqrt{s_{t,-i}}\left(\frac{n}{\sqrt{n-1}} - \frac{1}{\sqrt{n-1}}\right)$$
$$\implies x_{t+1,i} \geq \sqrt{(n-1)s_{t,-i}}.$$

Now, let $j \neq i$ be the agent who plays the next BR move, say at time τ (ignoring the redundant consecutive BR moves by i for now). At τ, $s_{\tau,-j} = s_{t+1,-j} \geq x_{t+1,i} \geq \sqrt{(n-1)s_{t,-i}}$.

Let $(y_t)_{t\geq 0}$ be a sequence where $y_0 = \gamma = \min_{j\in[n]} s_{T_{warm},-j} < \frac{1}{n-1}$ and $y_{t+1} \geq \sqrt{(n-1)y_t}$. Notice that y_t tracks the evolution of $s_{t,-i_t}$ for $t \geq T_{warm}$ assuming there are no redundant BR moves. By the definition of T_{warm}, there are at least two agents with strictly positive output, so $\gamma > 0$. We want to measure the time it takes for y_t to reach $\frac{1}{n-1}$.

$$y_t \geq \sqrt{(n-1)y_{t-1}} \geq (n-1)^{\frac{1}{2}+\frac{1}{4}} y_{t-2}^{\frac{1}{4}} \geq \ldots \geq (n-1)^{1-\frac{1}{2^t}} y_0^{\frac{1}{2^t}} = (n-1)^{1-\frac{1}{2^t}} \gamma^{\frac{1}{2^t}}.$$

We want to ensure $y_t \geq \frac{1}{n-1}$, which is implied by

$$(n-1)^{1-\frac{1}{2^t}} \gamma^{\frac{1}{2^t}} \geq \frac{1}{n-1} \impliedby \gamma^{\frac{1}{2^t}} \geq \frac{1}{2} \impliedby \frac{1}{2^t} \lg(\gamma) \geq -1 \impliedby t \geq \lg\lg(\frac{1}{\gamma}).$$

So, we need $\kappa = O(\log\log(\frac{1}{\gamma}))$ non-redundant moves. Now, let's provide a high probability bound on the time it takes to have κ non-redundant moves. At each time step, we make a non-redundant move with probability at least $(n-1)L$. So, we need to find the time it takes to get κ heads of a geometric random variable when the probability of getting a head is $p \geq (n-1)L$.

Let $m = \frac{2}{p}\max(\kappa, \frac{1}{p}\ln(\frac{1}{\delta}))$. Let Z_i be a random variable that takes value -1 (head) w.p. p and 0 (tail) w.p. $1-p$. $\mathbb{E}[\sum_{i=1}^m Z_i] = -pm$. Less than κ heads (-1 values) corresponds to having $\sum_{i=1}^m Z_i > -\kappa$. Using Hoeffding's inequality, we have

$$\begin{aligned}\mathbb{E}[\sum_{i=1}^m Z_i > -\kappa] &= \mathbb{E}[\sum_{i=1}^m Z_i - \mathbb{E}[\sum_{i=1}^m Z_i] > -\kappa + pm] \\ &\leq e^{\frac{-2(pm-\kappa)^2}{m}} \leq e^{\frac{-2(pm-\frac{pm}{2})^2}{m}} \leq e^{\frac{-mp^2}{2}} \leq \delta.\end{aligned}$$

So, after $\frac{2}{(n-1)L}\max(\kappa, \frac{1}{(n-1)L}\ln(\frac{1}{\delta}))$ steps, the probability of getting less than κ non-redundant moves is bounded above by δ, as required. $\square$

Proof of Lemma 2.16. If $T_{warm} = 0$, i.e., all required conditions for completion of the warm-up phase are satisfied by the initial state x_0, then we trivially get $\gamma = \min_j s_{0,-j} \geq \min(\mathcal{A}) \geq \gamma_{lb}$.

Let us assume that $T_{warm} > 0$. Let us focus on the time step $T_{warm} - 1$. Let $t = T_{warm} - 1 \iff T_{warm} = t + 1$.

As $t+1$ is the smallest time when all conditions for completion of the warm-up phase, Definition 2.5, are satisfied, therefore at time t, there must be at least one violation. We do a case analysis depending upon which condition was violated at t. Let $\kappa = \frac{n^2}{(n-1)c'(0)}$ for conciseness. Let $i = i_t$ be the agent who makes the transition at time t.

Case 1: Definition 2.5 condition (3) violated at t, i.e., there is only one agent j with $x_{t,j} > 0$.

First, notice that at time t an agent $i \neq j$ makes a transition to a positive $x_{t+1,i}$ to satisfy all three conditions of Definition 2.5. Check that the conditions (2) and (1) of Definition 2.5 must not have been violated at time t because, then, in a single step, we could not have satisfied all three conditions. This implies that $s_t = x_{t,j} \leq \frac{n^2}{4(n-1)} < \frac{n^2}{(n-1)c'(0)} = \kappa$.

Let us now trace our steps back from t to 0. We claim that all transitions before time t were made by agent j, i.e., for every $\tau < t$, $j = i_\tau$. If not, let $\tau < t$ be the most recent transition by an agent $k \neq j$. As $x_{\tau+1,k} = x_{t,k} = 0$, therefore $s_{\tau,-k} \geq \kappa$, but $s_{\tau,-k} = s_{t,-k} = x_{t,j} < \kappa$. Contradiction.

As j makes all the transitions before time t, we have either (i) if $t = 0$, then $x_{t,j} = x_{0,j}$; or (ii) if $t > 0$, then $x_{t,j} = a$ as a response to 0 output by everyone else. Notice that $\gamma \geq \min(x_{t+1,i}, x_{t+1,j})$. Further, we have either $x_{t+1,i} \geq 1$ or $x_{t+1,i} \leq 1 \implies c'(x_{t+1,i}) \leq 1$ and from the first-order condition

$$\begin{aligned} x_{t+1,i} \geq \sqrt{s_{t,-i}} \left(\frac{n}{\sqrt{n-1}} - \sqrt{s_{t,-i}} \right) &= \sqrt{x_{t,j}} \left(\frac{n}{\sqrt{n-1}} - \sqrt{x_{t,j}} \right) \\ &\geq \sqrt{x_{t,j}} \left(\frac{n}{\sqrt{n-1}} - \frac{n}{2\sqrt{n-1}} \right) \geq \sqrt{x_{t,j}} \geq \min(1, x_{t,j}). \end{aligned}$$

So, $\gamma \geq \min(\{a\} \cup \mathcal{A})$.

Case 2: Definition 2.5 condition (1) violated at t, i.e., total output $s_t \geq \kappa$, but condition (3) is satisfied.

Let i be the agent that makes the move at time t to decrease the total output from $s_t \geq \kappa$ to $s_t < \kappa$. We argued in Case 1 that condition (3) must have been satisfied at t. So, there are at least two agents with strictly positive output at t. Let $j \neq i$ be an agent other than i that has $x_{t,j} > 0$. We claim that $x_{t,j} = x_{0,j}$.

We prove our claim by contradiction. Let us trace our steps back from t. Let agent k make the transition at $t-1$. We show that $k \notin \{i, j\}$. Notice that $x_{t,k} = 0$ because if $x_{t,k} > 0$ then either: (i) $s_{t-1,-k} = 0$, but this is not possible as $s_{t-1,-k} \geq \min(x_{t,i}, x_{t,j}) > 0$; (ii) $0 < s_{t-1,-k} < \kappa$, but this is also not possible because then $s_t < \kappa$ by Lemma 2.11, but $s_t \geq \kappa$. As we already know that $x_{t,i} > 0$ and $x_{t,j} > 0$, so $k \notin \{i, j\}$. Repeating the same argument, for every $\tau < t$, we can show that $i_\tau \notin \{i, j\}$, which implies $x_{t+1,j} = x_{t,j} = x_{0,j}$. Using same argument as done for Case 1, we can show that $x_{t+1,i} \geq \min(1, x_{t,j})$ and $\gamma \geq \min(\{a\} \cup \mathcal{A})$.

Case 3: Definition 2.5 condition (2) violated at t, i.e., there is an agent i with $x_{t,i} > \frac{n^2}{4(n-1)}$, but conditions (3) and condition (1) are satisfied. After agent i made the move at time t, $x_{t+1,i} \le \frac{n^2}{4(n-1)}$ at time $t+1$.

We have $s_{t,-i} > 0$ (by condition (3)) and $s_{t,-i} \le s_{t+1} \le \frac{n^2}{2(n-1)}$ because $\gamma = \min_j s_{t+1,-j} < \frac{1}{n-1}$, and which implies that $s_{t+1} \le \gamma + \max_j x_{t+1,j} \le \frac{1}{n-1} + \frac{n^2}{4(n-1)} \le \frac{n^2}{2(n-1)}$.

We can lower bound γ as a function of $s_{t+1,-i}$. First, notice that $\gamma \ge \min(s_{t+1,-i}, x_{t+1,i}) = \min(s_{t,-i}, x_{t+1,i})$. Further, we can lower bound $x_{t+1,i}$ as $x_{t+1,i} \ge \min(1, \frac{s_{t,-i}}{2})$ because, if $x_{t+1,i} \le 1$, then from the first-order condition we have

$$x_{t+1,i} \ge \sqrt{s_{t,-i}} \left(\frac{n}{\sqrt{n-1}} - \sqrt{s_{t,-i}} \right) \ge \sqrt{s_{t,-i}} \left(\frac{n}{\sqrt{n-1}} - \frac{n}{\sqrt{2}\sqrt{n-1}} \right) \ge \frac{\sqrt{s_{t,-i}}}{2} \ge \min\left(1, \frac{s_{t,-i}}{2}\right).$$

Therefore, let us focus on lower bounding $s_{t,-i}$.

If $t = 0$, then we trivially have $s_{t,-i} = \sum_{j \ne i} x_{0,j} \ge \min(\mathcal{A})$ (and we know that $s_{t,-i} > 0$, so there is an agent $j \ne i$ with positive $x_{0,j}$).

Let $t > 0$. In the proof of Lemma 2.12, we argued that if an agent j plays at least one move before time τ, for any τ, then $x_{\tau,j} \le \frac{n^2}{4(n-1)}$. As $x_{t,i} > \frac{n^2}{4(n-1)}$, so $i_\tau \ne i$ for all $\tau < t$, which implies $x_{\tau,i} = x_{0,i}$ for all $\tau \le t$.

Let $j = i_{t-1} \ne i$ be the agent who made the move at time $t-1$. Let $\beta = \sum_{k \ne i,j} x_{t-1,k}$. So, $s_{t-1,-j} = \beta + x_{t-1,i} = \beta + x_{0,i}$ and $s_{t,-i} = \beta + x_{t,j}$. Let's do a case analysis on the value of β and $x_{t,j}$

- Let $x_{t,j} \ge 1$ or $\beta \ge 1$. Then $s_{t,-i} = \beta + x_{t,j} \ge 1$.
- Let $x_{t,j} < 1$ and $\beta < 1$. The first-order condition for agent j is

  $$\frac{s_{t-1,-j}}{(x_{t,j} + s_{t-1,-j})^2} = \frac{n-1}{n^2} c'(x_{t,j}).$$

  Notice that $\frac{s_{t-1,-j}}{(x_{t,j}+s_{t-1,-j})^2}$ decreases with $s_{t-1,-j}$ because the derivative w.r.t. $s_{t-1,-j}$ is

  $$\frac{\partial}{\partial s_{t-1,-j}} \left(\frac{s_{t-1,-j}}{(x_{t,j} + s_{t-1,-j})^2} \right) = \frac{(x_{t,j} + s_{t-1,-j}) - 2s_{t-1,-j}}{(x_{t,j} + s_{t-1,-j})^3} < 0,$$

  as $x_{t,j} < 1$ and $s_{t-1,-j} \ge x_{0,i} \ge \frac{n^2}{4(n-1)} \ge 1$. This implies that if β increases, then $s_{t-1,-j}$ increases, then $x_{t,j}$ decreases.

  If $\kappa = \infty$ (i.e., $c'(0) = 0$), then we get a lower bound $x_{t,j} = BR(x_{0,i} + \beta) \ge BR(x_{0,i} + 1) > 0$. On the other hand, if $\kappa < \infty$, then either $\beta \ge (\kappa - x_{0,i})/2 >$

0 or $x_{t,j} = BR(x_{0,i} + \beta) \geq BR((\kappa + x_{0,i})/2) > 0$. So, $s_{t,-i} = \beta + x_{t,j} \geq \max(\beta, x_{t,j}) \geq \min((\kappa - x_{0,i})/2, BR((\kappa + x_{0,i})/2))$.

$\square$

Proof of Lemma 2.17. Fix an arbitrary agent i. Let $\widehat{\epsilon} = 2(1 + K)\epsilon$ Let $u_+ = u_i(BR(s_{-i}), s_{-i})$ and $u_- = u_i(x_i, s_{-i})$. We want to prove that

$$u_- \geq (1 - \widehat{\epsilon})u_+ \iff \frac{u_-}{u_+} \geq 1 - \widehat{\epsilon}.$$

As $||\boldsymbol{x} - \boldsymbol{x}^*||_1 \leq \epsilon$, we have $|s_{-i} - s^*_{-i}| \leq \epsilon \implies s_{-i} \in [(n-1) - \epsilon, (n-1) + \epsilon]$. Also, as $||\boldsymbol{x} - \boldsymbol{x}^*||_1 \leq \epsilon$, we have $x_i \in [1 - \epsilon, 1 + \epsilon]$, and as $|BR(s_{-i}) - x^*_i| \leq \epsilon$, we have $BR(s_{-i}) \in [1 - \epsilon, 1 + \epsilon]$. We can lower bound u_- as

$$\begin{aligned}
u_- = u_i(x_i, s_{-i}) &= \frac{x_i}{x_i + s_{-i}} - \frac{n-1}{n^2}c(x_i) \\
&\geq \min_{\substack{y \in [1-\epsilon, 1+\epsilon] \\ z \in [(n-1)-\epsilon, (n-1)+\epsilon]}} \left(\frac{y}{y+z} - \frac{n-1}{n^2}c(y) \right) \\
&\geq \min_{\substack{y \in [1-\epsilon, 1+\epsilon] \\ z \in [(n-1)-\epsilon, (n-1)+\epsilon]}} \frac{y}{y+z} - \max_{y \in [1-\epsilon, 1+\epsilon]} \frac{n-1}{n^2}c(y) \\
&\geq \min_{y \in [1-\epsilon, 1+\epsilon]} \frac{y}{y + (n-1) + \epsilon} - \frac{n-1}{n^2}c(1+\epsilon) \\
&\geq \frac{1-\epsilon}{n} - \frac{n-1}{n^2}(c(1) + K\epsilon) \\
&\geq \frac{n - (n-1)c(1)}{n^2} - \frac{\epsilon(n + (n-1)K)}{n^2}
\end{aligned}$$

Similarly, we can upper bound u_+ as

$$\begin{aligned}
u_+ = u_i(BR(s_{-i}), s_{-i}) &= \frac{BR(s_{-i})}{BR(s_{-i}) + s_{-i}} - \frac{n-1}{n^2}c(BR(s_{-i})) \\
&\leq \max_{\substack{y \in [1-\epsilon, 1+\epsilon] \\ z \in [(n-1)-\epsilon, (n-1)+\epsilon]}} \left(\frac{y}{y+z} - \frac{n-1}{n^2}c(y) \right) \\
&\leq \max_{\substack{y \in [1-\epsilon, 1+\epsilon] \\ z \in [(n-1)-\epsilon, (n-1)+\epsilon]}} \frac{y}{y+z} - \min_{y \in [1-\epsilon, 1+\epsilon]} \frac{n-1}{n^2}c(y) \\
&\leq \max_{y \in [1-\epsilon, 1+\epsilon]} \frac{y}{y + (n-1) - \epsilon} - \frac{n-1}{n^2}c(1-\epsilon) \\
&\leq \frac{1+\epsilon}{n} - \frac{n-1}{n^2}(c(1) - K\epsilon) \\
&\leq \frac{n - (n-1)c(1)}{n^2} + \frac{\epsilon(n + (n-1)K)}{n^2}
\end{aligned}$$

Putting these two bounds together, we get

$$\frac{u_-}{u_+} \geq \frac{n-(n-1)c(1)-\epsilon(n+(n-1)K)}{n-(n-1)c(1)+\epsilon(n+(n-1)K)} = \frac{1-\epsilon\frac{n+(n-1)K}{n-(n-1)c(1)}}{1+\epsilon\frac{n+(n-1)K}{n-(n-1)c(1)}}$$

$$\geq \frac{1-\epsilon(n+(n-1)K)}{1+\epsilon(n+(n-1)K)},$$

because $c(1) \leq 1$ as $c(0) = 0$ and $c'(y) \leq c'(1) = 1$ for $y \leq 1$,

$$\geq \frac{1-\epsilon n(1+K)}{1+\epsilon n(1+K)} = \frac{1-\widehat{\epsilon}/2}{1+\widehat{\epsilon}/2} \geq (1-\widehat{\epsilon}/2)^2,$$

because $\widehat{\epsilon} \leq 1$ and $1 \geq 1-y^2 \iff \frac{1}{1+y} \geq 1-y$ for $0 \leq y \leq 1$,

$$\geq 1-\widehat{\epsilon}, \text{ as required.}$$

As $c'(z) \geq c'(1) = 1$ for $z \in [1, 1+\epsilon]$, so $K \geq 1$ and $2(1+K)n\epsilon \leq 4Kn\epsilon$. $\square$

Proof of Theorem 2.18. The proof for upper bound is straightforward as we already have done most of the work in the proof of Theorem 2.9. First, to complete the warm-up phase (Definition 2.5), we need every agent to play at least once, and additionally, at most one more transition, as argued in the proof of Lemma 2.12. This can be completed in $n+1$ time steps by selecting agents in a round robin manner. Second, the number of non-redundant transitions (consecutive transitions by the same agent are redundant) required after the warm-up phase to reach the a phase of the BR dynamics that corresponds to the discounted-sum dynamics is bounded above by $O(\log\log(\frac{1}{\gamma}))$ as argued in the proof of Lemma 2.15. And this can be completed in $O(\log\log(\frac{1}{\gamma}))$ time by again selecting the agents in a round robin manner. Third, using the upper bound for the best-case selection model for the discount-sum dynamics (Lemma 2.8), we reach close to the equilibrium output profile measured by ℓ_1-distance in $O(n\log(\frac{n}{\epsilon}))$ time. Finally, using Lemma 2.17 completes our proof.

Let us now prove the lower bound. We will separately show bounds of $\Omega(n)$, $\Omega(\log\log(\frac{1}{\gamma}) - \log\log(n))$, and $\Omega(\frac{1}{\log(n)}\log(\frac{1}{\epsilon}))$, which together imply a lower bound of $\Omega(\max(n, \log\log(\frac{1}{\gamma}) - \log\log(n), \frac{1}{\log(n)}\log(\frac{1}{\epsilon})) = \Omega(n + \log\log(\frac{1}{\gamma}) + \frac{1}{\log(n)}\log(\frac{1}{\epsilon}))$. Note that these bounds are in the worst case over the class of all convex cost functions. In particular, except for the lower bound of $\Omega(n)$ that holds for all cost functions, for every ϵ and γ, there exists a convex cost function that reaches an approximate equilibrium after just one BR transition by each agent. In particular, if $c'(z) = z^r$ for all $z \geq 0$ and $r \to \infty$, then $BR(s_{-i}) \to 1$ for any i and s_{-i}. So, we will prove the lower bounds w.r.t. γ and ϵ using the linear cost function $c'(z) = 1$ for all $z \geq 0$. We leave tighter analysis of lower and upper bounds for specific classes of convex cost functions for future work.

Proof for $\Omega(n)$ Observe that, to reach an approximate equilibrium, every agent must make at least one best-response move. For example, it is easy to check that there can be no approximate equilibrium with any agent producing an output of n: if an agent is playing n, then everyone else must play ≤ 1 in an equilibrium (see Lemma 2.13 for a formal argument); if everyone else is playing ≤ 1, then this agent would want to deviate and not play n. So, if every agent starts with an output of n, then each agent must make at least one move to reach an approximate equilibrium. So, we get a lower bound of $\Omega(n)$ as required. Notice that the same argument also implies that for the randomized model, using Lemma 2.10, we get the lower bound of $\Omega(\frac{1}{L}\log(n\delta))$ w.p. $1-\delta$.

Proof for $\Omega(\log\log(\frac{1}{\gamma}) - \log\log(n))$ Notice that, at the equilibrium, the total output is n. So, for an output profile $\boldsymbol{x}_t$ to be at an ϵ-approximate equilibrium for small enough ϵ, we need $s_t \geq \frac{n}{2}$ (Lemma 2.23 formally proves a similar result). Starting from small $s_0 = \gamma$, we lower bound the time it takes to reach $s_t \geq \frac{n}{2}$. At time t, for any agent i, the first order condition, equation (2.4), for an agent with a linear cost function is

$$\frac{s_{t-1,-i}}{s_t^2} = \frac{n-1}{n^2} \iff s_t = \frac{n}{\sqrt{n-1}}\sqrt{s_{t-1}}$$
$$\implies s_t \leq n\sqrt{s_{t-1}} \leq n^{1+\frac{1}{2}} s_{t-2}^{\frac{1}{2^2}} \leq n^{1+\frac{1}{2}+\frac{1}{2^2}} s_{t-3}^{\frac{1}{2^4}} \ldots \leq n^2 s_0^{\frac{1}{2^t}} \leq n^2 \gamma^{\frac{1}{2^t}}.$$

So, $s_t \geq \frac{n}{2} \implies n^2\gamma^{\frac{1}{2^t}} \geq \frac{n}{2}$, or

$$s_t < \frac{n}{2} \impliedby n^2\gamma^{\frac{1}{2^t}} < \frac{n}{2} \iff \frac{1}{2^t}\lg(\gamma) < \lg\left(\frac{1}{2n}\right) \iff t < \lg\lg\left(\frac{1}{\gamma}\right) - \lg\lg(2n).$$

So, we need $t \geq \lg\lg\left(\frac{1}{\gamma}\right) - \lg\lg(2n)$ to ensure $s_t \geq \frac{n}{2}$ and that we are close to the equilibrium.

Proof for $\Omega(\frac{1}{\log(n)}\log(\frac{1}{\epsilon}))$ Before we prove the lower bound w.r.t. ϵ, let us prove the following lemma.

Lemma 2.23. *Given an output profile $\boldsymbol{x} = (x_i)_{i\in[n]}$ and the equilibrium profile $\boldsymbol{x}^* = (1,\ldots,1)$, if $||\boldsymbol{x} - \boldsymbol{x}^*||_1 \leq \epsilon$ and $|BR(s_{-i}) - 1| \leq \epsilon$, then $u_i(x_i, s_{-i}) < \left(1 - \frac{(x_i - BR(s_{-i}))^2}{8n^2}\right) u_i(BR(s_{-i}), s_{-i})$, for $0 \leq \epsilon \leq 1$ and $x_i \neq BR(s_{-i})$.*

Proof. We prove the lemma using strong concavity of the utility function near the equilibrium point. From equation (2.3), we have

$$-\frac{\partial^2 u_i(z, s_{-i})}{\partial z^2} = \frac{2s_{-i}}{(z + s_{-i})^3} + \frac{n-1}{n^2} c_i''(z) \geq \frac{2s_{-i}}{(z + s_{-i})^3}.$$

As $||\boldsymbol{x} - \boldsymbol{x}^*||_1 \leq \epsilon$ and $|BR(s_{-i}) - 1| \leq \epsilon$, we have $x_i \in [1-\epsilon, 1+\epsilon]$, $BR(s_{-i}) \in [1-\epsilon, 1+\epsilon]$, and $s_{-i} \in [n-1-\epsilon, n-1+\epsilon]$. So, for any $z \in [1-\epsilon, 1+\epsilon]$, we have

$$-\frac{\partial^2 u_i(z, s_{-i})}{\partial z^2} \geq \min_{\substack{z \in [1-\epsilon, 1+\epsilon] \\ s_i \in [n-1-\epsilon, n-1+\epsilon]}} \frac{2s_{-i}}{(z + s_{-i})^3} = \frac{2(n-1+\epsilon)}{(n+2\epsilon)^3} \geq \frac{1}{4n^2},$$

for $n \geq 2$ and $\epsilon \in [0, 1]$. At $z = BR(s_{-i})$, by the first order condition, we have $\left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=BR(s_{-i})} = 0$. So, using the strong convexity of $-u_i(z, s_{-i})$ w.r.t. z, we can write

$$\begin{aligned}
-u_i(x_i, s_{-i}) &\geq -u_i(BR(s_{-i}), s_{-i}) + (x_i - BR(s_{-i})) \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=BR(s_{-i})} \\
&\quad + \frac{1}{8n^2}(x_i - BR(s_{-i}))^2 \\
u_i(x_i, s_{-i}) &\leq u_i(BR(s_{-i}), s_{-i}) - \frac{1}{8n^2}(x_i - BR(s_{-i}))^2 \\
&< \left(1 - \frac{(x_i - BR(s_{-i}))^2}{8n^2}\right) u_i(BR(s_{-i}), s_{-i}),
\end{aligned}$$

as $u_i(BR(s_{-i}), s_{-i}) < 1$ and $x_i \neq BR(s_{-i})$, which completes the proof. □

Given Lemma 2.23, we can prove that an action profile $\boldsymbol{x}$ is not an ϵ-approximate equilibrium profile by showing that there is some agent i who has large deviation, i.e., $(x_i - BR(s_{-i}))^2$ is large.

Next, in Lemma 2.24, we show that if there are two agents that have output at least α away from the equilibrium output of 1 at time t, then there are at least two agents that are at least $\alpha/(2n)$ away from 1 at time $t+1$.

Lemma 2.24. *Assuming $n \geq 3$. At time t, if there exists agents i and $j \neq i$ such that $|x_{t,i} - 1| \geq \alpha$ and $|x_{t,j} - 1| \geq \alpha$ for some $\alpha \in (0, 1)$, then there exists agents i' and $j' \neq i'$ such that $|x_{t+1,i'} - 1| \geq \frac{\alpha}{2n}$ and $|x_{t+1,j'} - 1| \geq \frac{\alpha}{2n}$.*

Proof. Let i and $j \neq i$ denote the agents that are at least α away from the equilibrium, i.e., $|x_{t,i} - 1| \geq \alpha$ and $|x_{t,j} - 1| \geq \alpha$ at time t. If the agent making the BR move at time t is not in $\{i, j\}$, then we trivially have two agents—agents i and j—at time $t+1$ who are at least $\alpha \geq \frac{\alpha}{2n}$ away from the equilibrium. With a similar argument, if there are three or more agents that are at least $\frac{\alpha}{2n}$ away from the equilibrium at time t, we will have the required condition at time $t+1$. So, let us assume w.l.o.g. that agent i makes the transition at time t and all agents $k \notin \{i,j\}$ have $|x_{t,k}-1| < \frac{\alpha}{2n}$.

As $|x_{t,k}-1| < \frac{\alpha}{2n}$ for all $k \notin \{i,j\}$, we have $|\sum_{k\notin\{i,j\}} x_{t,k} - (n-2)| < \frac{(n-2)\alpha}{2n} < \frac{\alpha}{2}$. We also know that $|x_{t,j}-1| \geq \alpha$. Putting together, we have $|s_{t,-i}-(n-1)| > \frac{\alpha}{2}$, or $s_{t,-i} = (n-1) \pm \beta$ for $\beta > \frac{\alpha}{2}$. Let us do a case analysis based on the sign of $s_{t,-i}-(n-1)$.

- $s_{t,-i}-(n-1) < 0 \iff s_{t,-i} = (n-1)-\beta$. Using the first order condition for agent i (with a linear cost function $c'(z) = 1$ for $z \geq 0$), we can write $x_{t+1,i}$ as

$$
\begin{aligned}
x_{t+1,i} = n\sqrt{\frac{s_{t,-i}}{n-1}} - s_{t,-i} &= n\sqrt{\frac{(n-1)-\beta}{n-1}} - ((n-1)-\beta) \\
&= n\sqrt{1-\frac{\beta}{n-1}} - (n-1) + \beta \\
&\geq n\left(1-\frac{\beta}{n-1}\right) - (n-1) + \beta, \text{ because } 0 < 1-\frac{\beta}{n-1} < 1, \\
&\geq 1 + \frac{n-2}{n-1}\beta \geq 1 + \frac{\beta}{2} \geq 1 + \frac{\alpha}{4} \geq 1 + \frac{\alpha}{2n}.
\end{aligned}
$$

- $s_{t,-i}-(n-1) > 0 \iff s_{t,-i} = (n-1)+\beta$. Again using the first order condition for agent i, we can write $x_{t+1,i}$ as

$$
\begin{aligned}
x_{t+1,i} = n\sqrt{\frac{s_{t,-i}}{n-1}} - s_{t,-i} &= n\sqrt{\frac{(n-1)+\beta}{n-1}} - ((n-1)+\beta) \\
&= n\sqrt{1+\frac{\beta}{n-1}} - (n-1) - \beta \\
&\leq n\left(1+\frac{\beta}{n-1}\right) - (n-1) - \beta, \text{ because } 1+\frac{\beta}{n-1} > 1, \\
&\leq 1 - \frac{n-2}{n-1}\beta \leq 1 - \frac{\beta}{2} \leq 1 - \frac{\alpha}{4} \leq 1 - \frac{\alpha}{2n}.
\end{aligned}
$$

Putting both cases together, we have $|x_{t+1,i}-1| \geq \frac{\alpha}{2n}$; we also know that $|x_{t+1,j}-1| = |x_{t,j}-1| \geq \alpha \geq \frac{\alpha}{2n}$. So, we have the two agents at time $t+1$ that satisfy the required condition. $\square$

Lemma 2.24 above implies that if we start from $\boldsymbol{x}_0 = (\frac{1}{2}, \frac{1}{2}, 1, \ldots, 1)$ that has two agents who are producing output at least $\frac{1}{2}$ away from equilibrium output of 1, then after t time steps, $\boldsymbol{x}_t$ will have at least two agents i and $j \neq i$ with $|x_{t+1,k}-1| \geq \frac{1}{2(2n)^t}$ for $k \in \{i,j\}$. Let $\alpha = \frac{1}{2(2n)^t}$ for conciseness. Now, let $z_{t,k} = x_{t,k} - 1$ for all $t \geq 0$ and $k \in [n]$. As $|x_{t+1,k}-1| \geq \alpha$ for $k \in \{i,j\}$, we have $|z_{t,k}| \geq \alpha$ for $k \in \{i,j\}$.

Further, from Lemmas 2.13, 2.14, and 2.15, we know that $\boldsymbol{z}_t = (z_{t,k})_{k\in[n]}$ follows the discounted-sum dynamics presented in Section 2.4. Then, as $|z_{t,k}| \geq \alpha$ for $k \in \{i, j\}$, the potential function given in equation (2.7) satisfies $f(\boldsymbol{z}_t) \geq \alpha$. In particular, one of the two sides of the potential function $f(\boldsymbol{z}_t)$ has size at least α, so the largest element on the larger side, say $|z_{t,k}|$, must have value at least $\frac{\alpha}{n}$. Moreover, if we select this agent k at time t, then $|z_{t+1,k}| \leq \frac{|z_{t,k}|}{2}$ as argued in the proof of Lemma 2.7, which implies

$$\begin{aligned}
&|z_{t+1,k}| \leq \frac{|z_{t,k}|}{2} \\
&\quad\implies |z_{t+1,k} - z_{t,k}| \geq \frac{|z_{t,k}|}{2} \geq \frac{\alpha}{2n} = \frac{1}{2(2n)^{t+1}} \\
&\quad\implies |BR(s_{t,-k}) - x_{t,k}| \geq \frac{1}{2(2n)^{t+1}} \\
&\quad\implies u_k(x_{t,k}, s_{t,-k}) < \left(1 - \frac{1}{2^{t+5}n^{t+3}}\right) u_i(BR(s_{t,-k}), s_{t,-k}) \text{ by Lemma 2.23.}
\end{aligned}$$

Setting $t \leq \frac{\log(1/\epsilon)}{\log(2n)} - 5 \implies \frac{1}{2^{t+5}n^{t+3}} \geq \epsilon$ completes the proof. $\square$

Chapter 3

Tighter Bounds for Linear Costs

In this chapter, we improve our convergence-rate upper bound for Tullock contests with linear cost functions, i.e., lottery contests. We use techniques—Markov chains and stochastic descent—that are different from the analysis in Chapter 2, which may also be of independent interest.

We keep the same model and notation as Chapter 2. For linear cost, the BR of an agent can be written explicitly as a function of the output of other agents. As the agents are homogeneous, let us normalize the linear cost as $c_i(z) = z$ for all $i \in [n]$ and $z \geq 0$. Notice that this normalization is slightly different from the one in Section 2.2.4, where it would have been $c_i(z) = \frac{n-1}{n^2} z$. The $c_i(z) = z$ normalization makes the analysis in this chapter cleaner. Using the first-order condition, equation (2.2), the BR of any agent i can be written explicitly as $BR(s_{-i}) = \max(0, \sqrt{s_{-i}} - s_{-i})$.

In our analysis in this chapter, we make an additional but mild technical assumption: $p_{t,i}(\boldsymbol{x}_t) \leq U < 1/2$ for all t, i, and $\boldsymbol{x}_t$. This assumption says that no single agent has a significantly high probability of making the BR move at any given time step. It is required in one of the steps in our analysis (Lemma 3.2), but is not unreasonable, especially for a large number of agents. For our current analysis, relaxing this assumption seems non-trivial and probably requires new ideas.

For the randomized agent selection model, for $n \geq 3$ agents with homogeneous linear cost function, our upper and lower bounds for the convergence-time in Chapter 2 were $\tilde{O}(\frac{1}{L^2})$ and $\tilde{\Omega}(\frac{1}{L})$, respectively. Here, we improve the upper bound to $\tilde{O}(\frac{1}{L})$. In particular, we improve our rate of convergence upper bound to $O(\hat{\alpha} \log(n/(\epsilon\delta)) + \hat{\beta} \log\log(1/\gamma))$ with probability $1 - \delta$, where $\hat{\alpha}$ and $\hat{\beta}$ are functions of the randomized agent selection process (Theorem 3.1), e.g., $\alpha = n$ and $\beta = 1$ if agents are selected uniformly at random each time step. The improvement from $\tilde{O}(\frac{1}{L^2})$ to $\tilde{O}(\frac{1}{L})$ comes primarily from the following change: In Chapter 2, we analysed the discounted-sum dynamics, which gave us the $\tilde{O}(\frac{1}{L^2})$ bound. In this chapter, we analyze a descent algorithm with respect to a smooth and strongly convex potential function (discussed below), which gives us the better $\tilde{O}(\frac{1}{L})$ bound.

Intuition for the Proof of Convergence It is divided into three parts. The first part (warm-up phase) is similar to the analysis in Section 2.5. The third part of our analysis requires the total output to be sufficiently large. In the second part of the analysis, we prove this requirement is satisfied often. We discretize the range of values that the total output can take into countably infinitely many intervals. We consider the discrete stochastic process that indicates the interval in which the total output lies at a given time point. We analyze this discrete stochastic process and prove that it frequently visits the states that correspond to a high total output, as required. We do this using suitable concentration bounds for Markov chains [80]. This part of the analysis also provides a few interesting insights about the BR dynamics, e.g., the total output never decreases by more than a constant factor in a single time step after the warm-up phase.

In the third part of the analysis, we use the potential function f (different from the one used in Section 2.4) to prove the rate of convergence. We show that the potential f is strongly convex and smooth when the total output is sufficiently large, which we then use to show that the potential decreases by a multiplicative factor each time step (given the total output is large). Our analysis here is motivated by the $\log(1/\epsilon)$ convergence rate achieved for the minimization of strongly convex and smooth functions using coordinate descent [110, 15, 25]. On the other hand, the potential function is always non-increasing. Putting these together, we get our rate of convergence.

3.1 Analysis

Technical proofs omitted to improve readability are provided in Section 3.2.

Theorem 3.1. *Best-response dynamics in lottery contests with $n \geq 3$ homogeneous agents and randomized selection of agents reaches an ϵ-approximate equilibrium with probability $1-\delta$ in $O\left(\frac{1}{1-2U}\log\log\left(\frac{1}{\gamma}\right) + \frac{1}{L(1-2U)}\log\left(\frac{n}{\epsilon\delta}\right) + \frac{1}{(1-2U)^2}\log\left(\frac{1}{\delta}\right)\right)$ steps for every $\epsilon, \delta \in (0,1)$, where γ is a function of the initial output profile as given in Lemma 2.16.*

The $(1-2U)$ term in the denominator shows up because of the analysis of the Markov chain in Lemma 3.2. The assumption of $U < 1/2$ is also directly related.

The Markov chain we analyze is finite on one side, say left, but is infinite on the right side, and we want the left-most state to be visited frequently. The Markov chain is biased to move left, and this bias is ensured by the $U < 1/2$ condition and shows up as the $(1 - 2U)$ term in the denominator in the analysis.

Notice that the Theorem 3.1 implies convergence in $O\left(\log\log\left(\frac{1}{\gamma}\right) + n\log\left(\frac{n}{\epsilon\delta}\right)\right)$ steps when the agents are selected uniformly at random (i.e., $L = U = 1/n$).

Proof of Theorem 3.1. We divide the analysis into three parts. The first part corresponds to a warm-up phase and is exactly the same as Theorem 2.9. The definition of γ is also the same as before (Lemma 2.16).[1] The remaining two parts of the analysis assume the completion of the warm-up phase. In the second part, we bound the time it takes for the total output s_t to be at least $\frac{1}{4}$ for a sufficiently large number of time steps. We do it by discretizing the space of s_t and analyzing the Markov chain associated with it. Finally, in the third part of the analysis, we show that, after $s_t \geq \frac{1}{4}$ for a sufficiently large number of steps, then the output profile is close to the equilibrium output profile with high probability. Here we use a well-behaved potential function to be introduced later.

A Biased Random Walk with a Wall Let us consider a Markov chain with countably infinitely many states denoted by positive integers $\mathbb{Z}_{>0}$. Let $y_t \in \mathbb{Z}_{>0}$ be a random variable that denotes the state of this Markov chain at time $t \geq 0$. The state moves left (or stays at 1 if it is already at 1) w.p. $(1-p)$ and moves right w.p. p, where $p < 1/2$. Formally, the transition function of the chain is: $\mathbb{P}[y_{t+1} = 1 \mid y_t = 1] = 1-p$, $\mathbb{P}[y_{t+1} = 2 \mid y_t = 1] = p$, $\mathbb{P}[y_{t+1} = i-1 \mid y_t = i] = 1-p$, and $\mathbb{P}[y_{t+1} = i+1 \mid y_t = i] = p$ for $i \in \mathbb{Z}_{\geq 2}$. Figure 3.1 shows the Markov chain. We bound the time it takes for the

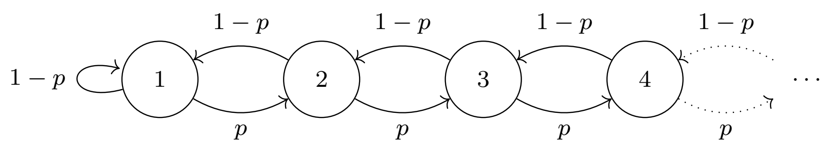

Figure 3.1: Markov chain

stochastic process y_t to visit state 1 a sufficiently large number of times with high probability, which we do using suitable concentration bounds in Lemma 3.2.

[1]The different normalization of the cost function in this section does not make a difference in the final upper bound because $1/\gamma$ is off by at most a $O(n)$ factor, which becomes an additive term after we take log.

Lemma 3.2. *Starting from initial state $y_0 = k$, the Markov chain visits state 1 at least m times w.p. at least $1 - \delta$ after $\frac{4}{1-2p} \max\left(m + k, \frac{1}{1-2p} \ln\left(\frac{1}{\delta}\right)\right)$ time steps.*

Analysis Part 2 (Total Output) We now use the Markov chain defined earlier to bound the time it takes after the warm-up phase for s_t to be above $\frac{1}{4}$ for a large number of time steps with high probability.

Lemma 3.3. *Within $T = T_{warm} + O\left(\frac{1}{1-2U}m + \frac{1}{1-2U} \log\log\left(\frac{1}{\gamma}\right) + \frac{1}{(1-2U)^2} \log\left(\frac{1}{\delta}\right)\right)$ time steps, w.p. $1 - \delta$, there must have been at least m times steps where $s_t \geq \frac{n^2}{4(n-1)}$, where γ is as defined in Lemma 2.16; formally, $\left|\left\{t \mid T_{warm} \leq t \leq T, s_t \geq \frac{1}{4}\right\}\right| \geq m$ w.p. $1 - \delta$.*

Proof. In this proof, we assume that the warm-up phase has completed. To make the analysis less cluttered, let us assume that $T_{warm} = 0$, i.e., x_t satisfies all conditions for completion of the warm-up phase for all $t \geq 0$ (rather than $t \geq T_{warm}$).

We map the evolution of s_t to a discrete stochastic process z_t that takes values in $\mathbb{Z}_{>0}$:

- We partition the unit interval $(0, 1)$, range of s_t, into the following intervals (written in decreasing order, starting closer to 1 and going to 0):

$$\left[\frac{1}{4}, 1\right), \left[\frac{1}{8}, \frac{1}{4}\right), \left[\frac{1}{32}, \frac{1}{8}\right), \ldots, \left[\left(\frac{1}{2}\right)^{2^{\ell-1}+1}, \left(\frac{1}{2}\right)^{2^{\ell-2}+1}\right), \ldots,$$

  where the ℓ-th interval for $\ell \geq 2$ is equal to $[(1/2)^{2^{\ell-1}+1}, (1/2)^{2^{\ell-2}+1})$. Notice that the values in the intervals are decreasing at a double exponential rate.

- The state of the stochastic process z_t is equal to ℓ if s_t lies in the ℓ-th interval. Formally,

$$z_t = \begin{cases} 1, & \text{if } s_t \in [1/4, 1), \\ \ell, & \text{if } s_t \in \left[(1/2)^{2^{\ell-1}+1}, (1/2)^{2^{\ell-2}+1}\right) \text{ where } \ell \geq 2. \end{cases}$$

We shall prove that the process $(z_t)_{t\geq 0}$ has a higher tendency to move left than y_t, where y_t follows the Markov chain defined earlier (Figure 3.1), given we set $p = U$ for y_t. Remember, y_t moves one step right ($y_{t+1} = y_t + 1$) w.p. $p = U$ and one step left ($y_{t+1} = \max(1, y_t - 1)$) w.p. $1 - p = 1 - U$. We shall show that z_t moves at least one step left (or stays at 1) w.p. at least $1 - p$ and moves at most one step right w.p. at most p. If we set $z_0 = y_0$, then this claim ensures that the probability that z_t visits state 1 at least m times is at least the probability that y_t visits state 1 at least m times.

Left transitions. We know that $s_{t+1} = BR(s_{t,-i}) + s_{t,-i} = (\sqrt{s_{t,-i}} - s_{t,-i}) + s_{t,-i} = \sqrt{s_{t,-i}}$, where i is selected from $[n]$ randomly. As $\sum_i x_{t,i} = s_t$, there can be at most one agent j such that $x_{t,j} > s_t/2$. This agent j is selected w.p. at most U. So, w.p. at least $1-U$ we select an agent i such that $x_{t,i} \le s_t/2 \implies s_{t,-i} \ge s_t/2 \implies s_{t+1} = \sqrt{s_{t,-i}} \ge \sqrt{s_t/2}$.

- If $z_t = 1$, then $s_t \ge 1/4$. So, $s_{t+1} \ge \sqrt{s_t/2} \ge \sqrt{1/8} \ge 1/4$ w.p. at least $1-U = 1-p$. So, $z_{t+1} = 1$ w.p. at least $1-p$.
- If $z_t = \ell$ for some $\ell \ge 2$, then $s_t \ge (1/2)^{2^{\ell-1}+1}$. So, $s_{t+1} \ge \sqrt{s_t/2} \ge \sqrt{(1/2)^{2^{\ell-1}+1}/2} = \sqrt{(1/2)^{2^{\ell-1}+2}} = (1/2)^{2^{\ell-2}+1}$ w.p. at least $1-p$. So, $z_{t+1} \le \ell - 1$ w.p. at least $1-p$.

Right transitions. We showed above that z_t makes a left transition w.p. at least $1-p$. So, z_t stays at its position or makes a right transition w.p. at most p, which is what we want. But we still need to show that z_t can move at most one step right, i.e., $z_{t+1} \le z_t + 1$ w.p. 1.

Let us focus on the transition at time t: $s_t \to s_{t+1}$, $z_t \to z_{t+1}$. Say agent $i = i_t$ makes this transition. If i made the best-response move at time $t-1$ as well, i.e., $i_{t-1} = i$, then $s_{t+1} = s_t$ because consecutive moves are redundant. Therefore, $z_{t+1} = z_t \le z_t + 1$, as required.

Let us assume that $i_{t-1} \ne i$. Let $j = i_{t-1} \ne i$ be the agent that made the move at time $t-1$. As $i_{t-1} \ne i$, therefore $x_{t-1,i} = x_{t,i}$. Similarly, as $i_t \ne j$, therefore $x_{t,j} = x_{t+1,j}$. On the other hand, for $k \ne i,j$, notice that $x_{t-1,k} = x_{t,k} = x_{t+1,k}$.

Let $\alpha = \sum_{k \ne i,j} x_{t,k}$. We can write $s_{t,-i}$ using α as

$$s_{t,-i} = x_{t,j} + \sum_{k \ne i,j} x_{t,k} = x_{t,j} + \alpha. \tag{3.1}$$

We can also write $x_{t,j}$ using α as

$$x_{t,j} = \sqrt{s_{t-1,-j}} - s_{t-1,-j} = \sqrt{x_{t-1,i} + \alpha} - (x_{t-1,i} + \alpha) = \sqrt{x_{t,i} + \alpha} - (x_{t,i} + \alpha), \tag{3.2}$$

and s_t using α as

$$s_t = \sqrt{s_{t-1,-j}} = \sqrt{x_{t-1,i} + \alpha} = \sqrt{x_{t,i} + \alpha}. \tag{3.3}$$

Combining (3.1) and (3.2), we get

$$s_t - x_{t,i} = x_{t,j} + \alpha = \sqrt{x_{t,i}+\alpha} - (x_{t,i}+\alpha) + \alpha = \sqrt{x_{t,i}+\alpha} - x_{t,i}. \tag{3.4}$$

We shall now lower bound the ratio between s_{t+1} and s_t. From (3.3) and (3.4), we have

$$\frac{s_{t+1}}{s_t} = \frac{\sqrt{s_t - x_{t,i}}}{s_t} = \frac{\sqrt{\sqrt{x_{t,i}+\alpha} - x_{t,i}}}{\sqrt{x_{t,i}+\alpha}} = \sqrt{\frac{1}{\sqrt{x_{t,i}+\alpha}} - \frac{x_{t,i}}{x_{t,i}+\alpha}}.$$

Note that $x_{t,i} + \alpha = s_t < 1$, so $\alpha < 1 - x_{t,i}$. Let us define the function $g(y) = \frac{1}{\sqrt{x_{t,i}+y}} - \frac{x_{t,i}}{x_{t,i}+y}$ for $y \in [0, 1-x_{t,i}]$. Notice that $\frac{s_{t+1}}{s_t} = \sqrt{g(\alpha)} \geq \min_{y\in[0,1-x_{t,i}]} \sqrt{g(y)}$. Let us differentiate $g(y)$ w.r.t. y, we get

$$g'(y) = \frac{-1}{2(x_{t,i}+y)^{3/2}} + \frac{x_{t,i}}{(x_{t,i}+y)^2} = \frac{2x_{t,i} - \sqrt{x_{t,i}+y}}{2(x_{t,i}+y)^2}.$$

As $x_{t,i} \leq 1/4$, we have

$$\sqrt{x_{t,i}} \leq \frac{1}{2} \implies 2x_{t,i} \leq \sqrt{x_{t,i}} \implies 2x_{t,i} - \sqrt{x_{t,i}} \leq 0 \implies 2x_{t,i} - \sqrt{x_{t,i}+y} \leq 0,$$

for every $y \geq 0$. Therefore, $g'(y) \leq 0$ for very $y \geq 0$. So,

$$\min_{y\in[0,1-x_{t,i}]} g(y) = g(1-x_{t,i}) = \frac{1}{\sqrt{x_{t,i}+(1-x_{t,i})}} - \frac{x_{t,i}}{x_{t,i}+(1-x_{t,i})} = 1 - x_{t,i} \geq \frac{3}{4},$$

as $x_{t,i} \leq 1/4$. As $\sqrt{\cdot}$ is a monotone increasing function, so

$$\frac{s_{t+1}}{s_t} = \sqrt{g(\alpha)} \geq \min_{y\in[0,1-x_{t,i}]} \sqrt{g(y)} = \sqrt{g(1-x_{t,i})} \geq \frac{\sqrt{3}}{2}.$$

We now use this bound on s_{t+1} to bound z_{t+1}. If $z_t = \ell$ for some $\ell \geq 1$, then $s_t \geq (1/2)^{2^{\ell-1}+1}$, which implies

$$s_{t+1} \geq \frac{\sqrt{3}}{2} s_t \geq \frac{\sqrt{3}}{2}\left(\frac{1}{2}\right)^{2^{\ell-1}+1} = \sqrt{3}\left(\frac{1}{2}\right)^{(2^{\ell-1}+1)+1} \geq \left(\frac{1}{2}\right)^{(2^{\ell-1}+1)+1} \geq \left(\frac{1}{2}\right)^{(2^{\ell})+1}$$
$$\implies z_{t+1} \leq z_t + 1.$$

To summarize, we have (i) $z_{t+1} = \max(1, z_t - 1)$ w.p. $\geq 1 - p = 1 - U$; (ii) $z_{t+1} \leq z_t + 1$ w.p. 1. So, z_t has a (weakly) higher tendency of moving left than y_t, as required. Finally, let us upper bound $z_0 = y_0$ as a function γ. We claim that $k \leq 1 + \lg\lg(1/\gamma)$ because

$$z_0 = k \leq 1 + \lg\lg(1/\gamma) \iff s_0 \geq \left(\frac{1}{2}\right)^{2^{k-1}+1} \impliedby s_0 \geq \left(\frac{1}{2}\right)^{\lg(\frac{1}{\gamma})+1} \geq \left(\frac{1}{2}\right)^{\lg(\frac{1}{\gamma})} = \gamma.$$

$\square$

Analysis Part 3 (Potential Function) In the third part of the analysis, we shall make use of the potential function f given in equation (3.5), which is based upon a similar potential function in [44]. Let $\boldsymbol{z} = (z_i)_{i\in[n]} \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i\in[n]} z_i < 1$. Let $\boldsymbol{z}^* = (z_i^*)_{i\in[n]}$ where $z_i^* = \frac{n-1}{n^2}$. Notice that $\boldsymbol{z}^*$ is the equilibrium action profile. Let $\sigma = \sum_i z_i$ and $\sigma^* = \sum_i z_i^*$.

$$f(\boldsymbol{z}) = \frac{1}{3}\left(\sum_i z_i\right)^3 - \sum_{i<j} z_i z_j + \frac{1}{6}\left(1 - \frac{1}{n}\right)^3, \tag{3.5}$$

Let us first prove some properties of the potential function f.

Lemma 3.4. *The potential function f satisfies:*

1. $f(\boldsymbol{z}^*) = 0$;
2. $0 \leq f(\boldsymbol{z}) \leq 1/2$;
3. $f(\boldsymbol{z}) \geq \frac{1}{40}\left(\frac{n-1}{n}\right)^3$ *if* $\sum_i z_i \leq \frac{3(n-1)}{4n}$;
4. $f(\boldsymbol{z}) \geq \frac{1}{2}||\boldsymbol{z} - \boldsymbol{z}^*||_2^2 = \frac{1}{2}\sum_i (z_i - z_i^*)^2$ *if* $\sum_i z_i \geq \frac{3(n-1)}{4n}$;
5. $f(\boldsymbol{z}) \leq \frac{1}{2}\sum_i (\sum_{j\neq i}(z_i - z_i^*))^2$.

In Lemma 3.4 (4), we showed that the ℓ_2-distance of the output profile is close to the equilibrium profile if the potential function is close to 0. In the Lemma 3.5, we show that we are at an approximate equilibrium if the output profile is close to the equilibrium profile.

Lemma 3.5. *If $||\boldsymbol{z} - \boldsymbol{z}^*||_2 \leq \epsilon$ for $0 \leq \epsilon < \frac{1}{4n\sqrt{n}}$, then $\boldsymbol{z}$ is an $(4n\sqrt{n}\epsilon)$-approximate equilibrium.*

Let us now go back to the BR dynamics. In the next lemma, we prove that the potential converges to 0 quickly with high probability.

Lemma 3.6. *$\mathbb{P}[f(\boldsymbol{x}_\tau) \leq \epsilon] \geq 1 - \delta$ for all $\tau \geq T$ after there have been at least $T = O\left(\frac{1}{L}\log\left(\frac{1}{\epsilon\delta}\right)\right)$ time steps t where the total output $s_t \geq \frac{1}{4}$ (after the warm-up phase).*

Proof. We assume that the warm-up phase has been completed. Let $f_i(\boldsymbol{x}_t)$ denote the value of the potential function at time $t+1$ given i played at time t. If i plays at time t, then $x_{t+1,i} = \sqrt{s_{t,-i}} - s_{t,-i}$ and $s_{t+1} = \sqrt{s_{t,-i}}$, and we have

$$\begin{aligned} f_i(\boldsymbol{x}_t) - \frac{1}{6}\left(\frac{n-1}{n}\right)^3 &= \frac{1}{3}s_{t+1}^3 - \sum_{j<k} x_{t+1,j}x_{t+1,k} \\ &= \frac{1}{3}s_{t+1}^3 - x_{t+1,i}\sum_{j\neq i} x_{t+1,j} - \sum_{j<k,j\neq i,k\neq i} x_{t+1,j}x_{t+1,k} \\ &= \frac{1}{3}(s_{t,-i})^{3/2} - (\sqrt{s_{t,-i}} - s_{t,-i})\sum_{j\neq i} x_{t,j} - \sum_{j<k,j\neq i,k\neq i} x_{t,j}x_{t,k} \\ &= \frac{1}{3}(s_{t,-i})^{3/2} - (\sqrt{s_{t,-i}} - s_t)(s_{t,-i}) - \sum_{j<k} x_{t,j}x_{t,k} \\ &= s_t(s_{t,-i}) - \frac{2}{3}(s_{t,-i})^{3/2} - \sum_{j<k} x_{t,j}x_{t,k}. \end{aligned} \tag{3.6}$$

We first show that $f_i(\boldsymbol{x}_t) \leq f(\boldsymbol{x}_t)$ w.p. 1 for every $i \in [n]$. Let $g(y) = \frac{1}{3}s_t^3 - y^2\left(s_t - \frac{2}{3}y\right)$ for $y \in [0,1]$. Notice that

$$f(\boldsymbol{x}_t) - f_i(\boldsymbol{x}_t) = \frac{1}{3}s_t^3 - \left(s_t(s_{t,-i}) - \frac{2}{3}(s_{t,-i})^{3/2}\right) = g(\sqrt{s_{t,-i}}).$$

Differentiating $g(y)$ w.r.t. y we get $g'(y) = 2y(y - s_t)$. As $g'(y) \leq 0$ for $y \in [0, s_t]$ and $g'(y) \geq 0$ for $y \in [s_t, 1]$, we have $\min_y g(y) = g(s_t) = \frac{1}{3}s_t^3 - s_t^2\left(s_t - \frac{2}{3}s_t\right) = 0$. As $s_{t,-i} \in [0,1]$, we have $g(\sqrt{s_{t,-i}}) \geq g(s_t) = 0 \implies f(\boldsymbol{x}_t) \geq f_i(\boldsymbol{x}_t)$, as required.

We next prove that the expected value of the potential $\mathbb{E}[f(\boldsymbol{x}_t)]$ decreases by a multiplicative factor in any given time step t to $t+1$ if $s_t \geq 1/4$. In particular, we show that $\mathbb{E}[f(\boldsymbol{x}_{t+1})|\boldsymbol{x}_t] \leq (1 - \kappa L)\, f(\boldsymbol{x}_t)$ for some constant $\kappa > 0$, or equivalently, $\mathbb{E}[f(\boldsymbol{x}_{t+1})|\boldsymbol{x}_t] - f(\boldsymbol{x}_t) \leq -\kappa L f(\boldsymbol{x}_t)$, given $s_t \geq 1/4$. As $f_i(\boldsymbol{x}_t)$ denotes the value of the potential at time $t+1$ given i played at time t, so $\mathbb{E}[f(\boldsymbol{x}_{t+1})|\boldsymbol{x}_t] = \sum_i p_{t,i}(\boldsymbol{x}_t) f_i(\boldsymbol{x}_t)$. Further, as $g(\sqrt{s_{t,-i}}) = f(\boldsymbol{x}_t) - f_i(\boldsymbol{x}_t)$, we can alternatively prove the slightly stronger condition

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{x}_{t+1})|\boldsymbol{x}_t] - f(\boldsymbol{x}_t) \leq -\kappa L f(\boldsymbol{x}_t) &\iff \sum_i p_{t,i}(\boldsymbol{x}_t)(f(\boldsymbol{x}_t) - f_i(\boldsymbol{x}_t)) \\ = \sum_i p_{t,i}(\boldsymbol{x}_t) g(\sqrt{s_{t,-i}}) \geq \kappa L f(\boldsymbol{x}_t) &\impliedby \sum_i g(\sqrt{s_{t,-i}}) \geq \kappa f(\boldsymbol{x}_t). \end{aligned} \tag{3.7}$$

Let us assume that $s_t \geq 1/4$. Going back to our function $g(y) = \frac{1}{3}s_t^3 - y^2\left(s_t - \frac{2}{3}y\right)$, we have $g''(y) = 4y - 2s_t$. We show that $g(y) \geq \kappa_g(y - s_t)^2$ for a positive constant $\kappa_g = \min\left(\frac{1}{20}, \left(\frac{1}{3} - \frac{27}{125}\right)\left(\frac{1}{4}\right)^3\right) > 0$ using the case analysis below:

- Case $y \geq 3s_t/5$. Then, $g''(y) = 4y - 2s_t \geq 2s_t/5 \geq 1/10$. As $g(s_t) = 0$ and $g'(s_t) = 0$, we get

$$g(y) \geq g(s_t) + g'(s_t)(y - s_t) + \frac{1}{2}\frac{2s_t}{5}(y - s_t)^2 \geq \frac{1}{20}(y - s_t)^2.$$

- Case $y \leq 3s_t/5$. As $g'(y) \leq 0$ for $y \in [0, s_t]$, we have

$$\begin{aligned} g(y) \geq g\left(\frac{3s_t}{5}\right) &= \frac{1}{3}s_t^3 - \left(\frac{3s_t}{5}\right)^2\left(s_t - \frac{2}{3}\frac{3s_t}{5}\right) \\ &= \left(\frac{1}{3} - \left(\frac{3}{5}\right)^3\right)s_t^3 \geq \left(\frac{1}{3} - \frac{27}{125}\right)\left(\frac{1}{4}\right)^3. \end{aligned}$$

  Now, as $y \in [0, 1]$, we have $(y - s_t)^2 \leq 1$, and we get our required result.

Plugging in $\sqrt{s_{t,-i}}$ into $g(y)$, we get

$$g(\sqrt{s_{t,-i}}) \geq \kappa_g(\sqrt{s_{t,-i}} - s_t)^2 \implies \sum_i g(\sqrt{s_{t,-i}}) \geq \kappa_g \sum_i (\sqrt{s_{t,-i}} - s_t)^2, \tag{3.8}$$

given $s_t \geq 1/4$. We now prove that $\sum_i(\sqrt{s_{t,-i}} - s_t)^2 \geq \kappa_h f(\boldsymbol{x}_t)$ for a positive constant $\kappa_h = \min\left(\left(\sqrt{\frac{6}{5}} - 1\right)^2 \frac{1}{8}, \frac{1}{12}\right) > 0$, assuming $s_t \geq 1/4$. We do a case analysis on the value of s_t:

- $\frac{1}{4} \leq s_t \leq \frac{5(n-1)}{6n}$. We claim that there exists an agent i such that $\sqrt{s_{t,-i}} \geq \sqrt{\frac{6}{5}}s_t$. Let this be false, i.e., for every $j \in [n]$, $\sqrt{s_{t,-j}} < \sqrt{\frac{6}{5}}s_t$. Then, we get a contradiction because

$$\begin{aligned} &\sqrt{s_{t,-j}} < \sqrt{\frac{6}{5}}s_t \iff s_{t,-j} < \frac{6}{5}s_t^2, \text{ for all } j \in [n] \\ &\quad \implies \sum_j (s_{t,-j}) < \sum_j \frac{6}{5}s_t^2 \implies (n-1)s_t < \frac{6}{5}ns_t^2 \implies s_t > \frac{5(n-1)}{6n}. \end{aligned}$$

  Now, as $\sqrt{s_{t,-i}} \geq \sqrt{\frac{6}{5}}s_t$ for an agent i, we have

$$\sum_j (\sqrt{s_{t,-j}} - s_t)^2 \geq (\sqrt{s_{t,-i}} - s_t)^2 \geq \left(\sqrt{\frac{6}{5}} - 1\right)^2 s_t^2 \geq \left(\sqrt{\frac{6}{5}} - 1\right)^2 \frac{1}{8}f(x_t), \tag{3.9}$$

  because $s_t \geq 1/4$ and $f(x_t) \leq 1/2$ (Lemma 3.4 (2)).

- $\frac{5(n-1)}{6n} \le s_t \le 1$. First, notice that

$$
\begin{aligned}
&(s_{t,-i} - s_t^2)^2 = (\sqrt{s_{t,-i}} - s_t)^2(\sqrt{s_{t,-i}} + s_t)^2 \\
&\qquad \implies \sum_i (\sqrt{s_{t,-i}} - s_t)^2 = \sum_i \frac{(s_{t,-i} - s_t^2)^2}{(\sqrt{s_{t,-i}} + s_t)^2} \ge \frac{1}{4}\sum_i (s_{t,-i} - s_t^2)^2, \quad (3.10)
\end{aligned}
$$

because $s_t \le 1$ and $\sqrt{s_{t,-i}} \le 1$. We do a change of variable. Let $y_k = s_{t,-k}$. Let $\rho = \sum_k y_k = \sum_k s_{t,-k} = (n-1)s_t \implies s_t = \rho/(n-1)$. Let us define the function $h(\boldsymbol{y}) = \frac{1}{4}\sum_k (y_k - (\frac{\rho}{n-1})^2)^2$. Taking derivative of h we get

$$
\begin{aligned}
\frac{\partial h}{\partial y_i} &= \frac{1}{2}\left(y_i - \left(\frac{\rho}{n-1}\right)^2\right) + \frac{1}{2}\sum_k \left(y_k - \left(\frac{\rho}{n-1}\right)^2\right)\left(\frac{-2}{n-1}\frac{\rho}{n-1}\right) \\
&= \frac{y_i}{2} - \frac{\rho^2}{2(n-1)^2} - \frac{\rho^2}{(n-1)^2} + \frac{n\rho^3}{(n-1)^4} = \frac{y_i}{2} + \frac{n\rho^3}{(n-1)^4} - \frac{3\rho^2}{2(n-1)^2}, \\
\frac{\partial^2 h}{\partial y_i^2} &= \frac{1}{2} + \frac{3n\rho^2}{(n-1)^4} - \frac{3\rho}{(n-1)^2}, \qquad \frac{\partial^2 h}{\partial y_i \partial y_j} = \frac{3n\rho^2}{(n-1)^4} - \frac{3\rho}{(n-1)^2}.
\end{aligned}
$$

So, the hessian $\nabla^2 h = \left(\frac{3n\rho^2}{(n-1)^4} - \frac{3\rho}{(n-1)^2}\right) + \frac{1}{2}I$, where I is the identity matrix. Let's find the eigenvalues of $\nabla^2 h$. Using the matrix determinant lemma, we have

$$
\begin{aligned}
\det(\nabla^2 h - \lambda I) &= \det\left(\left(\frac{3n\rho^2}{(n-1)^4} - \frac{3\rho}{(n-1)^2}\right) + \left(\frac{1}{2} - \lambda\right) I\right) \\
&= \left(1 + \frac{n\left(\frac{3n\rho^2}{(n-1)^4} - \frac{3\rho}{(n-1)^2}\right)}{\left(\frac{1}{2} - \lambda\right)}\right)\left(\frac{1}{2} - \lambda\right)^n.
\end{aligned}
$$

Setting $\det(\nabla^2 h - \lambda I) = 0$ we get $\lambda = \frac{1}{2}$ or $\lambda = \frac{1}{2} + \frac{3n^2\rho^2}{(n-1)^4} - \frac{3n\rho}{(n-1)^2} = \frac{1}{2} + 3\frac{np}{(n-1)^2}\left(\frac{np}{(n-1)^2} - 1\right)$. As $s_t \ge \frac{5(n-1)}{6n} \implies \rho = (n-1)s_t \ge \frac{5(n-1)^2}{6n} \implies \frac{np}{(n-1)^2} \ge \frac{5}{6}$, we get $\lambda \ge \frac{1}{2} + 3\frac{5}{6}\left(\frac{5}{6} - 1\right) = \frac{1}{12}$. So, h is (1/12)-strongly convex. Now, let $y_k^* = \frac{(n-1)^2}{n^2}$ for every $k \in [n]$, which corresponds to the equilibrium. Notice that we have $\rho^* = \frac{(n-1)^2}{n} \ge \frac{5(n-1)^2}{6n}$. Further, $h(\boldsymbol{y}^*) = 0$ because $\left(y_i^* - \left(\frac{\rho}{n-1}\right)^2\right)^2 = 0$ for very $i \in [n]$ and $\nabla h(\boldsymbol{y}^*) = 0$ because

$$
\begin{aligned}
(\nabla h(\boldsymbol{y}^*))_i &= \frac{y_i^*}{2} + \frac{n(\rho^*)^3}{(n-1)^4} - \frac{3(\rho^*)^2}{2(n-1)^2} \\
&= \frac{1}{2}\frac{(n-1)^2}{n^2} + \frac{n}{(n-1)^4}\frac{(n-1)^6}{n^3} - \frac{3}{2(n-1)^2}\frac{(n-1)^4}{n^2} = 0,
\end{aligned}
$$

for every $i \in [n]$. So, we can write

$$
h(\boldsymbol{y}) \ge h(\boldsymbol{y}^*) + (\nabla h(\boldsymbol{y}^*))^\intercal(\boldsymbol{y} - \boldsymbol{y}^*) + \frac{1}{24}||\boldsymbol{y} - \boldsymbol{y}^*||_2^2 = \frac{1}{24}||\boldsymbol{y} - \boldsymbol{y}^*||_2^2.
$$

From Lemma 3.4 (5), we know that $||\boldsymbol{y}-\boldsymbol{y}^*||_2^2 \geq 2f(\boldsymbol{x}_t)$. Putting everything together,

$$\sum_j (\sqrt{s_{t,-j}} - s_t)^2 \geq \frac{1}{4}\sum_j (s_{t,-j} - s_t^2)^2 = h(\boldsymbol{y}) \geq \frac{1}{24}||\boldsymbol{y}-\boldsymbol{y}^*||_2^2 \geq \frac{1}{12} f(\boldsymbol{x}_t). \quad (3.11)$$

Combining (3.7), (3.8), (3.9), and (3.11), we have $\mathbb{E}[f(\boldsymbol{x}_{t+1})|\boldsymbol{x}_t] \leq (1-\kappa_g\kappa_h L)\, f(\boldsymbol{x}_t)$ assuming $s_t \geq 1/4$.

Let $\kappa = \kappa_g\kappa_h$. Taking expectation w.r.t. $\boldsymbol{x}_t$, we get $\mathbb{E}[f(\boldsymbol{x}_{t+1})] \leq (1-\kappa L)\,\mathbb{E}[f(\boldsymbol{x}_t)]$ for all $\boldsymbol{x}_t$ conditioned on $s_t \geq 1/4$. Further, as $f_i(\boldsymbol{x}_t) \leq f(\boldsymbol{x}_t)$ for every $i \in [n]$ w.p. 1, we also have $\mathbb{E}[f(\boldsymbol{x}_{t+1})] \leq \mathbb{E}[f(\boldsymbol{x}_t)]$ w.p. 1. So, the sequence $\mathbb{E}[f(\boldsymbol{x}_t)]$ is a non-increasing sequence and decreases by a constant factor for the time steps when $s_t \geq 1/4$.

Assuming that there have been at least $m \geq \frac{1}{\kappa L}\ln\left(\frac{1}{2\epsilon\delta}\right)$ steps after the warm-up phase until time t where the total output was at least $1/4$, we get

$$\mathbb{E}[f(\boldsymbol{x}_t)] \leq (1-\kappa L)^m f(\boldsymbol{x}_{T_{warm}}) \leq \frac{(1-\kappa L)^m}{2} \leq \frac{e^{-\kappa L m}}{2} \leq \epsilon\delta.$$

As $\mathbb{E}[f(\boldsymbol{x}_t)] \leq \epsilon\delta$ and $f(\boldsymbol{x}_t) \geq 0$, using Markov inequality we have $\mathbb{P}[f(\boldsymbol{x}_t) \geq \epsilon] \leq \frac{\mathbb{E}[f(\boldsymbol{x}_t)]}{\epsilon} \leq \delta$, which implies $\mathbb{P}[f(\boldsymbol{x}_t) \leq \epsilon]$ w.p. $1-\delta$, as required. □

For $\epsilon \leq \frac{1}{4n\sqrt{n}}$ and $s_t \geq \frac{3(n-1)}{n}$, from Lemma 3.5, we can ensure that x_t is an ϵ-approximate equilibrium by ensuring $||\boldsymbol{x}_t - \boldsymbol{z}^*|| \leq \frac{1}{4n\sqrt{n}}\epsilon$, where z^* is the equilibrium profile; which we can further guarantee, using Lemma 3.4 (4), by ensuring that $f(\boldsymbol{x}_t) \leq \frac{1}{8n\sqrt{n}}\epsilon$. On the other hand, for $\epsilon \geq \frac{1}{4n\sqrt{n}}$ or $s_t \leq \frac{3(n-1)}{n}$, we can simply plug in $\frac{1}{4n\sqrt{n}}$ instead of ϵ and get an $(\frac{1}{4n\sqrt{n}})$-approximate equilibrium, which is also an ϵ-approximate equilibrium, by ensuring that $f(\boldsymbol{x}_t) \leq \frac{1}{16n^3}$.

Let us now combine all our lemmas and finish the proof. We set the probabilities of failure in each of the Lemmas 2.12, 3.3, and 3.6 to be $\delta/3$, ensuring total probability of failure is at most δ by union bound. So, we get an ϵ-approximate equilibrium w.p. $1-\delta$ in steps:

$$O\left(\frac{1}{1-2U}\log\log\left(\frac{1}{\gamma}\right) + \frac{1}{L(1-2U)}\log\left(\frac{n}{\epsilon\delta}\right) + \frac{1}{(1-2U)^2}\log\left(\frac{1}{\delta}\right)\right).$$

□

3.2 Omitted Proofs

Proof of Lemma 3.2. Consider the biased random walk $(z_t)_{t\geq 0}$, where $z_{t+1} = z_t + 1$ w.p. $p < 1/2$ and $z_{t+1} = z_t - 1$ w.p. $1 - p > 1/2$. Notice that the behavior of z_t is almost similar to the behavior of y_t defined earlier (Figure 3.1), except that y_t cannot move left from state 1 but z_t can. Let us assume that $z_0 = y_0 = k$ and z_t is coupled with y_t, i.e., z_t moves left if y_t moves left (or stays at state 1) and z_t moves right if y_t moves right. Notice that $y_t \geq z_t$, for all t, by construction.

It is easy to check that, if $z_\tau \leq 1$ at time t, then $y_t = 1$ for some $t \leq \tau$. Because, if $y_\tau = z_\tau$, then trivially $y_t = 1$ for $t = \tau$. Else $y_\tau > z_\tau$, which can only happen if $z_t = y_t = 1$ at some time t and $z_{t+1} = 0$ but $y_{t+1} = 1$. Similarly, we can show that, if $z_\tau \leq 2 - m$, then y_t has visited state 1 at least m times by time τ, i.e., $|\{t \leq \tau \mid y_t = 1\}| \geq m$. The argument is similar as before: The first time z_t is at 1, y_t is also at 1 (the first visit). Further, the difference between y_t and z_t increases only when y_t is at state 1 and tries unsuccessfully to move left but z_t moves left. As $y_\tau \geq 1$ and $z_\tau \leq 2 - m$, so $y_\tau - z_\tau \geq m - 1$, so y_t makes at least $m - 1$ additional visits to state 1 by time τ.

Next, we bound the time it takes for z_t fall below $2 - m$. Let Z_i be a random variable that takes value 1 w.p. p and -1 w.p. $1 - p$. z_t can be written as $z_t = \sum_{i=1}^{t} Z_i + k$. Notice that $\mathbb{E}[\sum_{i=1}^{t} Z_i] = t(2p - 1)$. Applying Hoeffding's inequality, we have

$$\begin{aligned}
\mathbb{P}[z_t > 2 - m] &= \mathbb{P}\left[\sum_{i=1}^{t} Z_i + k \geq 3 - m\right] \\
&= \mathbb{P}\left[\sum_{i=1}^{t} Z_i - \mathbb{E}\left[\sum_{i=1}^{t} Z_i\right] \geq 3 - m - k + t(1 - 2p)\right] \\
&\leq \exp\left(\frac{-2(t(1-2p) - m - k + 3)^2}{4t}\right) \\
&\leq \exp\left(\frac{-(t(1-2p) - m - k + 3)^2}{2t}\right).
\end{aligned}$$

Setting $t \geq \frac{4}{1-2p} \max\left(m + k, \frac{1}{1-2p} \ln\left(\frac{1}{\delta}\right)\right)$, we get

$$\begin{aligned}
\exp\left(\frac{-(t(1-2p) - m - k + 3)^2}{2t}\right) &\leq \exp\left(\frac{-3^2 t^2 (1-2p)^2}{4^2 2t}\right) \\
&\leq \exp\left(\frac{-9}{8}\frac{t(1-2p)^2}{4}\right) \leq \delta.
\end{aligned}$$

□

Proof of Lemma 3.4. Proof of Lemma 3.4 (1). As $z_i^* = \frac{n-1}{n^2}$, we have $\sigma^* = \frac{n-1}{n}$ and $\sum_{i<j} z_i^* z_j^* = \binom{n}{2}\left(\frac{n-1}{n^2}\right)^2 = \frac{(n-1)^3}{2n^3}$. Plugging this into f we get

$$\begin{aligned} f(\boldsymbol{z}^*) &= \frac{1}{3}(\sigma^*)^3 - \sum_{i<j} z_i^* z_j^* + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 \\ &= \frac{1}{3}\left(\frac{n-1}{n}\right)^3 - \binom{n}{2}\left(\frac{n-1}{n^2}\right)^2 + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 = 0. \end{aligned}$$

Proof of Lemma 3.4 (2). Notice that we can write $\sum_{i<j} z_i z_j$ as $\left(\sum_i z_i\right)^2 - \sum_i z_i^2$. Using this

$$f(\boldsymbol{z}) = \frac{1}{3}\sigma^3 - \sum_{i<j} z_i z_j + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 = \frac{1}{3}\sigma^3 - \frac{1}{2}\sigma^2 + \frac{1}{2}\sum_i z_i^2 + \frac{1}{6}\left(\frac{n-1}{n}\right)^3.$$

Now, as $\sum_i z_i^2$ is a convex function of $\boldsymbol{z}$, we have

$$\sum_i z_i^2 \geq \sum_i \left(\frac{\sigma}{n}\right)^2 = \frac{\sigma^2}{n} \implies f(\boldsymbol{z}) \geq \frac{1}{3}\sigma^3 - \left(\frac{n-1}{2n}\right)\sigma^2 + \frac{1}{6}\left(\frac{n-1}{n}\right)^3.$$

Let $g(y) = \frac{1}{3}y^3 - \frac{n-1}{2n}y^2$ for $y \in [0,1]$. Notice that $f(\boldsymbol{z}) \geq g(\sigma) - \frac{1}{6}\left(\frac{n-1}{n}\right)^3$. Let us now lower bound $g(y)$.

$$g'(y) = y^2 - \frac{n-1}{n}y = 0 \implies y \in \left\{0, \frac{n-1}{n}\right\}.$$

Notice that $g'(y) \leq 0$ if $y \in [0, (n-1)/n]$ and $g'(y) \geq 0$ for $y \geq (n-1)/n$. So,

$$\begin{aligned} \min_{y\in[0,1]} g(y) = g\left(\frac{n-1}{n}\right) &= \frac{1}{3}\left(\frac{n-1}{n}\right)^3 - \frac{n-1}{2n}\left(\frac{n-1}{n}\right)^2 = -\frac{1}{6}\left(\frac{n-1}{n}\right)^3 \\ \implies f(\boldsymbol{z}) &\geq g(\sigma) + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 \geq g\left(\frac{n-1}{n}\right) + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 = 0. \end{aligned}$$

On the other hand, $f(\boldsymbol{z}) \leq \frac{1}{2}$ because $f(\boldsymbol{z}) = \frac{1}{3}\sigma^3 - \sum_{i<j} z_i z_j + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 \leq \frac{1}{3}\sigma^3 + \frac{1}{6} \leq \frac{1}{2}$.

Proof of Lemma 3.4 (3). We are given that $\sigma \leq 3(n-1)/(4n)$. We showed earlier that $f(\boldsymbol{z}) \geq g(\sigma) + ((n-1)/n)^3/6$ and that $g(y)$ decreases for $y \in [0, (n-1)/n]$, so

$$\begin{aligned} \min_{y\in[0,\frac{3(n-1)}{4n}]} g(y) \geq g\left(\frac{3(n-1)}{4n}\right) &= \left(\frac{1}{3}\frac{3^3}{4^3} - \frac{1}{2}\frac{3^2}{4^2}\right)\left(\frac{n-1}{n}\right)^3 = -\frac{3^2}{4^3}\left(\frac{n-1}{n}\right)^3 \\ \implies f(\boldsymbol{z}) \geq g(\sigma) + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 &\geq g\left(\frac{3(n-1)}{4n}\right) + \frac{1}{6}\left(\frac{n-1}{n}\right)^3 \\ &= \left(\frac{1}{6} - \frac{3^2}{4^3}\right)\left(\frac{n-1}{n}\right)^3 \geq \frac{1}{40}\left(\frac{n-1}{n}\right)^3. \end{aligned}$$

Proof of Lemma 3.4 (4). We are given that $\sigma \geq 3(n-1)/(4n)$. Let's show that f is 1-strongly convex if $\sigma \geq 3(n-1)/(4n)$. Taking derivative of f we get

$$\frac{\partial f}{\partial z_i} = \sigma^2 - \sum_{j \neq i} z_j, \qquad \frac{\partial^2 f}{\partial z_i^2} = 2\sigma, \qquad \frac{\partial^2 f}{\partial z_i \partial z_j} = 2\sigma - 1, \qquad \text{for } i, j \in [n],\, i \neq j.$$

So, the hessian $\nabla^2 f$ has $(\nabla^2 f)_{i,j} = 2\sigma$ if $i = j$ and $(\nabla^2 f)_{i,j} = 2\sigma - 1$ if $i \neq j$. The hessian can be written as $\nabla^2 f = (2\sigma - 1) + I$, where I is the $n \times n$ identity matrix. Let us now find the eigenvalues of $\nabla^2 f$; λ is an eigenvalue if $\det(\nabla^2 f - \lambda I) = \det((2\sigma-1) + (1-\lambda)I) = 0$. Using matrix determinant lemma,[2] we have

$$\det((2\sigma-1)+(1-\lambda)I) = \left(1 + \frac{(2\sigma-1)n}{1-\lambda}\right)(1-\lambda)^n = (1-\lambda)^{n-1}(2n\sigma-(n-1)-\lambda) = 0,$$

which gives us $\lambda = 1$ or $\lambda = 2n\sigma-(n-1) = 2n(\sigma - \frac{n-1}{2n}) \geq 2n(\frac{3(n-1)}{4n} - \frac{n-1}{2n}) = \frac{n-1}{2} \geq 1$. As all eigenvalues of the hessian $\nabla^2 f$ are at least 1, so f is 1-strongly convex. Now, as f is 1-strongly convex, we have

$$f(\boldsymbol{z}) \geq f(\boldsymbol{z}^*) + (\nabla f(\boldsymbol{z}^*))^\intercal(\boldsymbol{z} - \boldsymbol{z}^*) + \frac{1}{2}||\boldsymbol{z} - \boldsymbol{z}^*||_2^2.$$

We have shown earlier that $f(\boldsymbol{z}^*) = 0$. Further, as

$$(\nabla f(\boldsymbol{z}^*))_i = \left.\frac{\partial f}{\partial z_i}\right|_{\boldsymbol{z}^*} = (\sigma^*)^2 - \sum_{j \neq i} z_j^* = \left(\frac{n-1}{n}\right)^2 - \sum_{j \neq i}\left(\frac{n-1}{n^2}\right) = 0,$$

for every i, we have $\nabla f(\boldsymbol{z}^*) = 0$, which gives our required result $f(\boldsymbol{z}) \geq \frac{1}{2}||\boldsymbol{z} - \boldsymbol{z}^*||_2^2$.

Proof of Lemma 3.4 (5). The proof is similar to that of Lemma 3.4 (4), but instead of strong convexity, we focus on Lipschitz smoothness.

Let $y_i = \sum_{j \neq i} z_i = \sigma - z_i$. Let $\rho = \sum_i y_i = \sum_i (\sigma - z_i) = (n-1)\sigma \leq (n-1)$. Further,

$$\begin{aligned}
\sum_{i<j} z_i z_j = \sum_{i<j}(\sigma - y_i)(\sigma - y_j) &= \binom{n}{2}\sigma^2 - \sigma\sum_{i<j}(y_i + y_j) + \sum_{i<j} y_i y_j \\
&= \frac{n(n-1)}{2}\sigma^2 - \sigma(n-1)\rho + \sum_{i<j} y_i y_j \\
&= \left(\frac{n}{2(n-1)} - 1\right)\rho^2 + \sum_{i<j} y_i y_j = -\frac{n-2}{2(n-1)}\rho^2 + \sum_{i<j} y_i y_j.
\end{aligned}$$

[2] https://en.wikipedia.org/wiki/Matrix_determinant_lemma

Let us define the function $h(\boldsymbol{y})$ such that $h(\boldsymbol{y}) = f(\boldsymbol{z})$; plugging in $\sigma = \frac{1}{(n-1)}\rho$ and $\sum_{i<j} z_i z_j = -\frac{n-2}{2(n-1)}\rho^2 + \sum_{i<j} y_i y_j$, we have

$$h(\boldsymbol{y}) = f(\boldsymbol{z}) = \frac{1}{3(n-1)^3}\rho^3 + \frac{n-2}{2(n-1)}\rho^2 - \sum_{i<j} y_i y_j + \frac{1}{6}\left(\frac{n-1}{n}\right)^3.$$

Taking derivative of h w.r.t. $\boldsymbol{y}$ we get

$$\frac{\partial h}{\partial y_i} = \frac{1}{(n-1)^3}\rho^2 + \frac{n-2}{n-1}\rho - (\rho - y_i) = \frac{1}{(n-1)^3}\rho^2 - \frac{1}{n-1}\rho + y_i,$$
$$\frac{\partial^2 h}{\partial y_i^2} = \frac{2}{(n-1)^3}\rho - \frac{1}{n-1}\rho + 1, \quad \frac{\partial^2 h}{\partial y_i \partial y_j} = \frac{2}{(n-1)^3}\rho - \frac{1}{n-1}, \quad \text{for } i, j \in [n],\ i \neq j.$$

So, the hessian can be written as $\nabla^2 h = \left(\frac{2}{(n-1)^3}\rho - \frac{1}{n-1}\right) + I$, where I is the $n \times n$ identity matrix. Let us now find the eigenvalues of $\nabla^2 h$. Using the matrix determinant lemma, we have

$$\det\left(\left(\frac{2}{(n-1)^3}\rho - \frac{1}{n-1}\right) + (1-\lambda)I\right) = \left(1 + \frac{n\left(\frac{2}{(n-1)^3}\rho - \frac{1}{n-1}\right)}{1-\lambda}\right)(1-\lambda)^n$$
$$= (1-\lambda)^{n-1}\left(\frac{2n}{(n-1)^3}\rho - \frac{n}{n-1} + 1 - \lambda\right),$$

which gives us $\lambda = 1$ or $\lambda = \frac{2n}{(n-1)^3}\rho - \frac{1}{n-1} \leq \frac{2n(n-1)}{(n-1)^3} - \frac{1}{n-1} = \frac{2n-(n-1)}{(n-1)^2} = \frac{n+1}{(n-1)^2} = \frac{1}{(n-1)} + \frac{2}{(n-1)^2} \leq 1$. As all eigenvalues of the hessian $\nabla^2 h$ are at most 1, so $h(\boldsymbol{y})$ is 1-Lipschitz smooth w.r.t. $\boldsymbol{y}$. Let $y_i^* = \sum_{j \neq i} z_i^* = \left(\frac{n-1}{n}\right)^2$ and $\rho^* = \sum_i y_i^* = \frac{(n-1)^2}{n}$. Notice that $h(\boldsymbol{y}^*) = f(\boldsymbol{z}^*) = 0$ and $\nabla h(\boldsymbol{y}^*) = 0$ because

$$(h(\boldsymbol{y}^*))_i = \left.\frac{\partial h}{\partial y_i}\right|_{\boldsymbol{y}^*} = \frac{1}{(n-1)^3}\frac{(n-1)^4}{n^2} - \frac{1}{n-1}\frac{(n-1)^2}{n} + \frac{(n-1)^2}{n^2}$$
$$= \frac{n-1}{n^2}(1 - n + (n-1)) = 0,$$

for every $i \in [n]$. This gives us

$$f(\boldsymbol{z}) = h(\boldsymbol{y}) \leq h(\boldsymbol{y}^*) + (\nabla h(\boldsymbol{y}^*))^\intercal(\boldsymbol{y} - \boldsymbol{y}^*) + \frac{1}{2}||\boldsymbol{y} - \boldsymbol{y}^*||_2^2$$
$$= \frac{1}{2}\sum_i (y_i - y_i^*)^2 = \frac{1}{2}\sum_i \left(\sum_{j \neq i}(z_i - z_i^*)\right)^2.$$

$\square$

Proof of Lemma 3.5. We will use Lemma 2.17, which showed a similar bound for ℓ_1-distance. As $||\boldsymbol{z} - \boldsymbol{z}^*||_2 \leq \epsilon$, using Cauchy–Schwarz inequality, $||\boldsymbol{z} - \boldsymbol{z}^*||_1 = \sum_i |z_i - z_i^*| = \sum_i |z_i - z_i^*| \cdot 1 \leq \sqrt{\sum_i |z_i - z_i^*|^2 \sum_i 1^2} = \sqrt{n}||\boldsymbol{z} - \boldsymbol{z}^*||_2 = \sqrt{n}\epsilon$. We need to adjust the cost function to satisfy the normalization in Lemma 2.17 (and Section 2.2 and 2.5). It is easy to check that the re-normalized cost function is $c_i(x_i) = \frac{n-1}{n^2}x_i$ for all i, so it is $\frac{n-1}{n^2}$-Lipschitz continuous. Similarly, if $||\boldsymbol{z} - \boldsymbol{z}^*||_1 \leq \sqrt{n}\epsilon$, then the re-normalized ℓ_1-distance is $\leq \frac{n^2}{n-1}\sqrt{n}\epsilon$. So, Lemma 2.17 tells us that we are at an $4(\frac{n-1}{n^2})n(\frac{n^2}{n-1}\sqrt{n}\epsilon) = 4n\sqrt{n}\epsilon$ equilibrium. □

Chapter 4

Continuous-Time Best-Response and Related Dynamics

Our results in Chapters 2 and 3 show convergence of the standard discrete-time BR dynamics for homogeneous agents but non-convergence for non-homogeneous agents. In contrast, in this chapter, we show that continuous-time BR and related learning dynamics converge even for non-homogeneous agents.

4.1 Introduction

In real-life situations, agents may not be able to instantly change their current actions to their BR actions. They are more likely to slowly move from their current actions to actions with better utility. As a concrete example, consider the game among miners in proof-of-work cryptocurrencies such as Bitcoin [31, 72]. A Bitcoin miner's expected reward for the next block is proportional to her costly computational effort, which is closely approximated by a Tullock contest. In this game, a miner may gradually buy more equipment and thereby increase their computational power if they see an opportunity to improve their utility by doing so. The change may be more rapid if the discrepancy is substantial, i.e., if the distance between the current action and the optimal BR action is larger, but we may still expect the change to be smooth enough that other agents can react before a given agent moves from their current action to their BR action. This observation also holds for other applications of Tullock contests: e.g., in a competition among firms for research and development of drugs and vaccines [35], a firm may generally be able to only slowly change their output by hiring/firing researchers or increasing/decreasing their investment.

The classic model for studying such smoothly changing actions is the continuous-time BR dynamics [52], which is the focus of this chapter. We show that continuous-time BR dynamics converges to the unique equilibrium when agents have arbitrary, possibly non-homogeneous, weakly convex cost functions.[1] Our rate-of-convergence bound is tight: the dynamics converges to an ϵ-approximate equilibrium in $\Theta(\log(1/\epsilon))$ time. We prove this result using a Lyapunov potential function that measures the total regret, as perceived by the agents, for playing their current action instead of their BR action. We then extend this analysis to show convergence in certain classes of discrete-time dynamics, e.g., when the agents take small steps towards the BR (where the step size is not necessarily limiting to 0 as assumed for continuous time) or when the agents best-respond to the empirical average action of other agents. These results also lend to a simple algorithm to compute an ϵ-approximate Nash equilibrium in time polynomial in $1/\epsilon$ and parameters of the model.

For related work, see Section 2.1.2 in Chapter 2. Additionally, the potential function in our Lyapunov arguments has previously been used in the literature in a different context of proving convergence of stochastic fictitious play in two-player zero-sum games [58, 95]; however, these results do not extend to zero-sum games with three or more players.[2] Our analysis works because of the special structure exhibited by Tullock contests.

4.2 Preliminaries

We keep the same notation as Chapters 2 and 3, defined in Section 2.2, except for the following minor changes: Instead of denoting the output of time t by $\boldsymbol{x}_t$, as done in Chapters 2 and 3, in this chapter we write it as $\boldsymbol{x}(t)$; this notation matches the literature in continuous dynamical systems and makes our analysis easier to read (e.g., the time derivative of $\boldsymbol{x}(t)$ reads better). Also, we will use an additive ϵ-approximate equilibrium[3] instead of the multiplicative version given in Definition 2.2. This is again for a cleaner presentation of the analysis. By suitable change of variables, we can recover our results for the multiplicative version (with slightly better dependency on the initial state).

[1] The primary focus in the literature has been on convex cost functions, because they model increasing marginal cost per unit utility, or, equivalently, decreasing marginal utility per unit cost, which is a feature of most economic environments. Also, Tullock contests with non-convex cost functions generally do not have a pure-strategy Nash equilibrium (PNE), and BR dynamics are not stable by definition when there does not exist a PNE; hence, such models need to be studied using alternative equilibrium concepts and learning dynamics.

[2] Any n-player game can be written as an $(n+1)$-player zero-sum game.

[3] An action profile $\boldsymbol{x}$ is an ϵ-approximate pure-strategy Nash equilibrium, for $\epsilon > 0$, if it satisfies $u_i(x_i, \boldsymbol{x}_{-i}) \geq u_i(x'_i, \boldsymbol{x}_{-i}) - \epsilon$, for every agent i and every action x'_i for agent i.

4.2.1 Continuous-Time BR Dynamics

Let $\boldsymbol{x}(t) = (x_i(t))_{i\in[n]}$ denote the action profile of the agents at time t in the BR dynamics. Similarly, let $s(t) = \sum_j x_j(t)$ and $s_{-i}(t) = \sum_{j\neq i} x_j(t)$. The continuous-time (or simply, continuous) BR dynamics starts from an initial profile $\boldsymbol{x}(0) = (x_i(0))_{i\in[n]} \in \mathbb{R}_{\geq 0}$. At time $t \geq 0$, each agent $i \in [n]$ continuously updates their action as

$$\frac{dx_i(t)}{dt} = BR_i(s_{-i}(t)) - x_i(t). \tag{4.1}$$

The continuous BR dynamics is a limiting case of the discrete BR dynamics when the step size goes to 0. In discrete BR, for a step size of $\Delta t = 1$, $x_i(t+1) = BR_i(s_{-i}(t)) \iff x_i(t+1) - x_i(t) = BR_i(s_{-i}(t)) - x_i(t)$ for any given agent i. For arbitrary $\Delta t \geq 0$, this dynamics can be generalized to $x_i(t + \Delta t) - x_i(t) = \Delta t(BR_i(s_{-i}(t)) - x_i(t)) \iff \frac{x_i(t+\Delta t)-x_i(t)}{\Delta t} = BR_i(s_{-i}(t)) - x_i(t)$. The continuous dynamics takes the limit $\Delta t \to 0$.

Such discrete to continuous limits are frequently used in economics and other mathematical sciences, where deterministic or stochastic difference equations are approximated as limit ordinary differential equations [52, 15, 86]. For example, gradient descent algorithms, including their stochastic versions, are studied by taking a limit on the step-size going to 0, and this continuous-time process is known as gradient flow [73, 16, 86].

4.3 Convergence of Continuous BR Dynamics

In this section, we prove that the continuous BR dynamics given in equation (4.1) rapidly converges to the unique pure-strategy Nash equilibrium, i.e., $\boldsymbol{x}(t) \to \boldsymbol{x}^*$ as $t \to \infty$, where $\boldsymbol{x}^*$ is the equilibrium. The conditions on the cost functions are the same as in Chapter 2, except that the Lipschitz condition is not required here. In particular, we assume that for every agent i, the cost function c_i is (a) twice differentiable; (b) zero cost for non-participation, $c_i(0) = 0$; (c) increasing, $c_i'(z) > 0$ for all $z \geq 0$; (d) weakly convex, $c_i''(z) \geq 0$ for all $z \geq 0$. These assumptions are standard in the literature, and without them, a pure-strategy Nash equilibrium may not exist (see, e.g., [106, Chapter 4], [46]).

Theorem 4.1. *The continuous best-response dynamics $\boldsymbol{x}(t)$ in Tullock contests with weakly convex cost functions converges to an ϵ-approximate pure-strategy Nash equilibrium in $O(\log(1/\epsilon))$ time. Further, there are instances where reaching an ϵ-approximate equilibrium takes $\Omega(\log(1/\epsilon))$ time.*

The $O(\log\log 1/\epsilon)$ bound is achieved in discrete-time setting only for the special case of 2 agents. For 3 or more agents, both results are $O(\log 1/\epsilon)$.

Note that linear dynamical systems converge in $\Theta(\log(1/\epsilon))$ time. In our dynamics, see equation (4.1), the $-x_i(t)$ term is linear but the $BR_i(s_{-i})$ term is non-linear, so a $\Theta(\log(1/\epsilon))$ convergence can be expected but is not obvious.

For $n \geq 3$ agents, the bound of $\Theta(\log(1/\epsilon))$ in Theorem 4.1 for non-homogeneous agents matches the bounds in Chapter 2 for discrete BR dynamics for homogeneous agents. For the special case of 2 homogeneous agents, the discrete BR dynamics has a faster convergence rate of $O(\log\log(1/\epsilon))$. The faster rate for 2 homogeneous agents is due to additional structure in the problem: both the first and second derivatives of the utility function of the agents vanish at the equilibrium, and this helps the discrete-time dynamics make rapid jumps towards the equilibrium.

We present the proof for the lower and upper bounds given in Theorem 4.1 separately. Let us start with the upper bound.

Proof of Theorem 4.1 (Upper Bound). To prove this convergence result, we use the potential function given in (4.2). This potential has been used previously to prove the convergence of stochastic fictitious play for two-player zero-sum games with finite action space; see the introduction for further discussion. For an action profile $\boldsymbol{x}$, let the potential function $V(\boldsymbol{x})$ be defined as:

$$\begin{aligned} V(\boldsymbol{x}) &= \sum_{i\in[n]} V_i(\boldsymbol{x}), \qquad (4.2)\\ \text{where } V_i(\boldsymbol{x}) &= \max_z u_i(z, s_{-i}) - u_i(x_i, s_{-i})\\ &= u_i(BR_i(s_{-i}), s_{-i}) - u_i(x_i, s_{-i}). \end{aligned}$$

Notice that $\max_z u_i(z, s_{-i})$ is not well-defined if $s_{-i} = 0$. In this case, as discussed in Section 4.2, we assume that $BR_i(0) \in (0, \frac{1}{2\max_i c_i'(0)}]$, and $V_i(\boldsymbol{x}) = u_i(BR_i(0), 0) - u_i(x_i, 0)$. We will prove that such profiles can only occur during a short initial phase of the BR dynamics.

$V_i(\boldsymbol{x})$ measures agent i's regret for playing x_i instead of the best possible action given s_{-i}, i.e., it is the amount of utility that agent i can increase by playing the BR instead of x_i. Notice that, by definition,

1. $V(\boldsymbol{x}) \geq 0$. Because $u_i(BR_i(s_{-i}), s_{-i}) \geq u_i(x_i, s_{-i}) \iff V_i(\boldsymbol{x}) \geq 0$ for every agent i and profile $\boldsymbol{x}$, which implies $V(\boldsymbol{x}) \geq 0$.

2. $V(\boldsymbol{x}) = 0$ at the equilibrium. $V(\boldsymbol{x}) = 0 \iff \boldsymbol{x} = \boldsymbol{x}^*$ because $V(\boldsymbol{x}) = 0 \iff V_i(\boldsymbol{x}) = 0, \forall i \in [n] \iff u_i(BR_i(s_{-i}), s_{-i}) = u_i(x_i, s_{-i}), \forall i \in [n]$.

Given the profile $\boldsymbol{x}(t)$, we can write the potential at time t as $V(\boldsymbol{x}(t)) = \sum_i V_i(\boldsymbol{x}(t))$. For conciseness, we denote $V_i(\boldsymbol{x}(t))$ by $V_i(t)$, or simply V_i, where the dependency on $\boldsymbol{x}(t)$ will be clear from the context; similarly, $V(\boldsymbol{x}(t))$ by $V(t)$ or V. Given the dynamics followed by $\boldsymbol{x}(t)$, equation (4.1), we can write the dynamics that $V(t)$ follows as

$$\frac{dV}{dt} = \sum_{i\in[n]} \frac{\partial V}{\partial x_i}\frac{dx_i}{dt}. \tag{4.3}$$

We next bound the time it takes to reach a state with two positive outputs, and we show that this property always holds thereafter.

Warm-Up Phase First, notice that if $x_i(\tau) > 0$ at some time point τ, then $x_i(t) > 0$ for all $t \geq \tau$ because $\frac{dx_i(t)}{dt} = BR(s_{-i}(t)) - x_i(t) \geq -x_i(t) \implies \frac{dx_i(t)}{x_i(t)} \geq -dt \implies x_i(t) \geq x_i(\tau)e^{-(t-\tau)} > 0$. So, once we reach a state with two agents i and $j \neq i$ with positive output, then these two agents will always have a positive output thereafter.

Say we start from a profile $\boldsymbol{x}(0) = \mathbf{0}$, i.e., all agents have 0 output initially. By our technical assumption discussed in Section 4.2, the action of an agent i is some small constant in $(0, \frac{1}{2\max_i c_i'(0)}]$, say η_i. So, $\frac{dx_i}{dt} = \eta_i > 0$ at time $t = 0$ for all i, which implies that at time $dt > 0$ with $dt \to 0$, we have $x_i(dt) > 0$, as required.

Let us now consider the case when there is only one agent i with $x_i(0) = \alpha > 0$, and all other agents $j \neq i$ have $x_j(0) = 0$. Now, if $x_i(0) = \alpha < \frac{1}{c_j'(0)}$ for some $j \neq i$, then $BR_j(s_{-j}(0)) = BR_j(\alpha) > 0$ because by the first-order condition, equation (2.2), for $z = 0$, we have

$$\frac{\partial u_i(z,\alpha)}{\partial z} = \frac{\alpha}{(z+\alpha)^2} - c_j'(z) = \frac{1}{\alpha} - c_j'(0) > 0.$$

So, at time 0, we have $\frac{dx_j}{dt} = BR_j(s_{-j}(0)) - x_j(0) = BR_j(\alpha) > 0$, which implies that at time $dt > 0$ with $dt \to 0$, we have $x_j(dt) > 0$.

Now, let us consider the case when $x_i(0) = \alpha \geq \max_{j\neq i} \frac{1}{c_j'(0)}$. Let $\beta = \max_{j\neq i} \frac{1}{c_j'(0)}$ for conciseness. Let us bound the time—denoted by T—it takes to reach $x_i(T) < \beta$. As $x_i(t) \geq \beta$ for all $t < T$, we have $\frac{dx_i(t)}{dt} = BR_i(s_{-i}(t)) - x_i(t) = BR_i(0) - x_i(t) = \eta_i - x_i(t) \leq \frac{\beta}{2} - x_i(t)$, where the last inequality holds because $\eta_i \leq \max_j \frac{1}{2c_j'(0)} \leq \max_{j\neq i} \frac{1}{2c_j'(0)} = \frac{\beta}{2}$. Using this, we get

$$\frac{dx_i(t)}{dt} \leq \frac{\beta}{2} - x_i(t) \implies \frac{dx_i(t)}{x_i(t) - \frac{\beta}{2}} \leq -dt \implies \ln\left(\frac{x_i(t) - \frac{\beta}{2}}{\alpha - \frac{\beta}{2}}\right) \leq -t,$$

which implies that to reach $x_i(T) < B$, it is sufficient to have $T > \ln\left(\frac{\alpha - \frac{\beta}{2}}{\beta - \frac{\beta}{2}}\right) = \ln\left(\frac{2\alpha}{\beta} - 1\right)$. Further notice that $V(\boldsymbol{x}(0)) = V_i(\boldsymbol{x}(0)) = (1 - c_i(\eta_i)) - (1 - c_i(x_i(0))) = x_i(0) - \eta_i \geq \alpha - \frac{\beta}{2}$. So, within $T = O(\log(V(0)))$ time, we get at least two agents with positive output, and this property holds thereafter.

Main Phase Given our analysis of the warm-up phase above, from here on we assume that there are always two agents with positive output. We next prove the following lemma about $V(t)$.

Lemma 4.2. *The potential $V(t) = V(\boldsymbol{x}(t))$ at any time t satisfies the following differential inequality*

$$\frac{dV}{dt} + V \leq -\sum_i \frac{y_i}{\sum_j y_j}\left(1 - \frac{\sum_j y_j}{y_i + s_{-i}}\right)^2 \leq 0,$$

where the dependency on t is suppressed, where $y_i(t) = BR_i(s_{-i}(t))$, and assuming that there are at least two agents with positive output in the profile $\boldsymbol{x}(t)$.[4]

Proof. Let us suppress the dependency on t, e.g., let us write $\boldsymbol{x}(t)$ as $\boldsymbol{x}$ and $V(t)$ as V. Let $y_i = BR_i(s_{-i})$ for conciseness.

Note that for general convex cost function c_i, we do not have a closed-form formula for BR_i. But from the first-order conditions, equation (2.4), we have

$$y_i > 0 \ \& \ \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=y_i} = 0, \text{ if } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=0} > 0;$$
$$y_i = 0 \ \& \ \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=y_i} \leq 0, \text{ if } \left.\frac{\partial u_i(z, s_{-i})}{\partial z}\right|_{z=0} \leq 0.$$

Now, $\frac{\partial u_i(z,s_{-i})}{\partial z} = \frac{s_{-i}}{(z+s_{-i})^2} - c_i'(z)$; plugging in $z = 0$ we get $\left.\frac{\partial u_i(z,s_{-i})}{\partial z}\right|_{z=0} = \frac{s_{-i}}{(0+s_{-i})^2} - c_i'(0) = \frac{1}{s_{-i}} - c_i'(0)$. The condition $\left.\frac{\partial u_i(z,s_{-i})}{\partial z}\right|_{z=0} > 0$ corresponds to $\frac{1}{s_{-i}} - c_i'(0) > 0 \iff s_{-i}c_i'(0) < 1$. Similarly, $\left.\frac{\partial u_i(z,s_{-i})}{\partial z}\right|_{z=0} \leq 0$ corresponds to $s_{-i}c_i'(0) \geq 1$. Using these, the first-order conditions at $y_i = BR_i(s_{-i})$ can be rewritten as

$$\frac{s_{-i}}{(y_i + s_{-i})^2} = c_i'(y_i), \text{ if } s_{-i}c_i'(0) < 1, \tag{4.4}$$

$$y_i = 0, \text{ if } s_{-i}c_i'(0) \geq 1. \tag{4.5}$$

[4] Only the $\frac{dV}{dt} + V \leq 0$ portion of the bound is used in the proof of this theorem; the detailed bound is stated explicitly because it is used in Lemma 4.3.

We can write V as

$$V = \sum_i V_i = \sum_i (u_i(y_i, s_{-i}) - u_i(x_i, s_{-i}))$$
$$= \sum_i \left(\frac{y_i}{y_i + s_{-i}} - c_i(y_i) - \frac{x_i}{x_i + s_{-i}} + c_i(x_i) \right)$$
$$= \sum_i \frac{y_i}{y_i + s_{-i}} - \sum_i c_i(y_i) + \sum_i c_i(x_i) - 1. \tag{4.6}$$

The time derivative of V w.r.t. t is $\frac{dV}{dt} = \sum_k \frac{\partial V}{\partial x_k} \frac{dx_k}{dt}$, where $\frac{dx_k}{dt} = y_k - x_k$. To write $\frac{\partial V}{\partial x_k}$, we need to know $\frac{\partial y_i}{\partial x_k}$ for all i and k. Note that $\frac{\partial y_i}{\partial x_k} = 0$ for $k = i$ and $\frac{\partial y_i}{\partial x_k} = \frac{dy_i}{ds_{-i}}$ for $k \neq i$. Due to the constraint that y_i (the best-response) is always non-negative, there is a point of non-differentiability at $y_i = 0$. In particular, if $c_i'(0) > 0$, then:

- If $s_{-i} > \frac{1}{c_i'(0)}$, then in the small neighborhood around s_{-i}, say $[s_{-i} - \delta, s_{-i} + \delta]$ for small $\delta > 0$, we will have the corresponding $y_i = 0$. So, $\frac{dy_i}{ds_{-i}} = 0$.

- If $s_{-i} < \frac{1}{c_i'(0)}$, then $y_i > 0$ and is governed by equation (4.4). We can differentiate this equation w.r.t. to s_{-i} to get

$$\frac{1}{(y_i + s_{-i})^2} - \frac{2s_{-i}}{(y_i + s_{-i})^3} = \left(c_i''(y_i) + \frac{2s_{-i}}{(y_i + s_{-i})^3} \right) \frac{dy_i}{ds_{-i}}$$
$$\implies \frac{dy_i}{ds_{-i}} = \frac{y_i - s_{-i}}{2s_{-i} + (y_i + s_{-i})^3 c_i''(y_i)}. \tag{4.7}$$

  If we take the limit $s_{-i} \uparrow \frac{1}{c_i'(0)}$, which implies that $y_i \downarrow 0$, we get

$$\frac{dy_i}{ds_{-i}} = \frac{-1}{2 + c_i''(0)/(c_i'(0))^2} < 0,$$

  where the last inequality is true because $c_i'(0) > 0$ and $c_i''(0) \geq 0$. On the other hand, we know that the magnitude is bounded $\frac{dy_i}{ds_{-i}} = \frac{-1}{2+c_i''(0)/(c_i'(0))^2} \geq \frac{-1}{2}$.

The above two cases tell us that $\frac{dy_i}{ds_{-i}}$ at $s_{-i} = \frac{1}{c_i'(0)}$ has a left limit strictly less than 0 but a right limit equal to 0, so we have non-differentiability at $s_{-i} = \frac{1}{c_i'(0)}$. But as $\frac{dy_i}{ds_{-i}}$ is bounded near $s_{-i} = \frac{1}{c_i'(0)}$, we can define it to be equal to some finite value in $[\frac{-1}{2+c_i''(0)/(c_i'(0))^2}, 0]$ at $s_{-i} = \frac{1}{c_i'(0)}$, which is sufficient for our analysis. An alternate analysis can be done using the *envelop theorem* of Milgrom and Segal [78, Theorem 2] to arrive at the same result. Taking partial derivative of V w.r.t. x_k we get

$$\frac{\partial V}{\partial x_k} = \frac{\partial}{\partial x_k} \left(\sum_i \frac{y_i}{y_i + s_{-i}} - \sum_i c_i(y_i) + \sum_i c_i(x_i) - 1 \right)$$
$$= \sum_i \left(\frac{s_{-i}}{(y_i + s_{-i})^2} - c_i'(y_i) \right) \frac{\partial y_i}{\partial x_k} - \sum_{i \neq k} \frac{y_i}{(y_i + s_{-i})^2} + c_k'(x_k).$$

From the discussion above, we know that either $\frac{\partial y_i}{\partial x_k} = 0$ or $\frac{s_{-i}}{(y_i+s_{-i})^2} = c_i'(y_i)$ and $\frac{\partial y_i}{\partial x_k}$ is bounded. In either case, we have $\left(\frac{s_{-i}}{(y_i+s_{-i})^2} - c_i'(y_i)\right)\frac{\partial y_i}{\partial x_k} = 0$, so

$$\frac{\partial V}{\partial x_k} = c_k'(x_k) - \sum_{i \neq k} \frac{y_i}{(y_i + s_{-i})^2}. \tag{4.8}$$

Putting together, we can write $\frac{dV}{dt}$ as

$$\begin{aligned}\frac{dV}{dt} = \sum_k \frac{\partial V}{\partial x_k}\frac{dx_k}{dt} &= \sum_k (y_k - x_k)c_k'(x_k) - \sum_k (y_k - x_k)\sum_{i \neq k}\frac{y_i}{(y_i + s_{-i})^2} \\ &= \sum_i (y_i - x_i)c_i'(x_i) - \sum_i \frac{y_i(\sigma - y_i - s_{-i})}{(y_i + s_{-i})^2},\end{aligned}$$

where $\sigma = \sum_k y_k$. Adding V and $\frac{dV}{dt}$ together, we get

$$\begin{aligned}&V + \frac{dV}{dt} \\ =&\sum_i \frac{y_i}{y_i + s_{-i}} - \sum_i c_i(y_i) + \sum_i c_i(x_i) - 1 + \sum_i (y_i - x_i)c_i'(x_i) - \sum_i \frac{y_i(\sigma - y_i - s_{-i})}{(y_i + s_{-i})^2} \\ =& -1 + \sum_i \underbrace{\left(-c_i(y_i) + c_i(x_i) + (y_i - x_i)c_i'(x_i)\right)}_{\leq 0 \text{ as } c_i \text{ is convex}} + \sum_i \frac{2y_i}{y_i + s_{-i}} - \sum_i \frac{y_i\sigma}{(y_i + s_{-i})^2}.\end{aligned}$$

Now, let $p_i = \frac{y_i}{\sum_j y_i} = \frac{y_i}{\sigma}$ and $q_i = \frac{s_{-i}}{\sigma}$. Notice that $\sum_i p_i = 1$ and, for all i, $p_i \geq 0$ and $q_i \geq 0$. Plugging this into the inequality above, we get

$$\begin{aligned}V + \frac{dV}{dt} &\leq -1 + \sum_i \frac{2y_i}{y_i + s_{-i}} - \sum_i \frac{y_i\sigma}{(y_i + s_{-i})^2} = -1 + \sum_i \frac{2p_i}{p_i + q_i} - \sum_i \frac{p_i}{(p_i + q_i)^2} \\ &= \sum_i p_i\left(-1 + \frac{2}{p_i + q_i} - \frac{1}{(p_i + q_i)^2}\right) = -\sum_i p_i\left(1 - \frac{1}{p_i + q_i}\right)^2 \leq 0. \qquad \square\end{aligned}$$

Let us now use Lemma 4.2 to get the desired rate of convergence upper bound. We use standard Lyapunov arguments:

$$\frac{dV(t)}{dt} + V(t) \leq 0 \implies \frac{dV(t)}{V(t)} \leq -dt \implies V(t) \leq V(0)e^{-t}.$$

For any $t \geq \ln\left(\frac{1}{\epsilon}\right) + \ln(V(0))$, we get $V(t) \leq \epsilon$, which implies that for every agent i we have $V_i(t) = V_i(\boldsymbol{x}(t)) \leq \epsilon \iff u_i(\boldsymbol{x}(t)) \geq u_i(BR_i(s_{-i}(t)), s_{-i}(t)) - \epsilon$. So, we are at an ϵ-approximate equilibrium. This completes the proof for the upper bound. $\square$

Proof of Theorem 4.1 (Lower Bound). We provide an example where it takes $\Omega(\log\left(\frac{1}{\epsilon}\right) + \log(V(0)))$ time to converge to an ϵ-approximate equilibrium.

Let there be $n = 2$ homogeneous agents with linear cost function $c_1(y) = c_2(y) = y/4$ for any $y \geq 0$. It can be easily derived that the unique equilibrium is $\boldsymbol{x}^* = (1, 1)$ and that there is a closed-form formula for the best-response $BR_i(s_{-i}) = 2\sqrt{s_{-i}} - s_{-i}$ (see, e.g., Vojnović [106]).

Let $\boldsymbol{x}(0) = (y(0), y(0))$, where we assume that $y(0)$ is sufficiently large and far away from the equilibrium value 1. As the two players are homogeneous and start from the same action, they will maintain the same action $\boldsymbol{x}(t) = (y(t), y(t))$, for some $y(t)$, for all $t \geq 0$. We suppress the dependency on t to avoid clutter. Let us track the evolution of y. From equation (4.1), we have

$$\frac{dy}{dt} = \frac{dx_i}{dt} = BR_i(s_{-i}) - x_i = BR_i(y) - y = (2\sqrt{y} - y) - y = 2(\sqrt{y} - y).$$

Let us now compute the potential function V. We have

$$\begin{aligned} V &= \sum_i (u_i(BR_i(s_{-i}), s_{-i}) - u_i(x_i, s_{-i})) = 2\left(u_1(2\sqrt{y} - y, y) - u_1(y, y)\right) \\ &= 2\left(\frac{2\sqrt{y} - y}{2\sqrt{y} - y + y} - \frac{2\sqrt{y} - y}{4} - \left(\frac{y}{y + y} - \frac{y}{4}\right)\right) = 1 + y - 2\sqrt{y} = (\sqrt{y} - 1)^2. \end{aligned}$$

Further, we can find the rate of change of V using the rate of change of y as

$$\begin{aligned} \frac{dV}{dy} &= \frac{d}{dy}(\sqrt{y} - 1)^2 = \frac{2(\sqrt{y} - 1)}{2\sqrt{y}} = \frac{\sqrt{y} - 1}{\sqrt{y}}, \\ \frac{dV}{dt} &= \frac{dV}{dy}\frac{dy}{dt} = \frac{\sqrt{y} - 1}{\sqrt{y}} 2(\sqrt{y} - y) = -2(\sqrt{y} - 1)^2 \\ &\implies \frac{dV}{dt} = -2V \implies V(t) = V(0)e^{-2t}. \end{aligned}$$

By the definition of V and the symmetry of the two agents, at an ϵ-approximate equilibrium, $V(t) = 2\epsilon$. So, it takes exactly $t = \frac{1}{2}\ln(\frac{1}{2\epsilon}) + \frac{1}{2}\ln(V(0)) = \Omega(\log\left(\frac{1}{\epsilon}\right) + \log(V(0)))$ time to reach the ϵ-approximate equilibrium. $\square$

4.4 Discrete-Time, Equilibrium Computation, and Related Dynamics

In this section, we consider discrete-time BR dynamics. We also provide an algorithm for computing an approximate equilibrium based on such dynamics. Proofs are given in Section 4.6.

Let us consider a modification to the original Tullock contest model we have studied till now. We assume that each agent i must always play an action $x_i \geq x_{\min}$

instead of $x_i \geq 0$, for some $x_{\min} \geq 0$. Notice that $x_{\min} = 0$ corresponds to the original model, while a $x_{\min} > 0$ says that any participant in the contest must have a positive minimum output. An assumption of $x_{\min} > 0$ may be plausible in practical scenarios where there is a positive cost of participation (*showing up* for the game). We also normalize the cost functions and assume that $\min_i c_i(1) = 1$ for all i; this ensures that any rational agent will always play an action ≤ 1. We also assume that the second derivative and the ratio of the first derivatives of the cost functions are bounded: $\frac{\max_{i,z\in[x_{\min},1]} c_i'(z)}{\min_{i,z\in[x_{\min},1]} c_i'(z)} = B_1$ and $\max_{i,z\in[x_{\min},1]} c_i''(z) = B_2$.

The first-order conditions for the case when $x_{\min} > 0$ are similar to the ones given in (2.4) except that the critical point above which the first-order condition is satisfied with equality is $x_{\min}$ instead of 0. The analysis for the continuous BR dynamics for this model is also analogous to the analysis for Theorem 4.1.

Let us now consider BR dynamics in this model with the step-size, say α_t, which may possibly vary with time t, and is small but not necessarily going to 0. In particular,

$$x_i(t + \alpha_t) = x_i(t) + \alpha_t \cdot (BR_i(s_{-i}(t)) - x_i(t)). \tag{4.9}$$

The continuous-time BR dynamics corresponds to equation (4.9) with $\alpha_t \to 0$. We aim to find bounds on α_t that ensure convergence.

Lemma 4.3. *For a profile $\boldsymbol{x}$, let $H(\boldsymbol{x})$ be defined as*

$$H(\boldsymbol{x}) = \frac{\frac{B_2}{2}\sum_i (y_i - x_i)^2 + \sum_{i\in E} \frac{(\sigma - y_i - s_{-i})^2}{s_{-i}^2}}{\sum_i \frac{y_i(\sigma - y_i - s_{-i})^2}{\sigma(y_i + s_{-i})^2}},$$

where $y_i = BR_i(s_{-i})$ and $\sigma = \sum_i y_i$. If the step-size at time t is bounded above by $1/\max(2, H(\boldsymbol{x}(t)))$, then the BR dynamics converges to the unique equilibrium. In particular, for $0 < \alpha_t \leq 1/\max(2, H(\boldsymbol{x}(t)))$, we have $V(t + \alpha_t) \leq (1 - \alpha_t)V(t)$.

If $x_{\min}$ is assumed to be strictly positive, then we can upper bound $H(\boldsymbol{x})$ as a function of $x_{\min}$.

Lemma 4.4. *If $x_{\min} > 0$, then $H(\boldsymbol{x}) = O\left(\frac{n(1+B_2)}{x_{\min}^3}\right)$ for all $\boldsymbol{x}$, which implies that the dynamics reaches an ϵ-approximate equilibrium in $O\left(\frac{1}{\alpha}\log\left(\frac{V(0)}{\epsilon}\right)\right)$ steps with a suitable step-size $\alpha = \Theta\left(\frac{x_{\min}^3}{n(1+B_2)}\right)$.*

In Lemma 4.4, the number of steps till convergence, $O\left(\frac{1}{\alpha}\log\left(\frac{V(0)}{\epsilon}\right)\right)$, has a $1/\alpha$ dependency on α. This follows directly from the $V(t+\alpha_t) \leq (1-\alpha_t)V(t)$ condition on the potential derived in Lemma 4.3.

Notice that the bound in Lemma 4.4 depends upon the n, B_2, and $x_{\min}$. The dependency on B_2 shows up as $(1+B_2)$, which is not very significant for most practical problems because: even for extremely convex functions like x_i^β for a large constant β, B_2 is bounded by the constant β^2. So, the convergence time has a dependency of $(1+\beta^2)$, which is a constant. A much more involved analysis may help improve this dependency. On the other hand, the dependency on n is essential, as highlighted by Lemma 4.5 below.

Lemma 4.5. *If the step-size is not $O(1/n)$, then there are instances with linear and homogeneous cost functions where the dynamics does not converge.*

Note that although the bound in Lemma 4.4 does not depend upon B_1 (the ratio of the first-derivatives of the cost functions of the agents, which measures the relative skills of the agents), B_1 may be implicit in $x_{\min}$. If $x_{\min}$ is not sufficiently small, e.g., if $x_{\min} = \omega(1/B_1^2)$, then the equilibrium when the agents are restricted to play $x_i \geq x_{\min}$ may be different from the equilibrium when the agents can play less than $x_{\min}$. For example, for two agents with linear cost functions $c_1(x_1) = x_1$ and $c_2(x_2) = \beta x_2$, where $\beta \geq 1$, the unique equilibrium is $\boldsymbol{x}^* = \left(\frac{\beta}{(1+\beta)^2}, \frac{1}{(1+\beta)^2}\right)$ if the agents can play any $x_i \geq 0$. Note that $B_1 = \beta$, so the equilibrium output of agent-2 is $\Theta(1/B_1^2)$. On the other hand, if the agents are restricted to play $x_i \geq x_{\min} = \omega(1/B_1^2)$, then the equilibrium will be forced to be different from $\boldsymbol{x}^*$. Given this observation, it would be natural to assume that $x_{\min}$ is small enough; in particular, $x_{\min} = O(1/B_1^2)$. Moreover, the dependency on B_1 is unavoidable, as formalized by Lemma 4.6 below.

Lemma 4.6. *If the step-size is not $O(1/B_1)$, then there are instances with two agents and linear cost functions where the dynamics does not converge.*

If $x_{\min} = 0$, then our results do not provide a lower bound on the step-size that is independent of the action profile $\boldsymbol{x}(t)$. In particular, Lemma 4.3 does not directly imply such a bound because $H(\boldsymbol{x})$ may be unbounded for some profiles $\boldsymbol{x}$, then the step-size recommended by the lemma at $\boldsymbol{x}$ to ensure convergence goes to 0. Indeed, such a lower bound might not exist. On the other hand, even in the case of $x_{\min} = 0$, we can simulate with a pseudo $\hat{x}_{\min} = \Theta(\epsilon)$ to compute an equilibrium in $poly(1/\epsilon, n, B_1, B_2)$ steps as shown below.

Algorithm Let us construct a modified game with a pseudo lower bound on the outputs of the agents: $\hat{x}_{\min} = \epsilon/(4B_1)$. We simulate the BR dynamics in this game with a step-size of $\alpha = \Theta\left(\frac{\hat{x}_{\min}^3}{n(1+B_2)}\right)$, as recommended by Lemma 4.4, to compute an $(\epsilon/2)$-approximate equilibrium of this modified game in $O\left(\frac{1}{\alpha}\log\left(\frac{V(0)}{\epsilon}\right)\right)$ steps. Let this approximate equilibrium be $\hat{\boldsymbol{x}}$. At $\hat{\boldsymbol{x}}$, all agents have a regret of at most $\epsilon/2$ assuming that they can only play above $\hat{x}_{\min}$. By playing below $\hat{x}_{\min}$, they can further increase their utility by at most $c_i(\hat{x}_{\min}) - c_i(0) \leq B_1\hat{x}_{\min}/(1-\hat{x}_{\min}) \leq 2B_1\hat{x}_{\min} \leq \epsilon/2$. So, at $\hat{\boldsymbol{x}}$, the total regret of any agent in the original game is at most ϵ, as required.

Best-Response to Empirical Average Let us consider a discrete-time dynamics with a step-size of $\alpha_t = 1$, but where the agents best-respond to the empirical average action of the other agents. Let $\overline{x}_i(t) = \frac{1}{t}\sum_{\tau=1}^{t} x_i(t)$ and $\overline{s}_{-i}(t) = \sum_{j\neq i}\overline{x}_j(t)$. Formally, the dynamics is defined as follows: the action of agent i at time $t+1$ is

$$x_i(t+1) = BR_i(\overline{s}_{-i}(t)), \qquad \forall i \in [n], t \in \mathbb{Z}_{\geq 0} \tag{4.10}$$

Given this, we can write the updated empirical average at time $t+1$ as

$$\overline{x}_i(t+1) = \overline{x}_i(t) + \frac{1}{t+1}(BR_i(\overline{s}_{-i}(t)) - \overline{x}_i(t)).$$

Notice that $\overline{x}_i(t)$ tracks a BR dynamics with a sequence of decreasing step-sizes that correspond to the harmonic sequence $(\frac{1}{t})_{t\in\mathbb{Z}_{\geq 1}}$. The harmonic sequence satisfies the following crucial properties: as $t \to \infty$, the sequence $\frac{1}{t} \to 0$ but the series $\sum_{k=1}^{t}\frac{1}{k} \to \infty$. These two properties ensure that the dynamics converges for the case $x_{\min} > 0$ using Lemma 4.3. In fact, we can generalize this dynamics to a weighted average, where the step-size at time t is η_t, and $\overline{x}_i(t)$ follows

$$\overline{x}_i(t+1) = \overline{x}_i(t) + \eta_t(BR_i(\overline{s}_{-i}(t)) - \overline{x}_i(t)). \tag{4.11}$$

Lemma 4.7. *If $x_{\min} > 0$, then a dynamics that evolves according to* (4.11) *converges if the sequence of step-sizes $(\eta_t)_{t\in\mathbb{Z}_{\geq 1}}$ satisfies: as $t \to \infty$, the sequence $\eta_t \to 0$ but the series $\sum_{k=1}^{t}\eta_k \to \infty$.*

Lemma 4.7 builds directly upon Lemma 4.4. In Lemma 4.4, we showed convergence for small but fixed step sizes. The $\eta_t \to 0$ condition ensures that the step size eventually becomes sufficiently small. The $\sum_{k=1}^{t}\eta_k \to \infty$ condition ensures that, even after the step size has become sufficiently small, it is large enough that the small steps accumulate to establish convergence. Similar results are known for many discrete dynamical systems, where convergence using small but fixed step sizes implies convergence using time dependent steps with the properties in Lemma 4.7, e.g., gradient descent [16, 86] and replicator dynamics [52].

In addition to $\eta_t = 1/t$, other examples of step-size sequences that lead to convergence are $\eta_t = 1/t^r$ for $r \in (0, 1]$ and $\eta_t = 1/\log(1+t)$. Note that convergence of $\overline{\boldsymbol{x}}(t)$ also implies convergence of $\boldsymbol{x}(t)$.

Agents Move at Different Rates When the agents move at different rates that can depend upon time, it is easy to see that the continuous-time dynamics may not converge, because we can simulate a discrete-time dynamics using the continuous-time dynamics by adjusting the rates. In particular, consider the example for two agents presented in Lemma 4.6 where a discrete-time dynamics with a step-size of 1/2 goes into a cycle. We can simulate this dynamics by slowing down the first agent and allowing the second agent to move, and then slowing down the second agent and allowing the first agent to move, and repeating the process. In a similar manner, the cycles presented in Chapter 2, where only one agent moves at a time, can also be simulated.

On the other hand, simple modifications to the continuous BR dynamics converge. For example, consider the following dynamics

$$\frac{dx_i}{dt} = f(\boldsymbol{x}, t)(BR_i(s_{-i}) - x_i),$$

where $f(\boldsymbol{x}, t)$ is some positive scaling factor that may depend upon the profile $\boldsymbol{x}$ and the time t but is the same for all agents. This dynamics follows exactly the same path as the original continuous BR dynamics (albeit at a different speed), and thus converges. Let us consider an additional modification to the dynamics

$$\frac{dx_i}{dt} = \eta_i f(\boldsymbol{x}, t)(BR_i(s_{-i}) - x_i),$$

where $\eta_i > 0$ is some positive constant specific to agent i. This dynamics speeds up certain agents relative to others. Studying this dynamics and similar extensions and showing their convergence (or non-convergence) is an important direction for future research.

4.5 Conclusion

We showed that the continuous BR dynamics, which is motivated by the observation that in certain applications the agents change their actions slowly compared to the feedback they receive from others, converges to the unique equilibrium in Tullock contests with convex costs. We then extended these convergence results to related discrete dynamics with small step sizes. These results indicate that we can expect Tullock contests with convex costs to reach equilibrium in a decentralized manner.

One open problem is to show convergence (or non-convergence) of the discrete dynamics with small step sizes when $x_{\min} = 0$ (see Section 4.4 for details). A different direction is to study the case when the agents move at rates. If the relative rates of the agents can change arbitrarily over time, we may have non-convergence. On the other hand, if the rates of the agents are similar to that of the continuous BR dynamics, we can show convergence. It would be valuable to analyze scenarios between these two extreme cases (see Section 4.4 for further discussion).

Another unexplored direction is to study learning in Tullock contests when the agents get only probabilistic feedback. In many practical applications, the proportional allocation function corresponds to the probability of allocating an indivisible item (and not the fraction of the item allocated to the agent). Here, an agent may only know their own actions and whether or not they won the indivisible item but may not know the actions of others. This line of work combines multi-armed bandits with game theory. It is related to learning in stochastic games [96], but most of the results on learning in general stochastic games indicate hardness/impossibility. The Tullock contest model is very structured, and we may expect some positive results. We are currently looking at a Bayesian best-response dynamics model, where each agent maintains a belief over the total output of other agents, best-responds to this belief, and updates their belief based on whether or not they win a round. We have preliminary results about this dynamics, particularly equivalence between this dynamics and other relatively simpler dynamics, but the problem is largely open.

4.6 Omitted Proofs

Proof of Lemma 4.3. We shall suppress the dependency on t for conciseness, e.g., write $\boldsymbol{x}(t)$ as $\boldsymbol{x}$ and $x_i(t)$ as x_i. Let $y_i = BR_i(s_{-i})$. Given the minimum output level $x_{\min} \geq 0$, y_i satisfies the first-order conditions given below, which can be derived following steps similar to the derivation of conditions (4.4) and (4.5).

$$\frac{s_{-i}}{(y_i + s_{-i})^2} = c_i'(y_i), \qquad \text{if } s_{-i}c_i'(x_{\min}) < 1, \tag{4.12}$$

$$y_i = x_{\min}, \qquad \text{if } s_{-i}c_i'(x_{\min}) \geq 1. \tag{4.13}$$

From equation (4.6), we know that the potential $V(t) = V(\boldsymbol{x}(t))$ is given by

$$V(t) = \sum_i \frac{y_i}{y_i + s_{-i}} - \sum_i c_i(y_i) + \sum_i c_i(x_i) - 1.$$

In Lemma 4.2, we showed that $\frac{dV}{dt} < 0$, i.e., $\lim_{\Delta t \to 0} \frac{V(t+\Delta t) - V(t)}{\Delta t} < 0$, which implies $\lim_{\Delta t \downarrow 0}(V(t+\Delta t) - V(t)) < 0$. We shall now consider small positive values for Δt, not necessarily limiting to 0, such that we can still guarantee that $V(t + \Delta t) - V(t) < 0$.

Let $\Delta t = \alpha$ be the step-size. Given the current profile $\boldsymbol{x}$ and the best-response profile $\boldsymbol{y} = (BR_i(s_{-i}))_{i\in[n]}$, the profile after the α-step is $\alpha(\boldsymbol{y} - \boldsymbol{x}) + \boldsymbol{x}$. We want to show that $V(t + \alpha) = V(\alpha(\boldsymbol{y} - \boldsymbol{x}) + \boldsymbol{x}) < V(t) = V(\boldsymbol{x})$. The multivariate Taylor's expansion of $V(\alpha(\boldsymbol{y} - \boldsymbol{x}) + \boldsymbol{x})$ w.r.t. $V(\boldsymbol{x})$ with a second-order error term is given by

$$V(\alpha(\boldsymbol{y} - \boldsymbol{x}) + \boldsymbol{x}) = V(\boldsymbol{x}) + \alpha(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla V(\boldsymbol{x}) + \frac{\alpha^2}{2}(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla^2 V(\hat{\boldsymbol{x}})(\boldsymbol{y} - \boldsymbol{x}), \quad (4.14)$$

where $\hat{\boldsymbol{x}} = \boldsymbol{x} + \beta(\boldsymbol{y} - \boldsymbol{x})$ for some value $\beta \in [0, \alpha]$, and where $\nabla V(\boldsymbol{z}) = (\frac{\partial V(\boldsymbol{z})}{\partial z_i})_{i\in[n]}$ and $\nabla^2 V(\boldsymbol{z}) = (\frac{\partial^2 V(\boldsymbol{z})}{\partial z_i \partial z_j})_{i,j\in[n]}$ for any given profile $\boldsymbol{z}$.

Let us first compute $\nabla^2 V(\boldsymbol{x})$.[5] As $s_{-i} = \sum_{j\neq i} x_j$, so $\frac{ds_{-i}}{dx_i} = 0$ and $\frac{ds_{-i}}{dx_j} = 1$ for all $j \neq i$. From equation (4.7), we have

$$\frac{\partial y_i}{\partial x_j} = \frac{y_i - s_{-i}}{2s_{-i} + (y_i + s_{-i})^3 c_i''(y_i)},$$

for $j \neq i$ and $y_i > x_{\min}$, and 0 otherwise. As derived in equation (4.8), we know that

$$\frac{\partial V}{\partial x_i} = c_i'(x_i) - \sum_{k\neq i} \frac{y_k}{(y_k + s_{-k})^2}.$$

Let $E = \{k \in [n] \mid y_k > x_{\min}\}$. Differentiating again w.r.t. x_j, if $j = i$, we have

$$\begin{aligned}\frac{\partial^2 V}{\partial x_i^2} &= c_i''(x_i) + \sum_{k\in E, k\neq i} \left(\frac{2y_k}{(y_k + s_{-k})^3} + \left(\frac{2y_k}{(y_k + s_{-k})^3} - \frac{1}{(y_k + s_{-k})^2} \right) \frac{\partial y_i}{\partial x_k} \right) \\ &= c_i''(x_i) + \sum_{k\in E, k\neq i} \frac{2y_k}{(y_k + s_{-k})^3} + \sum_{k\in E, k\neq i} \frac{\partial y_i}{\partial x_k} \frac{y_k - s_{-k}}{(y_k + s_{-k})^3}.\end{aligned}$$

[5] As discussed in Lemma 4.2, all points except when $s_{-i} = \frac{1}{c_i'(x_{\min})}$ are continuously differentiable (and can be shown to be twice continuously differentiable by extending the same argument). When $s_{-i} = \frac{1}{c_i'(x_{\min})}$, the derivative of y_i w.r.t. s_{-i} is not well-defined. However, the derivative is well-defined and bounded at all points in its neighborhood. This allowed us to extend the analysis for the differentiable points to this non-differentiable point. A similar but more detailed analysis can be done for computing $\nabla^2 V(\boldsymbol{x})$ to get the same results, but here let us restrict our focus on the *generic* case of $s_{-i} \neq \frac{1}{c_i'(x_{\min})}$.

Plugging in the value of $\frac{\partial y_i}{\partial x_k}$ for $k \in E$, we get

$$\frac{\partial^2 V}{\partial x_i^2} = c_i''(x_i) + \sum_{k\in E, k\neq i} \frac{2y_k}{(y_k+s_{-k})^3} + \sum_{k\in E,k\neq i} \frac{(y_k - s_{-k})^2}{(y_k+s_{-k})^3(2s_{-k} + (y_k+s_{-k})^3 c_k''(y_k))}.$$

Let $\eta_k = (y_k + s_{-k})^3 c_k''(y_k)$. The above equation can be rewritten as

$$\begin{aligned}\frac{\partial^2 V}{\partial x_i^2} &= c_i''(x_i) + \sum_{k\in E, k\neq i} \frac{2y_k}{(y_k+s_{-k})^3} + \sum_{k\in E,k\neq i} \frac{(y_k - s_{-k})^2}{(y_k+s_{-k})^3(2s_{-k} + \eta_k)}\\ &= c_i''(x_i) + \sum_{k\in E,k\neq i} \frac{(y_k+s_{-k})^2 + 2y_k\eta_k}{(y_k+s_{-k})^3(2s_{-k}+\eta_k)}.\end{aligned}$$

Let $a_i = c_i''(x_i)$. Let $b_i = \frac{(y_i+s_{-i})^2+2y_i\eta_i}{(y_i+s_{-i})^3(2s_{-i}+\eta_i)}$ if $i \in E$ and $b_i = 0$ if $i \notin E$. Notice that $a_i \geq 0$ and $b_i \geq 0$ for all $i \in [n]$. We have

$$\frac{\partial^2 V}{\partial x_i^2} = a_i + (\sum_i b) - b_i.$$

Following a similar sequence of steps, we can compute $\frac{\partial^2 V}{\partial x_i \partial x_j}$ for $j \neq i$ as

$$\frac{\partial^2 V}{\partial x_i \partial x_j} = (\sum_i b) - b_i - b_j.$$

Let $\boldsymbol{a} = (a_i)_{i\in[n]}$ and $\boldsymbol{b} = (b_i)_{i\in[n]}$. Let $\boldsymbol{e} = (1, 1, \ldots, 1)$ denote the n-dimensional vector of all 1s. Let A denote the $n \times n$ diagonal matrix with $A_{ii} = a_i + b_i$ and $A_{ij} = 0$ for $j \neq i$. Putting everything together, we can write $\nabla^2 V(\boldsymbol{x})$ as

$$\nabla^2 V(\boldsymbol{x}) = (\sum_i b)\boldsymbol{e}\boldsymbol{e}^\intercal - \boldsymbol{b}\boldsymbol{e}^\intercal - \boldsymbol{e}\boldsymbol{b}^\intercal + A. \tag{4.15}$$

Let us now consider a vector $\boldsymbol{w} \in \mathbb{R}^n$. We can compute $\boldsymbol{w}^\intercal \nabla^2 V(\boldsymbol{x})\boldsymbol{w}$ as

$$\begin{aligned}\boldsymbol{w}^\intercal \nabla^2 V(\boldsymbol{x})\boldsymbol{w} &= \boldsymbol{w}^\intercal((\sum_i b)\boldsymbol{e}\boldsymbol{e}^\intercal - \boldsymbol{b}\boldsymbol{e}^\intercal - \boldsymbol{e}\boldsymbol{b}^\intercal + A)\boldsymbol{w}\\ &= (\sum_i b_i)(\sum_i w_i)^2 - (\sum_i b_i w_i)(\sum_i w_i) - (\sum_i w_i)(\sum_i b_i w_i) + \sum_i (a_i + b_i) w_i^2\\ &= \sum_i b_i \left((\sum_j w_j)^2 - 2w_i(\sum_j w_j) + w_i^2 \right) + \sum_i a_i w_i^2\\ &= \sum_i b_i (\sum_{j\neq i} w_j)^2 + \sum_i a_i w_i^2.\end{aligned}$$

Now, we know that $b_i = \frac{(y_i+s_{-i})^2+2y_i\eta_i}{(y_i+s_{-i})^3(2s_{-i}+\eta_i)}$ for $i \in E$. Let us find a suitable upper bound on b_i. Let the function $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be defined as

$$g(z) = \frac{(y_i + s_{-i})^2 + 2y_i z}{(y_i + s_{-i})^3(2s_{-i} + z)}.$$

Notice that $b_i = g(\eta_i)$. Differentiating g w.r.t. z, we get

$$\begin{aligned}
\frac{dg}{dz} &= \frac{2y_i}{(y_i + s_{-i})^3(2s_{-i} + z)} - \frac{(y_i + s_{-i})^2 + 2y_i z}{(y_i + s_{-i})^3(2s_{-i} + z)^2} \\
&= \frac{2y_i(2s_{-i} + z) - (y_i + s_{-i})^2 - 2y_i z}{(y_i + s_{-i})^3(2s_{-i} + z)^2} \\
&= \frac{-(y_i - s_{-i})^2}{(y_i + s_{-i})^3(2s_{-i} + z)^2}.
\end{aligned}$$

So, $\frac{dg}{dz}$ is less than or equal to 0 for all $y_i \geq 0$, $s_{-i} \geq 0$, and $z \geq 0$, which are satisfied by the definition of y_i, s_{-i}, and z. So, $g(z)$ is maximized at $z = 0$, which implies that

$$b_i = g(\eta_i) \leq \max_{z \geq 0} g(z) = g(0) \leq \frac{1}{2(y_i + s_{-i})s_{-i}}.$$

Plugging back this bound on b_i and the value of a_i to $\boldsymbol{w}^\intercal \nabla^2 V(\boldsymbol{x})\boldsymbol{w}$, we get

$$\boldsymbol{w}^\intercal \nabla^2 V(\boldsymbol{x})\boldsymbol{w} \leq \sum_{i \in E} \frac{(\sum_{j \neq i} w_j)^2}{2(y_i + s_{-i})s_{-i}} + \sum_i c_i''(x_i) w_i^2.$$

In the Taylor's expansion in equation (4.14), $\nabla^2 V$ is evaluated at $\hat{\boldsymbol{x}}$. Let $\hat{s}_{-i} = \sum_{j \neq i} \hat{x}_j$ and $\hat{y}_i = BR_i(\hat{s}_{-i})$. Setting the vector $\boldsymbol{w} = \boldsymbol{y} - \boldsymbol{x}$, we get

$$(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla^2 V(\hat{\boldsymbol{x}})(\boldsymbol{y} - \boldsymbol{x}) \leq \sum_{i \in E} \frac{(\sum_{j \neq i}(y_i - x_i))^2}{2(\hat{y}_i + \hat{s}_{-i})\hat{s}_{-i}} + \sum_i c_i''(\hat{x}_i)(y_i - x_i)^2.$$

Notice that $\hat{s}_{-i} = \alpha \sum_{j \neq i} y_j + (1 - \alpha)s_{-i} \geq (1 - \alpha)s_{-i}$. So, $(\hat{y}_i + \hat{s}_{-i})\hat{s}_{-i} \geq (1 - \alpha)^2 s_{-i}$. As α is a small constant, $(1 - \alpha)^2$ is bounded away from 0, e.g., if $\alpha \leq 1/2$, then $(1 - \alpha)^2 \geq 1/4$. Let $\sigma = \sum_i y_i$. For $\alpha \leq 1/2$, we get

$$(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla^2 V(\hat{\boldsymbol{x}})(\boldsymbol{y} - \boldsymbol{x}) \leq \sum_i \frac{2(\sigma - y_i - s_{-i}))^2}{s_{-i}^2} + \sum_i B_2(y_i - x_i)^2,$$

as $c_i''(z) \leq B_2$ for any z. Let α be equal to the smaller among $1/2$ and

$$\alpha \leq \frac{\sum_i \frac{y_i(\sigma - y_i - s_{-i})^2}{\sigma(y_i + s_{-i})^2}}{\frac{B_2}{2}\sum_i (y_i - x_i)^2 + \sum_{i \in E} \frac{(\sigma - y_i - s_{-i})^2}{s_{-i}^2}}.$$

Notice that $\sum_i \frac{y_i(\sigma - y_i - s_{-i})^2}{\sigma(y_i + s_{-i})^2} \leq -((\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla V(\boldsymbol{x}) + V(\boldsymbol{x}))$ from Lemma 4.2. So, we get

$$\begin{aligned}
&\alpha \leq \frac{-2(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla V(\boldsymbol{x}) - 2V(\boldsymbol{x})}{(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla^2 V(\hat{\boldsymbol{x}})(\boldsymbol{y} - \boldsymbol{x})} \\
&\qquad \implies (\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla V(\boldsymbol{x}) + \frac{\alpha(\boldsymbol{y} - \boldsymbol{x})^\intercal \nabla^2 V(\hat{\boldsymbol{x}})(\boldsymbol{y} - \boldsymbol{x})}{2} \leq -V(\boldsymbol{x}).
\end{aligned}$$

Plugging this into the Taylor's expansion in equation (4.14), we get

$$V(\alpha(\boldsymbol{y}-\boldsymbol{x})+\boldsymbol{x}) \leq (1-\alpha)V(\boldsymbol{x}),$$

as required. □

Proof of Lemma 4.4. Note that for any $\boldsymbol{z} \in \mathbb{R}^n$ and $n \geq 2$, we have $\sum_i z_i^2 \leq \sum_i (\sum_{j \neq i} z_j)$ because

$$\sum_i (\sum_{j \neq i} z_j)^2 = \sum_i ((\sum_j z_j) - z_i)^2 = n(\sum_j z_j)^2 + \sum_i z_i^2 - 2(\sum_j z_j)(\sum_i z_i)$$
$$= (n-2)(\sum_j z_j)^2 + \sum_i z_i^2.$$

Using this, we have $\sum_i (y_i - x_i)^2 \leq \sum_i (\sigma - y_i - s_{-i})^2$.

As $x_i, y_i \in [x_{\min}, 1]$ for all i, we have $\frac{\sigma(y_i+s_{-i})^2}{y_i} \leq \frac{n^3}{x_{\min}}$ and $\frac{1}{s_{-i}^2} \leq \frac{1}{(n-1)^2 x_{\min}^2}$. So, $H(\boldsymbol{x}) \leq \frac{B_2 n^3}{2x_{\min}} + \frac{n^3}{(n-1)^2 x_{min}^3} = O(\frac{(1+B_2)n}{x_{\min}^3})$ assuming $x_{\min} \leq \frac{1}{n}$.

For any step-size of $\alpha > 0$ satisfying $\alpha \leq \frac{1}{H(\boldsymbol{x})}$ for all $\boldsymbol{x}$, from Lemma 4.3, we know that $V(t+\alpha) \leq (1-\alpha)V(t)$. After k steps, we have

$$V(k\alpha) \leq (1-\alpha)^k V(0) \leq e^{-k\alpha} V(0).$$

If we take $k = \frac{1}{\alpha}\log(\frac{V(0)}{\epsilon})$, then $V(k\alpha) \leq \epsilon$, which implies that we have reached an ϵ-approximate equilibrium, as required. □

Proof of Lemma 4.5. We construct a simple example with n homogeneous agents where the BR dynamics cycles. Each agent has a cost function $c_i(x_i) = \frac{n-1}{n^2} x_i$. It is not hard to derive that the unique equilibrium is at $(1, 1, \ldots, 1)$ (see, e.g., Vojnović [106]). We will construct a simple two-step cycle where the agents play $(x, x, \ldots, x)$ and $(y, y, \ldots, y)$ repeatedly, where $x < 1$ and $y > 1$. This implies that if the agents enter this cycle, then they will never reach an ϵ-approximate equilibrium for small enough ϵ.

Let us try to construct this cycle, i.e., we try to compute the values of x and y given Δt, which will then tell us the required bound on Δt to have such a cycle. From the first-order conditions, for a linear cost function, we have a closed-form formula for the BR as

$$\frac{\partial u_i(z, s_{-i})}{\partial z} = 0 \implies \frac{s_{-i}}{(z+s_{-i})^2} = \frac{n-1}{n^2} \implies BR_i(s_{-i}) = n\sqrt{\frac{s_{-i}}{n-1}} - s_{-i}.$$

As x is the next action of each agent if everyone is playing y, so

$$x = y + \Delta t \cdot (BR((n-1)y) - y) = y + \Delta t \cdot (n\sqrt{y} - (n-1)y - y) \\ = y + n\Delta t(\sqrt{y} - y).$$

Let $\beta = n\Delta t$, we get

$$x = \beta\sqrt{y} + (1-\beta)y.$$

By symmetry, as y is the next action for an agent if everyone is playing x, we get

$$y = \beta\sqrt{x} + (1-\beta)x \iff (y - (1-\beta)x)^2 = \beta^2 x.$$

Combining the two equations in x and y, in particular, replacing x to get an equation in a single variable y, we get

$$\begin{aligned} & (y - (1-\beta)(\beta\sqrt{y} + (1-\beta)y)))^2 = \beta^2(\beta\sqrt{y} + (1-\beta)y) \\ \iff & (\beta(2-\beta)y - \beta(1-\beta)\sqrt{y})^2 = \beta^2(\beta\sqrt{y} + (1-\beta)y) \\ \iff & (2-\beta)^2y^2 + (1-\beta)^2y - 2(2-\beta)(1-\beta)y\sqrt{y} = \beta\sqrt{y} + (1-\beta)y \\ \iff & \sqrt{y}((2-\beta)^2y\sqrt{y} - 2(2-\beta)(1-\beta)y - \beta(1-\beta)\sqrt{y} - \beta) = 0. \end{aligned}$$

From the above equation, we get our first root as $\sqrt{y} = 0 \implies y = 0$. This is the degenerate solution that corresponds to everyone playing 0, and therefore, the BR also being 0 (see Section 4.2 for a discussion on BR to 0). Let us focus on the other roots. For simplicity, let $z = \sqrt{y}$. We have

$$(2-\beta)^2z^3 - 2(2-\beta)(1-\beta)z^2 - \beta(1-\beta)z - \beta = 0.$$

Notice that $z = 1$ is a root because

$$\begin{aligned} & (2-\beta)^2 - 2(2-\beta)(1-\beta) - \beta(1-\beta) - \beta \\ & \quad = (2-\beta)((2-\beta) - 2(1-\beta)) - \beta + \beta^2 - \beta = (2-\beta)\beta - 2\beta + \beta^2 = 0. \end{aligned}$$

This root $\sqrt{y} = z = 1 \implies y = 1$ corresponds to the equilibrium $(1, 1, \ldots, 1)$. Let us now derive the two other roots, which are of our primary interest as they give us the cycle. First, let us factor out the root $z = 1$ to convert the cubic equation to a quadratic equation. Let $\widehat{z} = z - 1 \iff z = \widehat{z} + 1$. Replacing z by $\widehat{z} + 1$, we get

$$\begin{aligned} & (2-\beta)^2(\widehat{z}+1)^3 - 2(2-\beta)(1-\beta)(\widehat{z}+1)^2 - \beta(1-\beta)(\widehat{z}+1) - \beta = 0 \\ \iff & \widehat{z}^3(2-\beta)^2 + \widehat{z}^2(3(2-\beta)^2 - 2(2-\beta)(1-\beta)) \\ & \quad + \widehat{z}(3(2-\beta)^2 - 4(2-\beta)(1-\beta) - \beta(1-\beta)) \\ & \quad + (2-\beta)^2 - 2(2-\beta)(1-\beta) - \beta(1-\beta) - \beta = 0 \\ \iff & \widehat{z}(\widehat{z}^2(2-\beta)^2 + (2-\beta)(4-\beta)\widehat{z} + (4-\beta)) = 0. \end{aligned}$$

Solving the above equation, in addition to the root $\widehat{z} = 0$ that we have already considered, we get the following two roots

$$\widehat{z} = \frac{\beta - 4 \pm \sqrt{\beta(\beta-4)}}{2(2-\beta)} \iff z = \widehat{z} + 1 = \frac{\beta \pm \sqrt{\beta(\beta-4)}}{2(\beta-2)}$$
$$\iff y = z^2 = \frac{\beta(\beta-2) \pm \beta\sqrt{\beta(\beta-4)}}{2(\beta-2)^2}.$$

In particular, $x = \frac{\beta(\beta-2)-\beta\sqrt{\beta(\beta-4)}}{2(\beta-2)^2} < 1$ and $y = \frac{\beta(\beta-2)-\beta\sqrt{\beta(\beta-4)}}{2(\beta-2)^2} > 1$ are the two roots, assuming $\beta \geq 4$. In particular, if we take $\beta = 6$, we get $x = \frac{3(2-\sqrt{3})}{8} = 0.1005$ and $y = \frac{3(2+\sqrt{3})}{8} = 1.3995$. As $\beta = \Delta t \cdot n = 6$ implies that $\Delta t = \frac{6}{n}$, we get a two step cycle for $\Delta t = \frac{6}{n}$. We can create such cycles for all $\Delta t > \frac{4}{n}$, so $\Delta t \leq \frac{4}{n} = O(\frac{1}{n})$ is necessary for convergence. □

Proof of Lemma 4.6. We consider examples with two agents. Let the first agent have a cost function $c_1(z) = z$ and the second agent have a cost function $c_2(z) = z/d$ for all $z \geq 0$ and some $d \geq 1$. Notice that agent 2 has a lower cost for the same output than agent 1, and $c_1'(z) = 1$ and $c_2'(z) = 1/d$. So, $\frac{\min_{i,z} c_i'(z)}{\max_{i,z} c_i'(z)} = \frac{1}{d}$. Let $\alpha = 1/\Delta t$. We next show non-convergence for $\Delta t = \Omega(1/d) \iff \alpha = O(d)$, as required.

As the cost functions of the agents is linear, using the first-order conditions given in equation 2.2, the BR of agent 1 is given by the explicit formula $BR_1(s_{-1}) = BR_1(x_2) = \sqrt{x_2} - x_2$. Similarly, the BR of agent 2 is $BR_2(s_{-2}) = BR_2(x_1) = \sqrt{dx_1} - x_1$. The unique equilibrium can also be explicitly computed and is equal to $\left(\frac{d}{(1+d)^2}, \frac{d^2}{(1+d)^2}\right)$.

The non-linearity of the BR dynamics makes it tedious to analytically compute the cycles for non-homogeneous agents, especially when Δt is strictly less than 1 . For example, if $d = 16$ and $\Delta t = 0.5$, starting from the state $(1/10, 1/10)$, the dynamics ends up in the cycle in Table 4.1 where $\boldsymbol{x}(t) = \boldsymbol{x}(t+6)$. So, we compute the cycles using numerical simulations.

We simulate the BR dynamics starting from the profile $(1/10, 1/10)$. Given a value of d, we find the critical value of α, denoted by $\alpha^*(d)$, such that the dynamics converges if $\Delta t < \frac{1}{\alpha^*(d)}$ but goes into a cycle if $\Delta t \geq \frac{1}{\alpha^*(d)}$. We do this by a simple binary search over the values of $\Delta t = 1/\alpha$. We notice that $\alpha^*(d)$ is an exact linear function of d, as required; plotted in Figure 4.1. The code is included in the supplementary material. □

Proof of Lemma 4.7. From Lemma 4.3 and 4.4, we know that at time t, if the step-size $\eta_t \leq \alpha$, where $\alpha = \Theta\left(\frac{x_{\min}^3}{n(1+B_2)}\right)$, then the potential V decreases by a factor of

	$x_1(t)$	$x_2(t)$
t	0.021697	0.923661
$t+1$	0.029555	0.745583
$t+2$	0.073722	0.701843
$t+3$	0.104820	0.857095
$t+4$	0.086759	1.023655
$t+5$	0.043385	1.057547
$t+6$	0.021697	0.923661
$t+7$	0.029555	0.745583
...	...	...

Table 4.1: Cycle two non-homogeneous agents ($d = 10$).

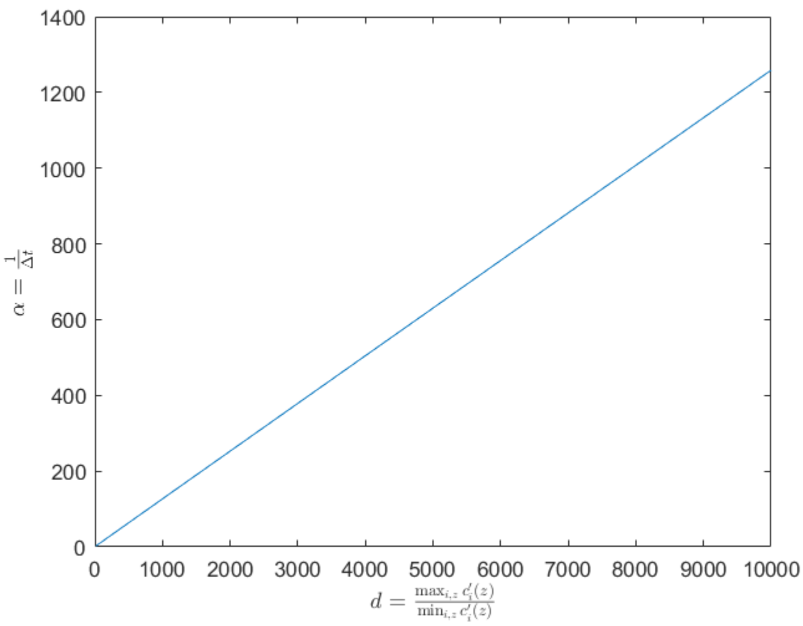

Figure 4.1: Dependency of step-size on cost ratio.

$(1 - \eta_t)$. As $\eta_t \to 0$, there exists a τ such that $\eta_t \leq \alpha$ and $V(t+1) \leq (1 - \eta_t)V(t)$ for all $t \geq \tau$. So, we get for $t \geq \tau$,

$$V(t+1) \leq (\prod_{k=\tau}^{t}(1 - \eta_k))V(\tau) \leq (\prod_{k=\tau}^{t} e^{-\eta_k})V(\tau) = e^{-\sum_{k=\tau}^{t}\eta_k}V(\tau).$$

As $\sum_{k=\tau}^{t} \eta_k \to \infty$, so $e^{-\sum_{k=\tau}^{t}\eta_k} \to 0$ and $V(t) \to 0$. $\square$

Part II

Contest Design

Chapter 5

Threshold Objectives

5.1 Introduction

In many applications, the rules of the contests may be intrinsic, but in others, they may be decided by a designer. We now shift our focus away from learning dynamics in contests to design of contests.

We study contests as incomplete information games. We assume that the players are self-interested and exert costly effort in order to win valuable prizes. Each player is associated with a private *ability* (or *quality*), and their cost, as a function of their output, is linear with a slope equal to the inverse of their ability. The players know the prize allocation scheme, their own ability, and the prior distributions of other players' abilities, and play strategically, reaching a Bayes–Nash equilibrium. On the other hand, the contest designer knows the prior distributions of the players' abilities, and can therefore compute the equilibrium behavior of the players. She wants to design the prize allocation scheme to elicit equilibrium behavior that optimizes her own objective.

The most widely studied designer's objective is the total output, i.e., the sum of the outputs generated by the players. Under the total output objective, the designer values equally the marginal output by a player producing a low or a high level of output. However, in several practical scenarios, the designer may want to focus on the output generated by a section of players producing low/middle/high level of output or to elicit an adequate output from several players instead of very high output from a few players.

Consider, for instance, a crowdsourcing task, such as a survey. It may be more valuable to get many contributors to give adequate responses than to get a few people to submit perfect responses. Similarly, a health insurance company promoting a fitness/exercise app aims to encourage many subscribers to start exercising regularly rather than to get a few fitness enthusiasts to log many hours each day. As another example, consider an instructor who is preparing the students for a standardized test that measures the school performance, but has a limited impact on the students' educational trajectories. The instructor/school may want to incentivize her students to perform better by awarding prizes. If a student is far below the pass/fail threshold, they are likely to fail even if they improve their performance a bit; similarly, there is no need to push students who are sure to excel to work even harder. The crucial students are the ones in between, and the instructor wants to award the prizes to elicit additional output from these students.

We focus on two types of objective functions. Under the *binary threshold objective*, there is a fixed threshold B, and the designer's objective value weakly increases with the number of *successful* players. A player is successful if she produces an output of at least B. In contrast, under the *linear threshold objective*, there is a lower threshold B_L and an upper threshold B_H. All players who produce an output below B_L make the same contribution to the designer's objective; similarly, all players who produce an output above B_H contribute the same amount. However, between these thresholds, a player's output contributes linearly to the designer's objective, just like in case of the total output objective.[1]

Depending upon the situation, a contest designer might be restricted to use only the relative value of the players' outputs to award prizes, which motivates the study of contests with a rank-order allocation of prizes (see, e.g., [81]). On the other hand, the contest designer might be allowed to use the numerical values of the players' outputs and design the optimal *general* contest. In the latter case, the contest design problem is similar to that of an all-pay auction design (see, e.g., [29]).

We assume that the prizes are non-negative, and we normalize them in two ways: *unit-sum* and *unit-range*. The unit-sum constraint is a budget constraint, which requires that the total prize money does not exceed 1. The unit-range constraint restricts the individual prizes awarded to the players to be between 0 and 1. Such a constraint is suitable when the designer is not restricted by a budget, or when the monetary value of the prizes is much less important to her compared to the players' outputs, but there are limits on the prize range. The unit-range constraint for the general contest allows us to optimize for an individual player independently of other players.

[1] In Section 5.6, for rank-order allocation of prizes, we also study a model where the designer's objective is the sum of a concave (or convex) transformation of players' outputs. A concave (convex) transformation captures decreasing (increasing) marginal utility for the designer as a player's output increases.

5.1.1 Overview of Results

At a high level, our results show that the optimal contests for the threshold objectives are strict but minor generalizations of the optimal contest for the well-studied total output objective. Also, these optimal contests for the threshold objectives are easy to interpret, compute, and implement.

Contests with a Rank-Order Allocation of Prizes For the binary threshold objective, the optimal contest equally distributes the prize among the top k players, where k depends upon the distribution F of the players' abilities but not upon the non-decreasing function, say $\rho : \{0\} \cup [n] \to \mathbb{R}_{\geq 0}$, that maps the number of successful players to the utility of the designer (Theorem 5.5). The independence on ρ follows from Lemma 5.3 that shows that the optimal contest for any non-decreasing ρ is the same.

For the linear threshold objective, the optimal contest has up to three levels of prizes: the top k players get the first-level prize, the next ℓ players get the second-level prize, the next m players get the third level prize, and the remaining $n - k - \ell - m$ players do not get anything; here, n is the number of players, and the values of k, ℓ, and m depend upon F (Theorem 5.7). We also prove that a simple contest that gives an equal prize to the first few players and nothing to others (recall that this format is optimal for the binary threshold) has an approximation ratio of 2 for the linear threshold (Theorem 5.8). Both for the binary and for the linear threshold, the results apply to both unit-sum and unit-range constraints on the prizes, although the numbers of players at the different prize levels depend upon the type of constraint.

General Contests For the binary threshold objective, the optimal contest equally distributes the prize to the players who produce an output above a reserve output level, and this reserve is equal to the threshold (Theorem 5.9). This is true for both unit-sum and unit-range constraints, and is independent of the function that maps the number of successful players to the utility of the designer (Lemma 5.3).

For the linear threshold objective, the optimal contest is an extension of the revenue-maximizing all-pay auction with a reserve. For the unit-range constraint, if the distribution of the players' abilities, F, is *regular*[2], then there is a reserve output between the lower and the upper threshold, and any player with an output above the reserve gets a prize of 1 (Theorem 5.14). For the unit-sum constraint and regular

[2]See Definition 5.3. This is a weaker assumption than the *monotone hazard rate* condition.

F, the contest has a reserve output and a *saturation* output. In this case, the prize allocation depends on whether the player with the highest output is: (i) below the reserve, (ii) between the reserve and the saturation level, or (iii) above the saturation level. In case (i), no one gets a prize; in case (ii), the player with the highest output gets the entire prize; and in case (iii), the prize is distributed equally among the players with outputs above the saturation level (Theorem 5.14). The reserve and the saturation levels depend upon F. For *irregular* F, following techniques from optimal auction design, we *iron* the *virtual ability function* to get an optimal contest that is a generalization of the optimal contest for the regular case (Theorem 5.16).

All proofs are provided at the end of the chapter in Section 5.6 to improve readability.

5.1.2 Related Work

Rank-order allocation of prizes is one of the widely studied models in contest theory. Our work is closely related to prior work on contest design with incomplete information and unit-sum constraints [54, 81, 29]. Glazer and Hassin [54] show that for linear cost functions and players' abilities sampled i.i.d. from a uniform distribution, the contest that maximizes the total output awards the entire prize to the top-ranked player. Moldovanu and Sela [81] give the symmetric Bayes–Nash equilibrium (Theorem 5.1) that we use in our analysis. They also generalize the result of Glazer and Hassin [54] and show that awarding the entire prize to the top-ranked player is optimal when the players have (weakly) concave cost functions with the abilities sampled i.i.d. from any distribution with continuous density function; however, with convex cost functions, the optimal mechanism can have multiple prizes.[3] Chawla et al. [29] optimize maximum individual output instead of total output, in a similar incomplete information setup with linear cost functions. Here too the optimal contest allocates the entire prize to the top-ranked player.

Our study of the optimal general contest design rests on the framework established in the seminal work by Myerson [85] on revenue optimal auction design. DiPalantino

[3]Specific cases of our problem (such as the linear threshold objective with only an upper threshold) are related, although not equivalent, to the problem of maximizing total output when players have non-linear cost functions. A non-linear cost function affects the players' equilibrium behavior, but a threshold objective is associated with the designer and does not directly affect the equilibrium behavior. For example, if every player has a convex cost function that is linear up to a certain threshold and then goes to infinity (this setting is similar in spirit to the linear threshold objective with an upper threshold), then no player would produce an output above the threshold, but we will see that the optimal rank-order allocation contest with a linear threshold objective may have players that produce an output above the upper threshold.

and Vojnovic [39] and Chawla et al. [29] connect crowdsourcing contests with all-pay auctions. For general contests with linear cost functions, the optimal contest for total output has been studied by Vojnović [106] and the optimal contest for maximum individual output has been considered by Chawla et al. [29]. The optimal contests for both these objectives have a structure similar to Myerson's optimal auction: they allocate the total budget to the player with the highest output above a reserve output level for regular distributions (and highest ironed output for irregular distributions).

Mechanisms that are similar to the binary threshold objective have been discussed in previous works [101, 56, 74]. For example, Taylor [101] mentions "in a research tournament, the terminal date is fixed, and the quality of innovations varies, while in an innovation race, the quality standard is fixed, and the date of discovery is variable." In the innovation race [74, 101], the objective is to get at least one very good outcome, and a winning criteria that corresponds to a binary threshold (with ties broken in favor of the player who first reaches this threshold) is a proposed mechanism. In our work, the binary threshold is the objective (and not the mechanism), and our aim is to find the optimal mechanism to maximize this objective. Also, even if we interpret the innovation race's mechanism as an objective, it would be a specific instance of the binary threshold objective where the designer's utility is 1 if there is at least one successful player and 0 otherwise (rather than an arbitrary increasing function of the number of successful players).

To the best of our knowledge, there has not been any work on maximizing the linear threshold objectives. In addition to maximizing the total output (e.g., [54, 81, 82, 79]) and the maximum individual output (e.g., [29, 81, 101, 2, 77]), other objectives that have been investigated include maximizing the cumulative output from the top k agents (e.g., [7, 53]).

On the technical side, our work on rank-order allocation builds upon the equilibrium characterization of Moldovanu and Sela [81]. As in the prior work, the *single-crossing* property (Definition 5.2) and properties of order statistics are useful for the characterization of the optimal rank-order allocation contest. For general contests, we build upon the work on revenue optimal auction design by Myerson [85] and on the implementability of auctions by Matthews [75]. Matthews [75] characterizes which interim allocation functions can be implemented by some allocation function, and therefore allows us to focus on interim allocation functions instead of allocation functions. Previous works in general contest design, such as the work of Chawla et al. [29], did not require the result of Matthews [75] because, unlike in our model, their objective functions were linear in the interim allocation.

There have also been several studies in the complete information settings (e.g., [12, 11]). We point the readers to the book by Vojnović [106], particularly Chapter 3, for a survey on related topics.

5.2 Notation and Preliminaries

There are n players. Let $\boldsymbol{v} = (v_1, v_2, \ldots, v_n)$ be the ability profile of the players, where the values v_i are drawn independently from a continuous and differentiable distribution F with support $[0, 1]$. Let f be the probability density function (PDF) of F. The n players simultaneously produce outputs $\boldsymbol{b} = (b_1, b_2, \ldots, b_n) \in \mathbb{R}_+^n$. Player i has a cost of b_i/v_i for producing output b_i; we shall use a scaled version of this cost for convenience, as given in (5.1). For a function $g(x)$ that is not one-to-one, let $g^{-1}(y)$ denote the minimum value of x such that $g(x) \geq y$, for y in the range of g.

5.2.1 Contests with a Rank-Order Allocation of Prizes

The contest has n prizes $\boldsymbol{w} = (w_1, w_2, \ldots, w_n)$, where $0 \leq w_{j+1} \leq w_j \leq 1$ for $j \in [n-1]$. For the unit-sum model, we additionally require that $\sum_j w_j \leq 1$. The prize w_1 is awarded to the highest-performing player, w_2 to the second highest, and so on; a player receives one of the prizes based on the rank of their outputs, with ties broken uniformly. Fix an output vector $\boldsymbol{b} = (b_i)_{i \in [n]}$ and suppose that $b_{i_1} \geq b_{i_2} \geq \ldots \geq b_{i_n}$. Then the utility of player i (scaled up by v_i for convenience) is given by:

$$u(v_i, \boldsymbol{b}) = v_i \sum_{j \in [n]} w_j \frac{\mathbb{1}\{b_i = b_{i_j}\}}{|\{k \mid b_k = b_{i_j}\}|} - b_i, \tag{5.1}$$

where $\mathbb{1}$ is the indicator function. To interpret this formula, observe that $\frac{\mathbb{1}\{b_i = b_{i_j}\}}{|\{k \mid b_k = b_{i_j}\}|}$ is the probability that player i receives the j-th prize, whereas $\frac{b_i}{v_i}$ is the cost of producing output b_i for player i.

Let $p_j(v)$ be the probability that a value $v \in [0, 1]$ is the jth highest among n i.i.d. samples from F, given by the expression:

$$p_j(v) = \binom{n-1}{j-1} F(v)^{n-j} (1 - F(v))^{j-1}. \tag{5.2}$$

Let $f_{n,j}$ be the PDF of the j-th highest order statistic out of n i.i.d. samples from F, given by the expression:

$$f_{n,j}(v) = \frac{n!}{(j-1)!(n-j)!} F(v)^{n-j} (1 - F(v))^{j-1} f(v). \tag{5.3}$$

A key role in the symmetric Bayes–Nash equilibrium of the contest is played by the order statistics of the abilities of the players.

Moldovanu and Sela [81] characterize the unique symmetric Bayes–Nash equilibrium in rank-order allocation contests. Chawla and Hartline [28] prove the uniqueness of this equilibrium in general (see also [29] for more details).

Theorem 5.1. *[81, 28] Consider the game that models the rank-order allocation contest with the values of placement prizes $w_1 \geq w_2 \geq \ldots \geq w_n \geq 0$. The unique Bayes–Nash equilibrium is given by*

$$\beta(v) = \sum_{j \in [n]} w_j \int_0^v t p_j'(t) dt, \tag{5.4}$$

where $\beta(v)$ is the output generated by a player with ability v.

Note that β depends on the prize vector $\boldsymbol{w}$, but we are suppressing it to keep the notation cleaner.

Definition 5.1 (Simple Contest)**.** A rank-order allocation contest is called *simple* if there exists a $j \in [n]$ such that it gives a positive prize of equal value to the first j players and 0 to the other $n - j$ players.

Definition 5.2 (Single-Crossing)**.** A function $f : [a, b] \to \mathbb{R}$ is *single-crossing with respect to a function* $g : [a, b] \to \mathbb{R}$ if there exists a point $x^* \in [a, b]$ such that $f(x) \leq g(x)$ for all $x \leq x^*$ and $f(x) > g(x)$ for all $x > x^*$; when this is the case, we will also say that f is single-crossing with respect to g at x^*.

5.2.2 General Contests

We utilize the connection between contests and all-pay auctions made by [39, 29]. Leveraging the *revelation principle* [85], for most of the analysis of general contests, we shall restrict our attention to direct revelation mechanisms and optimize over allocation rules $\boldsymbol{x}(\boldsymbol{v}) = (x_i(\boldsymbol{v}))_{i \in [n]}$ that determine the allocation based on the abilities of the players $\boldsymbol{v} = (v_i)_{i \in [n]}$. We interpret $x_i(\boldsymbol{v})$ as the expected value of the prize obtained by player i given the ability profile $\boldsymbol{v}$, where the expectation is taken over the choices of the tie-breaking mechanism. We recognize that, when running a contest, we do not have direct access to players' abilities; rather, the players must produce outputs, and the allocation function must be based upon the observed outputs. After deriving the optimal allocation function based on the players' abilities, we shall convert it to an allocation function based on the players' outputs.

For both unit-range and unit-sum settings, we have the restriction $0 \leq x_i(\boldsymbol{v}) \leq 1$ for all $i \in [n]$. For unit-sum, we additionally have $\sum_i x_i(\boldsymbol{v}) \leq 1$. The unit-range case is comparatively easier, because $x_i(\boldsymbol{v})$ can be optimized independently for every player i, whereas for unit-sum, for two players j and ℓ, $x_j(\boldsymbol{v})$ and $x_\ell(\boldsymbol{v})$ are not independent as we need to satisfy the $\sum_i x_i(\boldsymbol{v}) \leq 1$ constraint. While studying unit-sum contests, we shall by default assume that $n \geq 2$.

We assume that the allocation rule $\boldsymbol{x}(\boldsymbol{v})$ is symmetric with respect to the players. As $\boldsymbol{x}(\boldsymbol{v})$ is symmetric, we have $\mathbb{E}[x_i(\boldsymbol{v}) \mid v_i = v] = \mathbb{E}[x_j(\boldsymbol{v}) \mid v_j = v]$ for any $i, j \in [n]$. Let the interim (or expected) allocation function be $\xi(v) = \mathbb{E}[x_i(\boldsymbol{v}) \mid v_i = v]$. Using [85]'s characterization of allocation rules that allow a Bayes–Nash equilibrium, we conclude that $\xi(v)$ should be non-negative and non-decreasing in v, and in the equilibrium, the output of a player as a function of her ability is given by:

$$\beta(v) = v\xi(v) - \int_0^v \xi(t)dt. \tag{5.5}$$

We slightly abuse the notation by representing the output function as β for both rank-order allocation contests and general contests.

Let us make a few crucial observations about $\boldsymbol{x}(\boldsymbol{v})$ and ξ. Both our objective functions, binary threshold (5.4) and linear threshold (5.5), depend upon the output function $\beta(v)$, which further depends upon the interim allocation function $\xi(v)$, but not directly on the allocation function $\boldsymbol{x}(\boldsymbol{v})$. So, any allocation function $\boldsymbol{x}(\boldsymbol{v})$ that leads to the same interim allocation function $\xi(v)$ leads to the same objective value.

Observe that for any interim allocation function ξ that is induced by an allocation function with unit-sum constraints, the following condition holds (check [75] for more details):

$$\int_V^1 \xi(v)f(v)dv \leq \frac{1 - F(V)^n}{n}. \tag{5.6}$$

In plain words, inequality (5.6) says that the probability that any player with an ability above V gets a prize is at most the probability that any player has an ability above V. Now, we state a result by Matthews [75] that we shall use in our analysis.

Theorem 5.2. *[75] Any non-decreasing interim allocation function ξ that satisfies inequality (5.6) is implementable by some allocation function $\boldsymbol{x}$ that satisfies unit-sum constraints.*

Given this result, we can focus on finding a ξ that is non-decreasing and satisfies inequality (5.6) without worrying about the unit-sum constraint $\sum_i x_i(\boldsymbol{v}) \leq 1$, because there will be some $\boldsymbol{x}$ that implements ξ. Further, for the optimal ξ, as we shall see later, the optimal $\boldsymbol{x}$ will be easy to derive.

In our study of optimal general contests, we shall give special attention to regular distributions, defined below. This property of the distribution of the players' abilities allows for simpler and efficiently computable optimal contests (see [85] for their use in optimal auction design).

Definition 5.3 (Regular Distributions)**.** A distribution F is *regular* if if the virtual ability $\psi(v) = v - \frac{1-F(v)}{f(v)}$ is non-decreasing in v.[4]

5.2.3 Objective Functions

We formally define the binary threshold and linear threshold objective functions. They apply to both rank-order allocation contests and general contests. We use the same notation for the output function for both cases, β.

Definition 5.4 (Binary Threshold Objective)**.** Under the binary threshold objective, the contest designer gets a utility of $\rho(k) \in \mathbb{R}_{\geq 0}$ if there are $k \in \{0\} \cup [n]$ players with output equal to or above a specified threshold B. The expected utility of the designer is given by:

$$BT = \mathbb{E}_{\boldsymbol{v}}[\rho(\sum_{i \in [n]} \mathbb{1}\{\beta(v_i) \geq B\})]. \tag{5.7}$$

We assume that ρ is non-decreasing, i.e., $\rho(k-1) \leq \rho(k)$ for all $k \in [n]$, and normalize $\rho(0) = 0$.

The next lemma, Lemma 5.3, proves that the optimal contest for the binary threshold objective (for both rank-order allocation and general contests) is the same for any non-decreasing ρ. Further, given Lemma 5.3, we can maximize $\mathbb{E}_v[\mathbb{1}\{\beta(v) \geq B\}]$ to get the optimal contest for the binary threshold objective.

Lemma 5.3. *The contest that maximizes $\mathbb{E}_v[\mathbb{1}\{\beta(v) \geq B\}]$ maximizes the binary threshold objective $\mathbb{E}_{\boldsymbol{v}}[\rho(\sum_{i \in [n]} \mathbb{1}\{\beta(v_i) \geq B\})]$ for any non-decreasing ρ.*

Definition 5.5 (Linear Threshold Objective)**.** Under the linear threshold objective, the contest designer's utility increases linearly with a player's output if the player's

[4]In the auction theory literature, e.g., [85], v corresponds to the valuation and $\psi(v)$ to the virtual valuation, and, again, F is regular if $\psi(v)$ is non-decreasing. Further, in auctions, $\psi(v)$ may be interpreted as the slope of the "revenue curve" at v, and more coarsely, the v term in $\psi(v)$ as the maximum revenue obtainable and $\frac{1-F(v)}{f(v)}$ as the revenue loss due to not knowing v in advance (a.k.a. "information rent"). Similarly, in contests, $\psi(v)$ may be interpreted as the slope of the "output curve" at ability v.

output is between a lower threshold of B_L and an upper threshold of B_H. Formally, we define it as:

$$
\begin{aligned}
LT &= \mathbb{E}_{\boldsymbol{v}}[\sum_{i\in[n]} \max(0, \min(B_H, \beta(v_i)) - B_L)] \\
&= n\,\mathbb{E}_v[\max(0, \min(B_H, \beta(v)) - B_L)] = n\,\mathbb{E}_v[\max(B_L, \min(B_H, \beta(v)))] - nB_L \\
&\equiv \mathbb{E}_v[\max(B_L, \min(B_H, \beta(v)))] = \int_0^1 \max(B_L, \min(B_H, \beta(v)))f(v)dv, \quad (5.8)
\end{aligned}
$$

where equivalence above denotes the fact that maximizing the left hand side is equivalent to maximizing the right hand side.

Note that the thresholds—B for the binary threshold objective and B_L and B_H for the linear threshold objective—are exogenous.

5.3 Rank-Order Allocation of Prizes

In this section, we study contests that allocate prizes based on the players' ranks. We first present some useful properties of the equilibrium output function β given in Theorem 5.1. Then we study the two objective functions based on these properties.

From Theorem 5.1, we have

$$\beta(v) = \sum_{j\in[n]} w_j \int_0^v t p_j'(t)dt.$$

Writing $p_j'(t)$ using order statistics:

$$p_j'(t) = \begin{cases} f_{n-1,1}(t), & \text{if } j = 1 \\ f_{n-1,j}(t) - f_{n-1,j-1}(t), & \text{if } 1 < j < n \\ -f_{n-1,n-1}(t), & \text{if } j = n \end{cases}.$$

Substituting $p_j'(t)$ into the formula for $\beta(v)$, we get

$$\beta(v) = \sum_{j\in[n-1]} (w_j - w_{j+1}) \int_0^v t f_{n-1,j}(t)dt. \quad (5.9)$$

From equation (5.9) we can observe that decreasing w_n to 0 does not decrease $\beta(v)$ for any v. Changing $\beta(v)$ so that it becomes greater or higher for every v leads to an equal or higher utility for the designer for both binary and linear threshold objectives. So, from now on, we shall assume that $w_n = 0$.

Depending upon whether we are looking at the unit-range or the unit-sum constraint on prizes, we have different constraints on $\boldsymbol{w}$. We now transform $\beta(v)$ further to make it more convenient to work with.

Unit-Sum We have the constraints $\sum_j w_j \leq 1$ and $w_j \geq w_{j+1} \geq 0$. Let $\alpha_j = j(w_j - w_{j+1})$. The set of constraints on $\boldsymbol{\alpha} = (\alpha_j)_{j\in[n-1]}$ that are equivalent to the constraints on $\boldsymbol{w}$ are: $\alpha_j \geq 0$ for all $j \in [n-1]$ and $\sum_{j\in[n-1]} \alpha_j \leq 1$. Let β_S denote the output function β with unit-sum constraints. We can rewrite equation (5.9) as

$$\beta_S(v) = \sum_{j\in[n-1]} \alpha_j \frac{1}{j} \int_0^v t f_{n-1,j}(t)dt = \sum_{j\in[n-1]} \alpha_j \beta_{Sj}(v), \tag{5.10}$$

where $\beta_{Sj}(v) := \frac{1}{j}\int_0^v t f_{n-1,j}(t)dt$. Observe that the simple contest that awards a prize of $1/j$ to the first j players and 0 to others has $\alpha_j = 1$ and $\alpha_k = 0$ for $k \neq j$. Moreover, this contest induces an output of $\beta_{Sj}(v)$ from a player with ability v. Thus, any rank-order prize structure can be written as a convex combination of these $(n-1)$ simple contests where the first j players get awarded $1/j$, for $j \in [n-1]$.

Unit-Range We have the constraints $1 \geq w_j \geq w_{j+1} \geq 0$ for all $j \in [n-1]$. In this case, let $\alpha_j = w_j - w_{j+1}$. We have the same set of constraints on $\boldsymbol{\alpha}$ as with unit-sum: $\alpha_j \geq 0$ for all j and $\sum_{j\in[n-1]} \alpha_j \leq 1$. However, we have a slightly different formula for β, which we denote by β_R:

$$\beta_R(v) = \sum_{j\in[n-1]} \alpha_j \int_0^v t f_{n-1,j}(t)dt = \sum_{j\in[n-1]} \alpha_j \beta_{Rj}(v), \tag{5.11}$$

where $\beta_{Rj}(v) := \int_0^v t f_{n-1,j}(t)dt$. Thus, similarly to the unit-sum case, any unit-range contest and the respective $\beta_R(v)$ can be written as a convex combination of $(n-1)$ simple unit-range contests β_{Rj}, $j \in [n-1]$. However, the unit-range contest that induces β_{Rj} awards a prize of 1 to the first j players and 0 to others, whereas the unit-sum contest β_{Sj} awards a prize of $1/j$ to the first j players.

Using the characterization of β in equation (5.10) for unit-sum and (5.11) for unit-range, we can easily prove that the *total output* objective is maximized by β_{S1} for unit-sum contests (proved by [54, 81]) and by a simple contest for unit-range contests (the proof is provided in Section 5.6.2.1).

Most of our analysis in this section applies to both unit-range and unit-sum settings; we shall use β to denote either β_S or β_R, and β_j to denote either β_{Sj} or β_{Rj}. Also, we shall assume without loss of generality that $\sum_j \alpha_j = 1$, because increasing α_i for some $i \in [n-1]$ while keeping α_j constant for all $j \in [n-1] \setminus \{i\}$ does not decrease $\beta(v)$ for any $v \in [0,1]$, and therefore does not decrease either of our two objective functions.

Theorem 5.4. *Fix an $\boldsymbol{\alpha}$ and the corresponding output function β. Consider a pair of indices j, k s.t. $1 \le j < k \le n-1$, and $\epsilon > 0$. Suppose both the vector $\boldsymbol{\alpha}$ and the vector $\boldsymbol{\alpha}'$ given by $\alpha'_j = \alpha_j + \epsilon$, $\alpha'_k = \alpha_k - \epsilon$, $\alpha'_\ell = \alpha_\ell$ for $\ell \notin \{j,k\}$ satisfy the required constraints. Let β' be the output function that corresponds to $\boldsymbol{\alpha}'$. Then, β' is single-crossing w.r.t. β, and β^{-1} is single-crossing w.r.t. β'^{-1}.*

The proof of Theorem 5.4 for unit-sum is available in ([106], Chapter 3); in Section 5.6.2.3, for completeness, we provide a similar proof for both unit-sum and unit-range settings.

5.3.1 Binary Threshold Objective

We first focus on the binary threshold objective (Definition 5.4; Lemma 5.3): $\mathbb{E}_v[\mathbb{1}\{\beta(v) \ge B\}]$.

Theorem 5.5. *The rank-order allocation contest that optimizes the binary threshold objective is simple, and the output function for the optimal contest is β_{j^*}, where j^* is selected from the set*

$$\arg\min_j \beta_j^{-1}(B).$$

Given Theorem 5.5, we can design the optimal contest by first finding the root of the equation $\beta_j(v) - B = 0$ for each $j \in [n-1]$. We can do this efficiently using a root-finding algorithm such as the bisection method because β_j is continuous and monotone. Then, we select a j with the smallest $\beta_j^{-1}(B)$.

5.3.2 Linear Threshold Objective

We now consider the linear threshold objective: $\mathbb{E}[\max(B_L, \min(B_H, \beta(v)))]$ (Definition 5.5). For the binary threshold objective, the optimal contest was simple, but for the linear threshold this is not true in general. The following example illustrates this:

Example 5.6. Consider a contest with: unit-sum prizes; three players, $n = 3$; uniform distribution, $F(v) = v$, $f(v) = 1$; lower threshold $B_L = 0$; upper threshold $B_H = \frac{1}{4} \times \left(\frac{2}{3}\right)^2 + \frac{1}{6} \times \left(\frac{2}{3}\right)^3 \approx 0.1605$. The output function is $\beta(v) = \alpha_1\beta_1(v) + \alpha_2\beta_2(v)$ where $\alpha_1 + \alpha_2 = 1$ and $\beta_1(v) = \int_0^v t f_{2,1}(t)dt = \int_0^v 2t^2 dt = \frac{2}{3}v^3$ and $\beta_2(v) = \frac{1}{2}\int_0^v t f_{2,2}(t)dt = \int_0^v t(1-t)dt = \frac{v^2}{2} - \frac{v^3}{3}$. We consider three contests: the two simple contests and a mixed one.

- Simple Contest 1: $\alpha_1 = 1$ and $\beta(v) = \beta_1(v)$. We have $\beta_1^{-1}(B_H) \approx 0.6221$ and the objective value is $\int_0^{\beta_1^{-1}(B_H)} \beta_1(v)dv + B_H(1 - \beta_1^{-1}(B_H)) \approx 0.0857$.

- Simple Contest 2: $\alpha_2 = 1$ and $\beta(v) = \beta_2(v)$. The objective value is $\int_0^1 \min(B_H, \beta_2(v))dv \le \int_0^1 \beta_2(v)dv = 1/12 \approx 0.0833$.
- Mixed Contest: $\alpha_1 = \alpha_2 = 1/2$ and $\beta(v) = \beta_1(v)/2 + \beta_2(v)/2$. We have $\beta^{-1}(B_H) = 2/3$ and the objective value is $\int_0^{2/3}(\frac{v^2}{4} + \frac{v^3}{6})dv + B_H(1 - \frac{2}{3}) \approx 0.0864$.

We observe that the given mixed contest outperforms the two simple contests.

For the case where there is only an upper threshold, i.e., $B_L = 0$, there is an optimal contest that is a convex combination of only two simple contests.

Theorem 5.6. *For a linear threshold objective with upper threshold only, i.e., with $B_L = 0$, there is an optimal $\boldsymbol{\alpha}$ with at most two positive entries α_i and α_j, i.e., with $\alpha_k = 0$ for $k \in [n-1] \setminus \{i, j\}$. For this $\boldsymbol{\alpha}$, i and j, we also have:*

$$\int_0^{V_H} \beta_i(v)f(v)dv = \int_0^{V_H} \beta_j(v)f(v)dv = \int_0^{V_H} \beta(v)f(v)dv,$$

where $V_H = \beta^{-1}(B_H)$ and β is the output function induced by $\boldsymbol{\alpha}$.

Theorem 5.6 suggests an algorithm for finding the optimal $\boldsymbol{\alpha}$, sketched below:

1. For every $i, j \in [n-1]$, $i < j$, find a $V_{ij} > 0$ (if any) such that $\int_0^{V_{ij}} \beta_i(v)f(v)dv = \int_0^{V_{ij}} \beta_j(v)f(v)dv$. Note that there might be multiple such values for V_{ij}, but these values form an interval of $[0, 1]$ because $\int_0^{v} \beta_i(v)f(v)dv$ is single-crossing w.r.t. $\int_0^{v} \beta_j(v)f(v)dv$. Select any one of those values as V_{ij}.
2. If V_{ij} exists and $\beta_i(V_{ij}) \ge B_H \ge \beta_j(V_{ij})$ then this pair i, j is a candidate for being the optimal.
3. The objective value for this pair is $\int_0^{V_{ij}} \beta_i(v)f(v)dv + B_H(1 - F(V_{ij}))$.
4. Comparing $O(n^2)$ such pairs along with the $O(n)$ simple contests, we find the optimal contest.
5. $\boldsymbol{\alpha}$ corresponding to pair i, j can be calculated as $\alpha_i = \frac{B_H - \beta_j(V_{ij})}{\beta_i(V_{ij}) - \beta_j(V_{ij})}$, $\alpha_j = 1 - \alpha_i$, and $\alpha_k = 0$ for $k \in [n-1] \setminus \{i, j\}$. (Check the proof of Theorem 5.6 for more details about this step.)

We can prove a result analogous to Theorem 5.6 if we only have a lower threshold and no upper threshold (i.e., $B_H = 1$): there is an optimal contest that is a convex combination of at most two simple contests. Now, we give a result that applies for arbitrary thresholds.

Theorem 5.7. *For a linear threshold objective, there is an optimal $\boldsymbol{\alpha}$ with at most three positive entries α_i, α_j, and α_k, i.e., with $\alpha_\ell = 0$ for $\ell \in [n-1] \setminus \{i, j, k\}$. For this $\boldsymbol{\alpha}$ and i, j, k, we also have:*

$$\int_{V_L}^{V_H} \beta_i(v) f(v) dv = \int_{V_L}^{V_H} \beta_j(v) f(v) dv = \int_{V_L}^{V_H} \beta_k(v) f(v) dv = \int_{V_L}^{V_H} \beta(v) f(v) dv,$$

where $V_L = \beta^{-1}(B_L)$, $V_H = \beta^{-1}(B_H)$, and β is the output function induced by $\boldsymbol{\alpha}$.

The main ingredients used in the proof of Theorem 5.7 (and Theorem 5.6) are: first-order optimality condition of the objective w.r.t. $\boldsymbol{\alpha}$; single-crossing property of β_i w.r.t. β_j for $i < j$; and the fact that every linear programming problem has an optimal solution that lies at a corner of the feasible region.

Recall that for the case $B_L = 0$, we used Theorem 5.6 to obtain an algorithm that compares at most $O(n^2)$ contests to find an optimal one. In a similar spirit, for general B_L and B_H, we can use Theorem 5.7 to obtain an algorithm that finds an optimal contest by comparing at most $O(n^3)$ contests.

5.3.2.1 Simple vs Optimal

In Theorem 5.7, we proved that a convex combination of at most three simple contests is optimal. We now compare this optimal contest with the best simple contest.

Theorem 5.8. *For the linear threshold objective, the objective value of the optimal contest is at most* 2 *times that of the best simple contest.*

In the proof of Theorem 5.8, we use the expression given in Theorem 5.7 and the single-crossing property of β_i w.r.t. β_j for any $i < j$ to show that the objective value of the optimal contest is at most the sum of the objective values of two simple contests. So, one of these two simple contests gives us an objective value at least half of the optimal.

5.4 General Optimal Contests

In the previous section, we restricted our focus to contests that awarded prizes based on players' ranks only. In this section, we relax this restriction and consider contests that may use the numerical values of the players' outputs to award prizes.

As we discussed in the preliminaries (Section 5.2), for the unit-range constraint, the allocation function must satisfy $0 \leq x_i(\boldsymbol{v}) \leq 1$. Equivalently, the expected allocation function ξ must satisfy $0 \leq \xi(v) \leq 1$. For the unit-sum constraint, in addition to the constraints for the unit-range case, the allocation function must also satisfy $\sum_i x_i(\boldsymbol{v}) \leq 1$. Equivalently (by Theorem 5.2), the expected allocation function ξ must satisfy inequality (5.6): $\int_V^1 \xi(v)f(v)dv \leq \frac{1-F(V)^n}{n}$. From Myerson's characterization of allocation functions that allow a Bayes–Nash equilibrium, we know that ξ must be monotonically non-decreasing.

The following allocation rules award the entire prize to players with abilities above V, if there are such players in a given ability profile. Therefore, they maximize $\int_V^1 \xi(v)f(v)dv$, and the inequality (5.6) is satisfied with an equality.

1. Give the prize to the player with the highest ability. Then the expected allocation function is $\xi(v) = F(v)^{n-1}$, and $\int_V^1 \xi(v)f(v)dv = \int_V^1 F(v)^{n-1}f(v)dv = \int_{F(V)}^1 y^{n-1}dy = \frac{1-F(V)^n}{n}$. We can also observe that for this allocation rule, inequality (5.6) is tight for every $V \in [0,1]$.

2. Uniformly distribute the prize among the players with abilities above V. Then the expected allocation function is $\xi(v) = \frac{1-F(V)^n}{n(1-F(V))}$, and $\int_V^1 \xi(v)f(v)dv = \int_V^1 \frac{1-F(V)^n}{n(1-F(V))}f(v)dv = \frac{1-F(V)^n}{n}$.

5.4.1 Binary Threshold Objective

For the binary threshold objective (Definition 5.4; Lemma 5.3), we have:

$$\max \mathbb{E}[\mathbb{1}\{\beta(v) \geq B\}] = \max \int_0^1 \mathbb{1}\{\beta(v) \geq B\}f(v)dv \equiv \min(\beta^{-1}(B)).$$

Thus, we want to find a ξ that minimizes $\beta^{-1}(B)$.

Theorem 5.9. *The optimal contest with the binary threshold objective gives a prize of* 1 *to all players who produce an output above the threshold* B *in the unit-range model and equally distributes the total prize of* 1 *to all players who produce an output above the threshold* B *in the unit-sum model.*

5.4.2 Linear Threshold Objective

The linear threshold objective is $\mathbb{E}[\max(B_L, \min(B_H, \beta(v)))]$ (Definition 5.5).

5.4.2.1 Regular Distributions.

Let us first focus on the case when F is *regular*. We shall see that the optimal contest resembles a highest bidder wins all-pay auction with a *reserve bid* (if every player bids below this, no one gets the prize) and a *saturation bid* (all bids above this level are considered equal), and the ties are broken uniformly. We shall also give an efficient way to find the reserve and the saturation bids.

We can, w.l.o.g., make the following assumptions on the optimal expected allocation function:

Lemma 5.10. *If ξ is optimal, we can assume w.l.o.g. that $\xi(v) = 0$ for $v < \beta^{-1}(B_L)$.*

Lemma 5.11. *If ξ is optimal, we can assume w.l.o.g. that $\xi(v) = \xi(V_H)$ for $v \geq \beta^{-1}(B_H)$, i.e., $\xi(v)$ is constant for $v \geq \beta^{-1}(B_H)$.*

From now on, we shall assume that ξ satisfies the assumptions formulated in these two lemmas.

We can write our linear threshold objective, using the notation $V_L = \beta^{-1}(B_L)$ and $V_H = \beta^{-1}(B_H)$, as:

$$B_L F(V_L) + \int_{V_L}^{V_H} \beta(v) f(v) dv + B_H(1 - F(V_H)). \tag{5.12}$$

Intuitively, the following lemma says that we should push the area under the curve ξ to the right, as much as possible. We prove this result for unit-sum; for unit-range it holds, as a corollary, if we set $n = 1$.

Lemma 5.12. *Let F be a regular distribution and suppose that the allocation function has to satisfy unit-sum constraints. Then there is an optimal ξ such that $\int_v^1 \xi(t) f(t) dt = \frac{1}{n}(1 - F(v)^n)$ for all $v \in [0, 1]$ where $\beta(v) < B_H$ and $\xi(v) > 0$.*

The previous lemma effectively says that for $v \in [V_L, V_H)$ we have

$$\int_v^1 \xi(t) f(t) dt = \frac{1 - F(v)^n}{n}.$$

As both sides of the above equation are continuous, taking the limit $v \to V_H$, we have:

$$\int_{V_H}^1 \xi(t) f(t) dt = \frac{1 - F(V_H)^n}{n}.$$

We already know that ξ is 0 for $v < V_L$ and constant for $v \geq V_H$. So, we get

$$\xi(v) = \begin{cases} 0, & \text{if } v < V_L \\ F(v)^{n-1}, & \text{if } V_L \leq v < V_H \\ \frac{1-F(V_H)^n}{n(1-F(V_H))}, & \text{if } v \geq V_H \end{cases} \tag{5.13}$$

We now prove the following lemma, which says that for an optimal allocation rule, the output generated by a player never goes strictly above the upper threshold B_H almost surely in $v \in [0, 1]$.

Lemma 5.13. *If ξ is an optimal expected allocation function, then the induced output function β almost surely satisfies $\beta(v) \leq B_H$ for $v \in [0,1]$. This result holds whether or not the distribution F is regular, and both for unit-range and for unit sum constraints.*

Combining the optimal expected allocation rule given in (5.13) with Lemma 5.13, we get

$$\beta(V_H) = \frac{V_H(1-F(V_H)^n)}{n(1-F(V_H))} - \int_{V_L}^{V_H} F(v)^{n-1}dv = B_H.$$

Note that the above formula is applicable only if there exists a V_H. It may be possible that $\beta(v)$ never reaches the threshold B_H, i.e., $\beta(v) < B_H$ for $v \in [0,1]$. The above formula also tells us that given V_L (or V_H), V_H (or V_L) can be calculated efficiently because, keeping V_L (or V_H) fixed, $\beta(v)$ is continuous and monotone in V_H (or V_L).

We now discuss how to efficiently compute the reserve V_L and the saturation V_H. We shall use the following notation in the next theorem: $\eta(x) = \frac{1-x^n}{n(1-x)}$, $\psi_u(v) = v - \frac{F(u)-F(v)}{f(v)}$, V_{low} is the solution of $V_{\text{low}}F(V_{\text{low}})^{n-1} = B_L$, V_{mid} is the solution of $\int_{V_{\text{mid}}}^1 F(v)^{n-1}dv = 1 - B_H$, and V_{up} is the solution of $B_H = V_{up}\eta(F(V_{\text{up}}))$.

Theorem 5.14. *The contest that optimizes the linear threshold objective for regular distributions has the following allocation and expected allocation functions:*

1. *For unit-range, the expected allocation function ξ and the allocation function $\boldsymbol{x}(\boldsymbol{v})$ are given as:*

$$\xi(v) = \begin{cases} 0, & \text{if } v < V \\ 1, & \text{if } v \geq V \end{cases}, \quad x_i(\boldsymbol{v}) = \begin{cases} 0, & \text{if } v_i < V \\ 1, & \text{if } v_i \geq V \end{cases}$$

where $V = \max(B_L, \min(B_H, \psi^{-1}(B_L)))$.

2. *For unit-sum, the optimal solution is given by one of the two cases below:*

   (a) *The expected allocation function ξ is:*

$$\xi(v) = \begin{cases} 0, & \text{if } v < V_L \\ F(v)^{n-1}, & \text{if } v \geq V_L \end{cases} \tag{5.14}$$

and the allocation function $\boldsymbol{x}(\boldsymbol{v})$ is:

$$x_i(\boldsymbol{v}) = \begin{cases} 0, & \text{if } \max_j(v_j) < V_L \text{ or } i \notin W \\ 1/|W|, & \text{if } i \in W \text{ and } \max_j(v_j) \geq V_L \end{cases} \tag{5.15}$$

where $W = \{k \mid v_k = \max_j(v_j)\}$ and $V_L = \min(V_{\text{mid}}, \max(V_{\text{low}}, \overline{V_L}))$; where $\overline{V_L}$ is the solution of $\frac{B_L}{F(\overline{V_L})^{n-1}} - \psi(\overline{V_L}) = 0$.

(b) *The expected allocation function* ξ *is:*

$$\xi(v) = \begin{cases} 0, & \text{if } v < V_L \\ F(v)^{n-1}, & \text{if } V_L \leq v < V_H \\ \eta(F(V_H)), & \text{if } v \geq V_H \end{cases} \tag{5.16}$$

and the allocation function $\boldsymbol{x}(\boldsymbol{v})$ *is:*

$$x_i(\boldsymbol{v}) = \begin{cases} 0, & \text{if } \max_j(v_j) < V_L \text{ or } i \notin W \\ 1/|W|, & \text{if } i \in W \text{ and } V_L \leq \max_j(v_j) < V_H \\ 1/|\widehat{W}|, & \text{if } i \in \widehat{W} \text{ and } \max_j(v_j) \geq V_H \end{cases} \tag{5.17}$$

where $W = \{k \mid v_k = \max_j(v_j)\}$, $\widehat{W} = \{k \mid v_k \geq V_H\}$, $V_L = \min(V_{\text{up}}, \max(V_{\text{mid}}, \overline{V_L}))$ *and* $V_H = \min(1, \max(V_{\text{up}}, \overline{V_H}))$*; where* $\overline{V_L}$ *and* $\overline{V_H}$ *are the solutions of equations* $\frac{B_L}{F(\overline{V_L})^{n-1}} - \psi_{\overline{V_H}}(\overline{V_L}) = 0$ *and* $\overline{V_H}\eta(F(\overline{V_H})) - \int_{\overline{V_L}}^{\overline{V_H}} F(v)^{n-1} dv = B_H$*.*

The values of V_L and V_H derived in the theorem above can be efficiently computed if the distribution F is known, see the proof for more details. Also note that, although the contest may seem complicated, it is reasonably simple from a player's perspective. The optimal contest that maximizes the total output has a reserve [106], here we have a saturation value in addition to the reserve. A player need not know how these reserve and saturation values are computed.

Implementing this mechanism in practice, i.e., finding the allocation as a function of the outputs of the players, is not difficult. Given a player's output b, map it to $g(b)$, where:

$$g(b) = \begin{cases} 0, & \text{if } b < F(V_L)^{n-1} \\ b, & \text{if } F(V_L)^{n-1} \leq b \leq F(V_H)^{n-1} \\ F(V_H)^{n-1}, & \text{if } F(V_H)^{n-1} < b < \frac{1-F(V_H)^n}{n(1-F(V_H))} \\ \frac{1-F(V_H)^n}{n(1-F(V_H))}. & \text{if } b \geq \frac{1-F(V_H)^n}{n(1-F(V_H))} \end{cases} \tag{5.18}$$

One can then distribute the prize equally among the players who have the maximum positive $g(b)$.

5.4.2.2 Irregular Distributions.

In the study of the optimal linear threshold contest for regular distributions, we used the regularity condition at two places: first, in Lemma 5.12, to pack the area under ξ to the right; second, in Theorem 5.14, to prove that we obtain the optimal values for V_L and V_H by solving for the roots of the specific equations given in the theorem statement using an efficient root-finding method. For irregular F, first, we give Lemma 5.15 analogous to Lemma 5.12; second, we find an approximate solution by discretizing the feasible space of V_L, V_H, and $\xi(V_H)$.[5]

We first introduce some additional notation. Consider the function $\psi_{U,V}(v) = v - \frac{F(V)-F(v)}{f(v)}$, defined on the interval $[U, V]$. We now define $\overline{\psi}_{U,V}$, which is the ironed version of $\psi_{U,V}$, also defined on $[U, V]$. Our definition proceeds in several steps.

1. Let $h_{U,V}(y) = \psi_{U,V}(F^{-1}(y))$ and $H_{U,V}(y) = \int_{F^{-1}(U)}^{y} h_{U,V}(y)dy$;

2. Let $\overline{H}_{U,V}(y)$ be the point-wise maximum convex function less than or equal to $H_{U,V}(y)$. Note that at the boundary $\overline{H}_{U,V}(F^{-1}(U)) = H_{U,V}(F^{-1}(U))$ and $\overline{H}_{U,V}(F^{-1}(V)) = H_{U,V}(F^{-1}(V))$;

3. Let $\overline{h}_{U,V}(y) = \overline{H}'_{U,V}(y)$ and $\overline{\psi}_{U,V}(v) = \overline{h}_{U,V}(F(v))$.

Let $l_{U,V}(v) = \min_{u \in [U,V], \overline{\psi}_{U,V}(u) = \overline{\psi}_{U,V}(v)} u$ and $r_{U,V}(v) = \max_{u \in [U,V], \overline{\psi}_{U,V}(u) = \overline{\psi}_{U,V}(v)} u$, and let $l(u) = l_{0,1}(u)$ and $r(u) = r_{0,1}(u)$.

Lemma 5.15. *For an irregular distribution F, in the unit-sum setting there is an optimal ξ such that $\int_v^1 \xi(t)f(t)dt = \frac{1}{n}(1-F(v)^n)$ for all v where $\beta(v) < B_H$, $\xi(v) > 0$, and $\overline{H}(F^{-1}(v)) = H(F^{-1}(v))$.*

Note that Lemma 5.15 does not apply to points $v \in [0, 1]$ where $\overline{H}(F^{-1}(v)) < H(F^{-1}(v))$. So, unlike the case for regular distributions, where we had $\int_{V_H}^1 \xi(t)f(t)dt = \frac{1}{n}(1 - F(V_H)^n)$ by applying Lemma 5.12 for $v \to V_H$, we may not have a similar result for irregular distributions.

For regular F, the allocation function $\boldsymbol{x}$ has three cases depending upon the highest ability, as given in (5.15). On the other hand, for irregular F, the allocation function $\boldsymbol{x}$ has five cases depending upon the highest ability, as given in the theorem below:

[5] We would like to note here that the algorithm we provide finds an approximate solution in time polynomial in the reciprocal of the parameter used to discretize V_L, V_H, and $\xi(V_H)$. One could have discretized the entire optimization problem (the functions are evaluated at discrete values, the integrals are written as finite summations, etc.) and directly found an approximately optimal discretized ξ and β using linear programming, also in polynomial time. The advantage of our analysis is that it characterizes the optimal solution and gives a more intuitive algorithm.

Theorem 5.16. *The contest that optimizes the linear threshold objective for irregular distributions and unit-sum constraints has the following allocation function:*

$$x_i(\boldsymbol{v}) = \begin{cases} 0, \qquad \textit{if } \max_j(v_j) < V_L \textit{ or } i \notin W \\ \frac{n\xi(V_L)(F(r_{V_L,V_H}(V_L))-F(V_L))}{|W_L|(F(r_{V_L,V_H}(V_L))^n-F(V_L)^n)}, \\ \qquad\qquad \textit{if } V_L \le \max_j(v_j) < r_{V_L,V_H}(V_L) \textit{ and } i \in W_L \\ 1/|W|, \qquad \textit{if } r_{V_L,V_H}(V_L) \le \max_j(v_j) < l_{V_L,V_H}(V_H) \textit{ and } i \in W \\ \frac{n\xi(l_{V_L,V_H}(V_H))(F(V_H)-F(l_{V_L,V_H}(V_H)))}{|W_H|(F(V_H)^n-F(l_{V_L,V_H}(V_H))^n)}, \\ \qquad\qquad \textit{if } l_{V_L,V_H}(V_H) \le \max_j(v_j) < V_H \textit{ and } i \in W_H \\ \frac{n\xi(V_H)(1-F(V_H))}{|\widehat{W}|(1-F(V_H)^n)}, \qquad \textit{if } V_H \le \max_j(v_j) \textit{ and } i \in \widehat{W} \end{cases}, \tag{5.19}$$

where $W = \{k \mid r_{V_L,V_H}(V_L) \le v_k < l_{V_L,V_H}(V_H), \psi_{V_L,V_H}(v_k) = \max_j(\psi_{V_L,V_H}(v_j))\}$, $W_L = \{k \mid v_k \in [V_L, r_{V_L,V_H}(V_L)]\}$, $W_H = \{k \mid v_k \in [l_{V_L,V_H}(V_H), V_H]\}$, *and* $\widehat{W} = \{k \mid v_k \ge V_H\}$.

In Section 5.6.3, we have sketched an approximation algorithm to find the parameters used in the theorem, like V_L, V_H, $\xi(V_L)$, $\xi(V_H)$, etc.

We can transform the allocation function $\boldsymbol{x}(\boldsymbol{v})$ given in the theorem to an allocation function based on outputs in a manner analogous to (5.18). The optimal contest for unit-range can be derived by combining ideas from the solution for unit-range with regular F and the solution for unit-sum with irregular F.

5.5 Conclusion

We looked at two natural and practically useful objectives for a contest designer and described optimal contests for the objectives. An interesting open problem is to find how well can a contest without a reserve and/or a saturation output, or a rank-order allocation contest, approximate the optimal general contest, possibly with an additional player; we may expect to get a result along the lines of [26]. Another extension of this work would be to study other practically relevant objective functions for the designer, monotone transformations other than the threshold transformations we studied in this chapter. Combining the objectives for the designer studied in this chapter with non-linear utility and cost functions for the players is also an important research direction.

5.6 Omitted Proofs and Additional Results

5.6.1 Supplement to Section 5.2

5.6.1.1 Expected Allocation

We use the expression given below in Section 5.4. For completeness, we provide its derivation.

Consider a player i with ability v, and fix a and b so that $0 \le a \le v \le b \le 1$. For each $k \in [n-1]$, let $p(v,a,b,k)$ be the probability that all players other than player i have ability at most b, with k of them having ability in $[a,b]$. We have:

$$p(v,a,b,k) = \binom{n-1}{k} F(a)^{n-1-k}(F(b)-F(a))^k.$$

Now, if the player with ability v is allocated $1/(k+1)$ with probability $p(v,a,b,k)$, then the expected prize allocation for the player is:

$$\begin{aligned}
\sum_{k=0}^{n-1} \frac{p(v,a,b,k)}{k+1} &= \sum_{k=0}^{n-1} \binom{n-1}{k} \frac{F(a)^{n-1-k}(F(b)-F(a))^k}{k+1} \\
&= \sum_{k=0}^{n-1} \binom{n}{k+1} \frac{F(a)^{n-1-k}(F(b)-F(a))^k}{n} \\
&= \sum_{k=1}^{n} \binom{n}{k} \frac{F(a)^{n-k}(F(b)-F(a))^{k-1}}{n} \\
&= \frac{1}{n(F(b)-F(a))} \sum_{k=1}^{n} \binom{n}{k} F(a)^{n-k}(F(b)-F(a))^k \\
&= \frac{1}{n(F(b)-F(a))} \left(\sum_{k=0}^{n} \binom{n}{k} F(a)^{n-k}(F(b)-F(a))^k - F(a)^n \right) \\
&= \frac{1}{n(F(b)-F(a))} \left(F(b)^n - F(a)^n\right) = \frac{1}{n} \cdot \frac{F(b)^n - F(a)^n}{F(b)-F(a)}.
\end{aligned}$$

5.6.1.2 Properties of Single-Crossing Functions

The following are some well-known properties of single-crossing functions that we shall use to prove some of our results. (Proofs are given for the reader's convenience.)

Lemma 5.17. *For two functions $f, g : [a,b] \to \mathbb{R}$, if f is single-crossing with respect to g, then $F(x) = \int_a^x f(x)dx$ is single-crossing with respect to $G(x) = \int_a^x g(x)dx$.*

Proof. Let $y^* \in [a,b]$ be the maximum value such that $F(y) \le G(y)$ for all $y \le y^*$. Such a value exists because the set $\{y : F(y) \le G(y)\}$ is closed and $F(a) = G(a) = 0$.

Suppose that f is single-crossing with respect to g at x^*. Then we have $y^* \geq x^*$, because for every $x \leq x^*$ we have $g(x) - f(x) \geq 0$, and therefore for every $y \leq x^*$ we have $G(y) - F(y) = \int_a^y (g(x) - f(x))dx \geq 0$. Now, for $y > x^*$, we know that $f(y) > g(y)$ so $F(y) - G(y)$ is strictly increasing. So, if $F(x)$ goes above $G(x)$ at some point in $[a, b]$, then it will remain that way. $\square$

Lemma 5.18. *For two continuous and strictly increasing functions $f, g : [a, b] \to \mathbb{R}$ with $f(a) = g(a) = 0$, if f is single-crossing with respect to g at x^*, then g^{-1} is single-crossing with respect to f^{-1} (in the intersection of the domains of f^{-1} and g^{-1}) at $y^* = f(x^*) = g(x^*)$.*

Proof. Consider a point y that belongs to the domains of f^{-1} and g^{-1}. Suppose $y \leq y^*$. Since f and g are continuous and strictly increasing, there exists a unique point x with $y = f(x)$ and a unique point z with $y = g(z)$. Further, $y \leq y^*$ implies $x \leq x^*$ and hence $f(x) \leq g(x)$. It follows that $z \leq x$, i.e., $g^{-1}(y) \leq f^{-1}(y)$. In a similar manner, we can prove that $g^{-1}(y) > f^{-1}(y)$ for each point $y > y^*$ in the intersection of the domains of f^{-1} and g^{-1}. $\square$

5.6.1.3 Omitted Proofs

Proof of Lemma 5.3. Recall that we are in an incomplete information model where the ability of every agent is picked independently from the distribution F, and where an agent i only knows her ability v_i but not the ability v_j of any other agent $j \neq i$. Moreover, as we are looking at symmetric mechanisms, from equations (5.4) and (5.5) we know that $\beta(v_i) = \beta(v_j)$ if $v_i = v_j$ for any agents i and j, i.e., the agents use the same function to map their ability v to the output $\beta(v)$.

Let $q = \mathbb{E}_v[\mathbb{1}\{\beta(v) \geq B\}] = \mathbb{P}_v[\beta(v) \geq B]$ be the probability that any given agent has output above B. The probability that k agents have output above B is equal to $\binom{n}{k} q^k (1-q)^{n-k}$. The binary threshold objective can be written as

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{v}}[\rho(\sum_{i\in[n]} \mathbb{1}\{\beta(v_i) \geq B\})] &= \sum_{k\in[n]} \rho(k) \binom{n}{k} q^k (1-q)^{n-k} \\
&= \sum_{k\in[n]} (\rho(k) - \rho(k-1)) \sum_{\ell=k}^{n} \binom{n}{\ell} q^\ell (1-q)^{n-\ell} = \sum_{k\in[n]} (\rho(k) - \rho(k-1)) g_k(q),
\end{aligned}$$

where, let, $g_k(q) = \sum_{\ell=k}^{n} \binom{n}{\ell} q^\ell (1-q)^{n-\ell}$. As ρ is non-decreasing, so $\rho(k) - \rho(k-1) \geq 0$ for every $k \in [n]$. Now, if we show that $g_k(q)$ is a non-decreasing function of q, then maximizing q will maximize the objective, which shall complete the proof. Let us differentiate $g_k(q)$ w.r.t. q, we get

$$\begin{aligned}
\frac{dg_k(q)}{dq} &= \sum_{\ell=k}^{n} \binom{n}{\ell} \frac{d(q^\ell(1-q)^{n-\ell})}{dq} \\
&= nq^{n-1} + \sum_{\ell=k}^{n-1} \binom{n}{\ell} \left(\ell q^{\ell-1}(1-q)^{n-\ell} - (n-\ell)q^\ell(1-q)^{n-\ell-1}\right) \\
&= nq^{n-1} + n\sum_{\ell=k}^{n-1} \left(\binom{n-1}{\ell-1} q^{\ell-1}(1-q)^{n-\ell} - \binom{n-1}{\ell} q^\ell (1-q)^{n-\ell-1}\right) \\
&= n\binom{n-1}{k-1} q^{k-1}(1-q)^{n-k} \geq 0 \text{ as } q \in [0,1].
\end{aligned}$$

□

5.6.2 Supplement to Section 5.3

5.6.2.1 Total Output

If $B_L = 0$ and $B_H = 1$ in the linear threshold objective (Definition 5.5), then the objective becomes simply $\mathbb{E}[\beta(v)]$, i.e., the well-studied objective of maximizing the total output. For the unit-sum case, it has been shown by [54] that the optimal allocation of prizes is to award the entire prize to the top-ranked player, i.e., set $\alpha_1 = 1$ and $\alpha_j = 0$ for $j > 1$. This can also be observed from the following analysis:

$$\begin{aligned}
\int_0^1 \beta_S(v) f(v) dv &= \int_0^1 \left(\sum_{j\in[n-1]} \frac{\alpha_j}{j} \int_0^v x f_{n-1,j}(x)dx \right) f(v)dv \\
&= \sum_{j\in[n-1]} \frac{\alpha_j}{j} \int_0^1 \left(\int_0^v x f_{n-1,j}(x)dx \right) f(v)dv \\
&= \sum_{j\in[n-1]} \frac{\alpha_j}{j} \int_0^1 \left(\int_x^1 f(v)dv \right) x f_{n-1,j}(x)dx \\
&= \sum_{j\in[n-1]} \alpha_j \int_0^1 \frac{x}{j} \left(1 - F(x)\right) f_{n-1,j}(x)dx \\
&= \frac{1}{n} \sum_{j\in[n-1]} \alpha_j \int_0^1 x f_{n,j+1}(x)dx.
\end{aligned}$$

As $\int_0^1 x f_{n,j+1}(x)dx$ is the expectation of the $(j+1)$-st highest value out of n i.i.d. samples from F, we have

$$\int_0^1 x f_{n,j+1}(x)dx \geq \int_0^1 x f_{n,k+1}(x)dx \text{ for } j \leq k.$$

Therefore, the optimal output function is β_{S1}.
For the unit-range objective, following similar steps, we get:

$$\int_0^1 \beta_R(v) f(v) dv = \frac{1}{n} \sum_{j \in [n-1]} \alpha_j \cdot j \cdot \int_0^1 x f_{n,j+1}(x) dx.$$

While for unit-sum, $\int_0^1 x f_{n,j+1}(x) dx$ decreases as j increases, for unit-range, $j \int_0^1 x f_{n,j+1}(x) dx$ does not necessarily decrease as j increases. Nevertheless, we can see that the optimal contest is a simple contest and the output function in the optimal contest is β_{Rj^*} where:

$$j^* \in \arg\max_{j \in [n-1]} \left(j \int_0^1 x f_{n,j+1}(x) dx \right).$$

5.6.2.2 Concave and Convex Transformations

Let $h : [0,1] \to [0,1]$ be a monotone non-decreasing differentiable function, either concave or convex. The contest designer's objective function is:

$$\mathbb{E}_{\boldsymbol{v}}[\sum_{i \in [n]} h(\beta(v_i))] \equiv \mathbb{E}_v[h(\beta(v))] = \int_0^1 h(\beta(v)) f(v) dv. \tag{5.20}$$

In Section 5.3, for both unit-sum and unit-range, we derived that $\beta(v) = \sum_{j \in [n-1]} \alpha_j \beta_j(v)$ where $\boldsymbol{\alpha} = (\alpha_j)_{j \in [n-1]}$ s.t. $\sum_{j \in [n-1]} \alpha_j = 1$ and $\alpha_j \geq 0$ for $j \in [n-1]$. For unit-sum, $\beta_j(v) = \beta_{Sj}(v) = \frac{1}{j} \int_0^v t f_{n-1,j}(t) dt$ (equation (5.10)), and for unit-range, $\beta_j(v) = \beta_{Rj}(v) = \int_0^v t f_{n-1,j}(t) dt$ (equation (5.11)). We shall use these results in this section.

Lemma 5.19. *If h is convex (concave), then the objective function of the contest designer is also convex (concave) w.r.t. $\boldsymbol{\alpha}$.*

Proof. Take $\gamma \in [0,1]$ and two vectors $\boldsymbol{\alpha}^{(1)}$ and $\boldsymbol{\alpha}^{(2)}$, and let $\boldsymbol{\alpha} = \gamma \boldsymbol{\alpha}^{(1)} + (1-\gamma)\boldsymbol{\alpha}^{(2)}$. Let $OBJ(\boldsymbol{\alpha})$ denote the objective value corresponding to $\boldsymbol{\alpha}$, similarly, $OBJ(\boldsymbol{\alpha}^{(1)})$ and $OBJ(\boldsymbol{\alpha}^{(2)})$.

$$\begin{aligned} OBJ(\boldsymbol{\alpha}) &= \int_0^1 h(\sum_{j \in [n-1]} \alpha_j \beta_j(v)) f(v) dv \\ &= \int_0^1 h(\gamma \sum_{j \in [n-1]} \alpha_j^{(1)} \beta_j(v) + (1-\gamma) \sum_{j \in [n-1]} \alpha_j^{(1)} \beta_j(v)) f(v) dv. \end{aligned}$$

If h is convex, then $h(\gamma x + (1-\gamma) y) \geq \gamma h(x) + (1-\gamma) h(y)$ for any x and y, and therefore, we get $OBJ(\boldsymbol{\alpha}) \geq \gamma OBJ(\boldsymbol{\alpha}^{(1)}) + (1-\gamma) OBJ(\boldsymbol{\alpha}^{(2)})$, so OBJ is convex. Similarly, we can prove that if h is concave, then OBJ is concave. $\square$

For convex h, the results are similar to the results for the total output objective (Section 5.6.2.1). For both unit-sum and unit-range, there is a simple contest that is also optimal. Moreover, for unit-sum, the simple contest that awards the entire prize to the first ranked player is also optimal. These results are similar to the results for a model where the players have concave cost functions, as studied by [81]. There is a similar correspondence for concave h and convex cost functions.

Theorem 5.20. *If h is convex, then there is a simple and optimal rank-order allocation contest.*

Proof. From Lemma 5.19, we know that if h is convex then the objective function is also convex w.r.t. $\boldsymbol{\alpha}$. Any $\boldsymbol{\alpha}$ can be written as a convex combination of the corner points, where a corner point has $\alpha_j = 1$ for some $j \in [n-1]$ and $\alpha_k = 0$ for all $k \neq j$. So, the optimal value of the objective is achieved at some of the corner points (it may be achieved at other points also). □

Theorem 5.21. *If h is convex, and the prizes are unit-sum, then the simple contest that awards the entire prize to the first-ranked player is optimal.*

Proof. From Theorem 5.20, we know that there is an optimal simple contest. Let this simple contest have $\alpha_j = 1$ for some $j \in [n-1]$ and $\alpha_k = 0$ for all $k \neq j$, so $\beta(v) = \beta_{Sj}(v)$. Comparing the objective value of the contest where $\alpha_1 = 1$ with this contest, we have:

$$\begin{aligned} \int_0^1 h(\beta_{S1}(v)) f(v) dv - \int_0^1 h(\beta_{Sj}(v)) f(v) dv &= \int_0^1 (h(\beta_{S1}(v)) - h(\beta_{Sj}(v))) f(v) dv \\ &\geq \int_0^1 h'(\beta_{Sj}(v))(\beta_{S1}(v) - \beta_{Sj}(v)) f(v) dv, \qquad \text{as } h \text{ is convex.} \end{aligned}$$

Using (i) $h'(\beta_{Sj}(v))$ is monotonically increasing and non-negative because $\beta_{Sj}(v)$ is monotonically increasing and h is convex, (ii) $\beta_{S1}(v)$ is single-crossing w.r.t. $\beta_{Sj}(v)$, and (iii) $\int_0^1 (\beta_{S1}(v) - \beta_{Sj}(v)) f(v) dv \geq 0$, as proved while studying total output (Section 5.6.2.1), we have our required result. □

Theorem 5.22. *If h is concave, then the optimal rank-order allocation contest need not be simple, but can be efficiently found by solving a concave maximization problem.*

Proof. The linear threshold objective with only an upper threshold is an example of a concave h (we smooth the non-differentiable point), and for this objective, we already gave an example that shows that a simple contest may not be optimal (Section 5.3.2). From Lemma 5.19, we know that the objective is concave if h is concave, so the optimal contest can be solved efficiently by solving a concave maximization problem (equivalently a convex minimization problem). □

5.6.2.3 Omitted Proofs

Proof of Theorem 5.4. First, let us prove that $tf_{n-1,j}(t)$ is single-crossing w.r.t. $tf_{n-1,k}(t)$ for $j < k$. We will argue that the following inequality is true for sufficiently large values of t:

$$\begin{aligned}
&tf_{n-1,j}(t) > tf_{n-1,k}(t) \\
&\Leftrightarrow t\frac{(n-1)!}{(j-1)!(n-j-1)!}F(t)^{n-j-1}(1-F(t))^{j-1}f(t) \\
&\quad > t\frac{(n-1)!}{(k-1)!(n-k-1)!}F(t)^{n-k-1}(1-F(t))^{k-1}f(t) \\
&\Leftrightarrow \frac{(k-1)!(n-k-1)!}{(j-1)!(n-j-1)!} > \left(\frac{1-F(t)}{F(t)}\right)^{k-j}.
\end{aligned}$$

Indeed, for fixed values of j, k, and n, the left-hand side of the above inequality is a positive constant. As t increases, the right-hand side monotonically decreases from ∞ to 0. Thus, the above inequality is true for all t above some t^*, and we have proved that $tf_{n-1,j}(t)$ is single-crossing w.r.t. $tf_{n-1,k}(t)$.

Following similar steps, we can prove that $t\frac{1}{j}f_{n-1,j}(t)$ is single-crossing w.r.t. $t\frac{1}{k}f_{n-1,k}(t)$ for $j < k$. Instead of the inequality $\frac{(k-1)!(n-k-1)!}{(j-1)!(n-j-1)!} > (\frac{1-F(t)}{F(t)})^{k-j}$ above, we have the inequality $\frac{k!(n-k-1)!}{j!(n-j-1)!} > (\frac{1-F(t)}{F(t)})^{k-j}$, but the subsequent argument applies.

Now, let us prove that β is single-crossing w.r.t. β'. Let us assume that the following inequality is true:

$$\beta'(v) > \beta(v).$$

For unit-range, we have:

$$\begin{aligned}
&\beta'_R(v) > \beta_R(v) \\
\Leftrightarrow &\sum_{l\in[n-1]} \alpha'_l \int_0^v xf_{n-1,l}(x)dx > \sum_{l\in[n-1]} \alpha_l \int_0^v xf_{n-1,l}(x)dx \\
\Leftrightarrow &\int_0^v xf_{n-1,j}(x)dx > \int_0^v xf_{n-1,k}(x)dx.
\end{aligned}$$

Using Lemma 5.17 and the fact that $tf_{n-1,j}(t)$ is single-crossing w.r.t. $tf_{n-1,k}(t)$ for $j < k$, we obtain that β'_R is single-crossing w.r.t. β_R for unit-range. Similarly, for unit-sum, we use Lemma 5.17 together with the result that $t\frac{1}{j}f_{n-1,j}(t)$ is single-crossing w.r.t. $t\frac{1}{k}f_{n-1,k}(t)$ for $j < k$ to prove that β'_S is single-crossing w.r.t. β_S.

As β' is single-crossing w.r.t. β, using Lemma 5.18, we get that β^{-1} is single-crossing w.r.t. β'^{-1}. □

Proof of Theorem 5.5. We need to prove that for any threshold value B, the optimal $\boldsymbol{\alpha}$ has $\alpha_j = 1$ for some j and $\alpha_k = 0$ for $k \neq j$. In other words, $\beta = \beta_j$ for some $j \in [n-1]$.

From Definition 5.4, and equations (5.10) and (5.11), we have:

$$\int_0^1 \mathbb{1}\{\beta(v) \geq B\} f(v) dv = \int_0^1 \mathbb{1}\{v \geq \beta^{-1}(B)\} f(v) dv = \int_{\beta^{-1}(B)}^1 f(v) dv$$
$$= 1 - F(\beta^{-1}(B)).$$

Thus, to maximize the binary threshold objective, we need to minimize $F(\beta^{-1}(B))$, and as F is non-decreasing, we need to minimize $\beta^{-1}(B)$. For any $\boldsymbol{\alpha}$ we have $\alpha_j \geq 0$ and $\sum_j \alpha_j = 1$, and for every j, the β_j functions are monotone. Therefore $\min_j \beta_j^{-1}(B) \leq \beta^{-1}(B)$. Thus, the optimal contest is simple, and the output function for the optimal contest is β_{j^*}, where j^* is:

$$j^* = \arg\min_j \beta_j^{-1}(B). \qquad \square$$

Proof of Theorem 5.6. Let us assume that $\boldsymbol{\alpha}$ is optimal and β is the induced output function. Let $V_H = \beta^{-1}(B_H)$. Let $I = \{i \in [n-1] \mid \alpha_i > 0\}$. If $|I| < 2$, we are trivially done so we can assume w.l.o.g. that $|I| \geq 2$. Select arbitrary $i, j \in I$, $i \neq j$. Let $\alpha_{ij} = \alpha_i + \alpha_j$ and $\gamma = \alpha_i/\alpha_{ij}$. Observe that $\alpha_i = \gamma\alpha_{ij}$ and $\alpha_j = (1-\gamma)\alpha_{ij}$. Also, as $\alpha_i, \alpha_j > 0$, we have $0 < \gamma < 1$.

Now, let us fix α_{ij} and α_k ($k \notin \{i, j\}$) and focus on γ. As $\boldsymbol{\alpha}$ is optimal and γ is strictly between 0 and 1, we have $\frac{\partial LT}{\partial \gamma} = 0$:

$$\frac{\partial LT}{\partial \gamma} = \frac{\partial \int_0^1 \min(B, \sum_{k \in [n-1]} \alpha_k \beta_k(v)) f(v) dv}{\partial \gamma} = 0$$
$$\implies \frac{\partial \int_0^1 \min(B, \sum_{k \in [n-1] \setminus \{i,j\}} \alpha_k \beta_k(v) + \alpha_{ij}(\gamma\beta_i(v) + (1-\gamma)\beta_j(v))) f(v) dv}{\partial \gamma} = 0$$
$$\implies \int_0^{V_H} (\beta_i(v) - \beta_j(v)) f(v) dv = 0$$
$$\implies \int_0^{V_H} \beta_i(v) f(v) dv = \int_0^{V_H} \beta_j(v) f(v) dv. \tag{5.21}$$

For any $i, j \in I$, we get $\int_0^{V_H} \beta_i(v) f(v) dv = \int_0^{V_H} \beta_j(v) f(v) dv = \int_0^{V_H} \beta(v) f(v) dv$.

Now, we will construct an $\boldsymbol{\alpha}'$ with at most two of its components strictly greater than 0. As β is a convex combination of the β_i output functions, we must either have (i) $\beta_i(V_H) = \beta(V_H)$ for some $i \in I$, or (ii) $\beta_i(V_H) > \beta(V_H)$ and $\beta_j(V_H) < \beta(V_H)$ for some $i, j \in I$. For (i) set $\alpha'_i = 1$ and $\alpha'_k = 0$ for $k \notin [n-1] \setminus \{i\}$, and for (ii) set $\alpha'_i = \frac{B - \beta_j(V_H)}{\beta_i(V_H) - \beta_j(V_H)}$ and $\alpha'_j = (1 - \alpha'_i)$ and $\alpha'_k = 0$ for $k \in [n-1] \setminus \{i, j\}$.

We can check that $\sum_{k \in [n-1]} \alpha_k \beta_k(V_H) = \sum_{k \in [n-1]} \alpha'_k \beta_k(V_H) = B$. So, we get $\beta'^{-1}(B_H) = \beta^{-1}(B_H) = V_H$, where β' is the output function induced by $\boldsymbol{\alpha}'$. Also, it is easy to verify that the objective value does not change:

$$
\begin{aligned}
&\int_0^1 \min(B, \sum_{j \in [n-1]} \alpha_j \beta_j(v)) f(v) dv = \int_0^1 \min(B, \sum_{j \in I} \alpha_j \beta_j(v)) f(v) dv \\
&= \int_0^{V_H} \sum_{j \in I} \alpha_j \beta_j(v) f(v) dv + B(1 - F(V_H)) \\
&= \sum_{j \in I} \alpha_j \int_0^{V_H} \beta_j(v) f(v) dv + B(1 - F(V_H)) \\
&= \sum_{j \in I} \alpha'_j \int_0^{V_H} \beta_j(v) f(v) dv + B(1 - F(V_H)) \quad \text{(using equation (5.21))} \\
&= \int_0^1 \min(B, \sum_{j \in [n-1]} \alpha'_j \beta_j(v)) f(v) dv.
\end{aligned}
$$

□

Proof of Theorem 5.7. Let us assume that $\boldsymbol{\alpha}$ is optimal and β is the induced output function. Let $V_L = \beta^{-1}(B_L)$ and $V_H = \beta^{-1}(B_H)$. Let $V_H = \beta^{-1}(B_H)$. Let $I = \{i \in [n-1] \mid \alpha_i > 0\}$; if $|I| < 2$, we are trivially done, so assume w.l.o.g. that $|I| \geq 2$. Select arbitrary $i, j \in I$, $i \neq j$. Let $\alpha_{ij} = \alpha_i + \alpha_j$ and $\gamma = \alpha_i / \alpha_{ij}$. Observe that $\alpha_i = \gamma \alpha_{ij}$ and $\alpha_j = (1 - \gamma) \alpha_{ij}$. Also, as $\alpha_i, \alpha_j > 0$, we have $0 < \gamma < 1$.

Let us fix α_{ij} and α_l ($l \notin \{i, j\}$) and focus on γ. As $\boldsymbol{\alpha}$ is optimal and γ is strictly between 0 and 1, we have $\frac{\partial LT}{\partial \gamma} = 0$:

$$
\begin{aligned}
&\frac{\partial LT}{\partial \gamma} = \frac{\partial \int_0^1 \max(B_L, \min(B_H, \sum_{l \in [n-1]} \alpha_l \beta_l(v))) f(v) dv}{\partial \gamma} = 0 \\
&\implies \frac{\partial \int_0^1 \max(B_L, \min(B_H, \sum_{k \in [n-1] \setminus \{i,j\}} \alpha_k \beta_k(v) + \alpha_{ij}(\gamma \beta_i(v) + (1 - \gamma) \beta_j(v)))) f(v) dv}{\partial \gamma} = 0 \\
&\implies \int_{V_L}^{V_H} (\beta_i(v) - \beta_j(v)) f(v) dv = 0 \\
&\implies \int_{V_L}^{V_H} \beta_i(v) f(v) dv = \int_{V_L}^{V_H} \beta_j(v) f(v) dv.
\end{aligned}
$$

Thus, for every $i \in I$, we have $\int_{V_L}^{V_H} \beta_i(v) f(v) dv = \int_{V_L}^{V_H} \beta(v) f(v) dv$.

Now, let us look at the constraints that $\boldsymbol{\alpha}$ satisfies:

1. $\beta(V_L) = \sum_{i \in I} \alpha_i \beta_i(V_L) = B_L$;

2. $\beta(V_H) = \sum_{i \in I} \alpha_i \beta_i(V_H) = B_H$;

3. $\sum_{i \in I} \alpha_i = 1$;

4. $\alpha_i \geq 0$ for $i \in I$;

5. $\alpha_i = 0$ for $i \notin I$.

Observe that any other $\boldsymbol{\alpha}'$ that satisfies the $3 + |I| + (n - 1 - |I|) = n + 2$ constraints given above will also be optimal (because $\int_{V_L}^{V_H} \beta_i(v) f(v) dv = \int_{V_L}^{V_H} \beta(v) f(v) dv$ for all $i \in I$ and any $\boldsymbol{\alpha}$ that satisfies the $(n + 2)$ constraints will have the same value for the objective). The feasible region of the $(n + 2)$ constraints is bounded on all sides and therefore has corner points. At a corner point, at least $(n - 1)$ of the constraints must be satisfied with an equality, so at least $(n - 4)$ of them are of the type $\alpha_i = 0$ for some $i \in [n - 1]$. Hence, selecting such a corner point, we will get a solution with at most 3 of the coordinates of $\boldsymbol{\alpha}$ strictly greater than 0. □

Proof of Theorem 5.8. Let us take an optimal solution $\boldsymbol{\alpha}$ with the minimum number of non-zero entries, i.e., the minimum number of indices i with $\alpha_i > 0$.

1. If there is only one such index, then we have an approximation ratio of 1 and we are done.

2. Now suppose there are three such indices. Let i, j, k be the three indices for which $\alpha_i, \alpha_j, \alpha_k > 0$ in the optimal solution. W.l.o.g. let $\beta_i(V_L) \leq \beta_j(V_L) \leq \beta_k(V_L)$. We claim that $\beta_i(V_H) \geq \beta_j(V_H) \geq \beta_k(V_H)$. If this were not true, then for a pair, say i, j, we would have had $\beta_i(V_L) \leq \beta_j(V_L)$ and $\beta_i(V_H) \leq \beta_j(V_H)$, with at least one of the two inequalities strict. As the output functions are single-crossing, we would have $\beta_i(v) \leq \beta_j(v)$ for all $v \in [V_L, V_H]$, and we could increase α_j by α_i and decrease α_i to 0 to get a better solution.

   As $\beta_i(V_L) \leq \beta_j(V_L) \leq \beta_k(V_L)$ and $\beta_i(V_H) \geq \beta_j(V_H) \geq \beta_k(V_H)$, their convex combination, β, has: $\beta_i(V_L) \leq \beta(V_L) \leq \beta_k(V_L)$ and $\beta_i(V_H) \geq \beta(V_H) \geq \beta_k(V_H)$.

3. Finally, if there are only two positive entries, say α_i and α_k, then also we can prove a similar condition: $\beta_i(V_L) \leq \beta(V_L) \leq \beta_k(V_L)$ and $\beta_i(V_H) \geq \beta(V_H) \geq \beta_k(V_H)$.

Look at a generic plot of β_i, β_k, β given in Figure 5.1. With reference to the figure, we have:

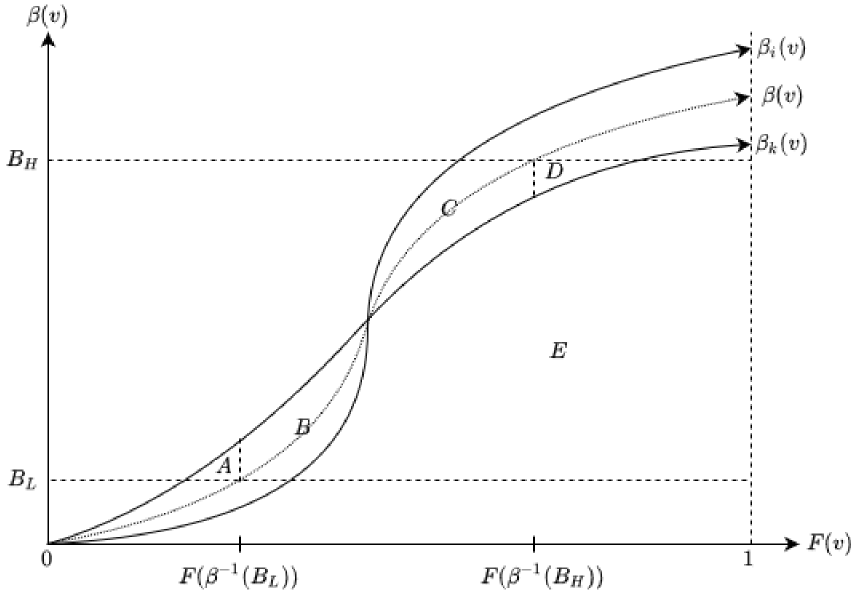

Figure 5.1: Plot of β_i, β_k, β where $i < k$.

- $\int_0^1 \max(B_L, \min(B_H, \beta_i(v)))f(v)dv = C + D + E + B_L$;
- $\int_0^1 \max(B_L, \min(B_H, \beta_k(v)))f(v)dv = A + B + E + B_L$;
- $\int_0^1 \max(B_L, \min(B_H, \beta(v)))f(v)dv = B_L F(V_L) + \int_{V_L}^{V_H} \beta(v)f(v)dv + B_H(1 - F(V_H))$, which is, by Theorem 5.7, equal to: $B_L F(V_L) + \int_{V_L}^{V_H} \beta_k(v)f(v)dv + B_H(1 - F(V_H)) = B + D + E + B_L$.

The approximation ratio is at most

$$\frac{B + D + E + B_L}{\max(A + B + E + B_L, C + D + E + B_L)} = \frac{E + B_L + B + D}{E + B_L + \max(A + B, C + D)} \leq \frac{B + D}{\max(A + B, C + D)} \leq \frac{B + D}{\max(B, D)} \leq 2.$$

□

5.6.3 Supplement to Section 5.4

5.6.3.1 Linear Threshold Objective: Irregular Distributions (continued)

Here, we provide an algorithm to find the (approximately) optimal contest. We perform an approximate search on the three parameters V_H, $\xi(V_H)$, and V_L, to maximize the linear threshold objective, using an algorithm sketched below:

1. Select values for V_H and $\xi(V_H)$.

2. Assume that ξ is constant in the interval $[l(V_H), V_H)$. Applying Lemma 5.15 to the point $l(V_H)$, compute $\xi(l(V_H))$:

$$\begin{aligned}\int_{l(V_H)}^{1} \xi(t)f(t)dt &= \xi(l(V_H))(F(V_H) - F(l(V_H))) + \xi(V_H)(1 - F(V_H)) \\ &= \frac{1 - F(l(V_H))^n}{n(1 - F(l(V_H)))} \\ \implies \xi(l(V_H)) &= \frac{1 - F(l(V_H))^n}{n(1 - F(l(V_H)))(F(V_H) - F(l(V_H)))} - \frac{\xi(V_H)(1 - F(V_H))}{F(V_H) - F(l(V_H))}.\end{aligned}$$

3. Compute $\overline{V_L}$ by solving $\beta(V_H) = V_H\xi(V_H) - \int_{\overline{V_L}}^{V_H} \xi(v)dv$ by the bisection method, where the value of $\xi(v)$ for $\overline{V_L} \leq v < l(V_H)$ is given as

$$\xi(v) = \begin{cases} F(v)^{n-1}, & \text{if } \overline{H}(F^{-1}(v)) = H(F^{-1}(v)) \text{ or } v \in [l(\overline{V_L}), r(\overline{V_L})] \\ \frac{F(r(v))^n - F(l(v))^n}{n(F(r(v)) - F(l(v)))}, & \text{if } \overline{H}(F^{-1}(v)) < H(F^{-1}(v)) \text{ and } v > r(\overline{V_L}) \end{cases}$$

4. If $\overline{H}(F^{-1}(\overline{V_L})) = H(F^{-1}(\overline{V_L}))$, then set $V_L = \overline{V_L}$, otherwise search for the $V_L \in [l(\overline{V_L}), r(\overline{V_L})]$ (and automatically for $\xi(V_L) \geq B_L$) that satisfies $\int_V^{r(\overline{V_L})} F(v)^{n-1}dv = \int_{V_L}^{r_{V_L,V_H}(V_L)} \xi(V_L)dv + \int_{r_{V_L,V_H}(V_L)}^{r(\overline{V_L})} F(v)^{n-1}dv$ and maximizes $B_L F(V_L) + \int_{V_L}^{r(\overline{V_L})} \beta(v)f(v)dv$. This step selects V_L to optimally redistribute the area under ξ in the interval $[\overline{V_L}, r(\overline{V_L})]$. [6]

5.6.3.2 Omitted Proofs

Proof of Theorem 5.9. For unit-range allocations, we optimize ξ subject to the constraints: $0 \leq \xi(v) \leq 1$. We have

$$\beta(v) = v\xi(v) - \int_0^v \xi(t)dt \leq v\xi(v) \leq v,$$

where the first inequality holds because $\xi(t) \geq 0$ and the second inequality holds because $\xi(t) \leq 1$ for all $t \in [0, 1]$. We have $\beta(v) \leq v \implies B \leq \beta^{-1}(B)$. Set $\xi(v) = 0$ for $v < B$ and $\xi(v) = 1$ for $v \geq B$. We have $\xi(B) = 1$ and $\int_0^B \xi(t)dt = 0$, so we get $\beta(B) = B\xi(B) - \int_0^B \xi(t)dt = B \implies \beta^{-1}(B) \leq B$. As we have already seen that $\beta^{-1}(B) \geq B$, this is optimal.

For unit-sum, we have an additional constraint on ξ, inequality (5.6): $\int_V^1 \xi(v)f(v)dv \leq \frac{1-F(V)^n}{n}$ for every V.

[6] We skipped a corner case: if $l(V_H)\xi(l(V_H)) < B_L$, then V_L must be greater than $l(V_H)$. Find the V_L in $[l(V_H), V_H]$ that optimizes the objective following a procedure similar to step (4). Also, we need to redistribute ξ in $[l(V_H), V_H)$ according to $\overline{\psi}_{V_L,V_H}$ to find $\xi(l_{V_L,V_H}(V_H))$ in a manner similar to step (2).

Lemma 5.23. *If ξ is optimal, we can assume w.l.o.g. that $\xi(v) = 0$ for $v < \beta^{-1}(B)$.*

Proof. We know that $\beta(V) = V\xi(V) - \int_0^V \xi(x)dx \leq V\xi(V)$. Hence, by setting $\xi(v) = 0$ for $v < V$ we still have $\beta^{-1}(B) = V$. $\square$

Lemma 5.24. *If ξ is optimal, we can assume w.l.o.g. that $\xi(v) = \xi(V)$ for $v \geq \beta^{-1}(B)$, i.e., $\xi(v)$ is constant for $v \geq \beta^{-1}(B)$.*

Proof. We define a transformed expected allocation function $\bar{\xi}$, where $\bar{\xi}(v) = \xi(v)$ for $v < V$, and $\bar{\xi}(v) = \frac{1}{1-F(V)} \int_V^1 \xi(t)f(t)dt$ for $v \geq V$. Let $\bar{\beta}$ be the output function for $\bar{\xi}$. As ξ is monotone, $\bar{\xi}(V) \geq \xi(V)$, and therefore $\bar{\beta}(V) \geq \beta(V) \geq B$. So, we still have $\bar{\beta}^{-1}(B) = V$. $\square$

From the previous two lemmas, we know that the allocation function equally distributes the prize among the players who have an ability above some value V, where $\beta(v) \geq B$ for $v \geq V$ and $\beta(v) < B$ otherwise. As ξ satisfies $\int_V^1 \xi(v)f(v)dv \leq \frac{1-F(V)^n}{n}$, we have:

$$\xi(V) \leq \frac{1 - F(V)^n}{n(1 - F(V))} \implies B \leq \beta(V) \leq \frac{V(1 - F(V)^n)}{n(1 - F(V))}.$$

As the expression $\frac{V(1-F(V)^n)}{n(1-F(V))}$ increases with V, and we want to minimize V, it is optimal to satisfy the above inequality with an equality. This gives us $B = \frac{V(1-F(V)^n)}{n(1-F(V))}$.

Also, as $\frac{V(1-F(V)^n)}{n(1-F(V))}$ is continuous and non-decreasing in $[0, 1]$, we can efficiently find V. However, we do not need to explicitly compute V because the contest that equally distributes the prize to all players who generate an output above B, if there are any such players in a given ability profile, automatically induces the required contest. $\square$

Proof of Lemma 5.12. We will show that pushing the area under ξ to the right does not decrease the objective. For the initial portion of the proof, let us disregard these two constraints: ξ is monotone and $\int_v^1 \xi(t)f(t)dt \leq \frac{1}{n}(1 - F(v)^n)$. At the end, we will briefly explain how to incorporate these constraints into the proof.

Take two points u and v such that $u < v$, $\xi(u) > 0$ and $\beta(v) < B$. For very small δ and ϵ greater than 0, let us decrease the area under ξ in a small neighborhood $[u - \epsilon, u + \epsilon)$ of u by $\delta/f(u)$ and increase the area under ξ in a small neighborhood $[v - \epsilon, v + \epsilon)$ of v by $\delta/f(v)$. Let $\bar{\xi}$ and $\bar{\beta}$ be the transformed ξ and β after the update. Select ϵ and δ small enough to maintain $\bar{\xi}(u - \epsilon) \geq 0$ and $\bar{\beta}(v + \epsilon) \leq B_H$.

Let us compute the change in the linear threshold objective value. To do that, first, let us look at $\bar{\beta}(y) - \beta(y)$ for $y \in [0,1]$:

$$\bar{\beta}(y) - \beta(y) = \begin{cases} 0, & \text{if } y \in [0, u-\epsilon) \\ \frac{-u\delta}{2\epsilon f(u)}, & \text{if } y \in [u-\epsilon, u+\epsilon) \\ \frac{\delta}{f(u)}, & \text{if } y \in [u+\epsilon, v-\epsilon) \\ \frac{v\delta}{2\epsilon f(v)} + \frac{\delta}{f(u)}, & \text{if } y \in [v-\epsilon, v+\epsilon) \\ \frac{\delta}{f(u)} - \frac{\delta}{f(v)}, & \text{if } y \in [v+\epsilon, 1] \end{cases}$$

As we are moving area from left to right, i.e., $u < v$, it is easy to check that for the lower threshold B_L, $V_L = \bar{\beta}^{-1}(B_L)$ does not decrease, so neither does $B_L F(V_L)$. For the remaining portion of the linear threshold objective (5.12), we have the following:

$$\begin{aligned}
&\int_{y \geq V_L} (\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y))) f(y) dy \\
&= \int_{y=u-\epsilon}^{u+\epsilon} \frac{-u\delta}{2\epsilon f(u)} f(y) dy + \int_{y=u+\epsilon}^{v-\epsilon} \frac{\delta}{f(u)} f(y) dy + \int_{y=v-\epsilon}^{v+\epsilon} \left(\frac{v\delta}{2\epsilon f(v)} + \frac{\delta}{f(u)}\right) f(y) dy \\
&\quad + \int_{y=v+\epsilon}^{1} (\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y))) f(y) dy, \qquad \text{because } \bar{\beta}(v+\epsilon) \leq B_H \\
&= -u\delta + \frac{\delta}{f(u)}(F(v) - F(u)) + v\delta \\
&\quad + \int_{y=v}^{1} (\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y))) f(y) dy, \qquad \text{as } \epsilon \to 0.
\end{aligned}$$

We will now consider two cases: (1) $f(u) \leq f(v)$ and (2) $f(u) > f(v)$.

Case (1): $f(u) \leq f(v) \implies \delta/f(u) \geq \delta/f(v)$. For $y \geq v + \epsilon$, $\bar{\beta}(y) - \beta(y) = \delta/f(u) - \delta/f(v) \geq 0$, so $\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y)) \geq 0$. The total change in the objective is:

$$\begin{aligned}
&- u\delta + \frac{\delta}{f(u)}(F(v) - F(u)) + v\delta + \int_{y=v}^{1} (\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y))) f(y) dy \\
&\qquad \geq 0, \qquad \text{as } v > u, F(v) \geq F(u), \text{ and } \min(B_H, \bar{\beta}(y)) \geq \min(B_H, \beta(y)).
\end{aligned}$$

Case (2): $f(u) > f(v) \implies \delta/f(u) \leq \delta/f(v)$. For $y \geq v + \epsilon$, $\bar{\beta}(y) - \beta(y) = \delta/f(u) - \delta/f(v) \leq 0$, so $\min(B_H, \bar{\beta}(y)) - \min(B_H, \beta(y)) \geq \frac{\delta}{f(u)} - \frac{\delta}{f(v)}$. The total change in the objective is:

$$
\begin{aligned}
& -u\delta + \frac{\delta}{f(u)}(F(v)-F(u)) + v\delta + \int_{y=v}^{1}(\min(B_H,\bar{\beta}(y)) - \min(B_H,\beta(y)))f(y)dy \\
&= -u\delta + \frac{\delta}{f(u)}(F(v)-F(u)) + v\delta + \int_{y=v}^{1}(\frac{\delta}{f(u)} - \frac{\delta}{f(v)})f(y)dy \\
&= -u\delta + \frac{\delta}{f(u)}(F(v)-F(u)) + v\delta + (\frac{\delta}{f(u)} - \frac{\delta}{f(v)})(1-F(v)) \\
&= \delta(-u + \frac{F(v)-F(u)}{f(u)} + v + \frac{1-F(v)}{f(u)} - \frac{1-F(v)}{f(v)}) \\
&= \delta(v - \frac{1-F(v)}{f(v)} - (u - \frac{1-F(u)}{f(u)})) \\
&= \delta(\psi(v) - \psi(u)) \geq 0, \qquad \text{because } F \text{ is regular.}
\end{aligned}
$$

We proved that transforming ξ in this manner does not decrease the objective value.

For the rest of the proof, let $a(y) = \int_y^1 \xi(t)f(t)dt$ and $b(y) = (1-F(y)^n)/n$.

Now, we explain how to incorporate the two constraints we disregarded in the beginning: non-decreasing property of ξ and $a(y) \leq b(y)$. For the $a(y) \leq b(y)$ constraint: select v to be close to the rightmost[7] point for which $a(v) < b(v)$ and has $\beta(v) < B_H$, select u to be close to the rightmost point less that v for which $a(u) = b(u)$ and $\beta(u) > 0$, if there is no such point u less than v with $a(u) = b(u)$, then select any point with $\beta(u) > 0$. If we select a small enough area to move from u to v (parameterized by δ in the first part of the proof) and select the neighborhood of u and v suitably (parameterized by ϵ earlier) we can satisfy the constraint. For the monotonically non-decreasing property, selecting v to be close to the rightmost point with $a(v) < b(v)$ and $\beta(v) < B_H$, and increasing ξ in very small increments, and repeating until convergence, maintains the non-decreasing property of ξ in the aggregate.

Starting from an arbitrary ξ, one can reach a ξ that satisfies the condition given in the statement of the lemma by transformations to ξ as given above. This completes the proof.

□

Proof of Theorem 5.13. Let us assume that the lemma is false, which gives us $\int_{V_H}^1 \beta(v)f(v)dv > B_H(1-F(V_H))$.

We average out ξ in the interval $v \in [V-\delta, 1]$, for some $\delta > 0$, while maintaining $\beta(v) \geq B$ for $v \geq V - \delta$. This will improve our objective. We can check that the

[7] If the set is open, we select v in limit. Same for u.

transformed ξ satisfies $\xi(v) = \frac{1}{1-F(V-\delta)} \int_{V-\delta}^{1} \xi(v)f(v)dv$ for $v \geq V - \delta$. We can find the δ by solving the following equation:

$$B = \beta(V-\delta) = (V-\delta)\frac{1}{1-F(V-\delta)} \int_{V-\delta}^{1} \xi(v)f(v)dv - \int_{0}^{V-\delta} \xi(v)dv.$$

Observe that the right-hand side is a non-increasing continuous function of δ. As δ goes from 0 to V, the right-hand side goes from strictly above B to 0, so we obtain the required solution for δ. $\square$

Proof of Theorem 5.14. We provide separate proofs for unit-range and unit-sum settings.

Unit-Range. Note that Lemmas 5.10,5.11, and 5.13 are applicable for unit-range, and Lemma 5.12 is applicable with slight modification. Together, they imply that $\xi(v) = 0$ for $v < V$ and $\xi(v) = 1$ for $v \geq V$, for some $V \in [0, 1]$. From Lemma 5.10 we also have $V_L = V \geq B_H$, and from Lemma 5.13, $V \leq V_H \leq B_H$ if there exists a V_H, or $V < B_H$ if not. The objective value can be written as:

$$B_L F(V) + V(1 - F(V)).$$

Differentiating w.r.t. V and equating to 0, we obtain

$$\begin{aligned} B_L f(V) - V f(V) + (1 - F(V)) &= f(V)(B_L - (V - \frac{1-F(V)}{f(V)})) \\ &= f(V)(B_L - \psi(V)) = 0 \\ &\implies \psi(V) = B_L \implies V = \psi^{-1}(B_L). \end{aligned}$$

We can also observe that the solution to the above equation is a global maximum because the derivative of the objective is greater that 0 for $V < \psi^{-1}(B_L)$ and less than 0 afterwards. Plugging in the constraints on V, $B_L \leq V \leq B_H$, the optimal solution is $V = \max(B_L, \min(B_H, \psi^{-1}(B_L)))$.

Unit-Sum. We divide the analysis into two cases, depending on whether β hits the upper threshold B_H.

1. $\beta(v) < B_H$ for $v \in [0, 1]$. We do not have a V_H, and for V_L we have the following inequality:

$$\begin{aligned} \beta(1) < B_H \implies 1\xi(1) - \int_{V_L}^{1} \xi(v)dv < B_H \implies \int_{V_L}^{1} F(v)^{n-1}dv > 1 - B_H \\ \implies V_L < V_{\text{mid}}, \end{aligned}$$

where V_{mid} is the solution of the equation $\int_{V_{\text{mid}}}^{1} F(v)^{n-1}dv = 1 - B_H$. We also know that

$$\beta(V_L) \geq B_L \implies V_L\xi(V_L) = V_L F(V_L)^{n-1} \geq B_L \implies V_L \geq V_{\text{low}},$$

where V_{low} is the solution to $V_{\text{low}} F(V_{\text{low}})^{n-1} = B_L$. Thus, a value of V_L in $[V_{\text{low}}, V_{\text{mid}})$ maximizes the objective:

$$OBJ = B_L F(V_L) + \int_{V_L}^{1} \beta(v) f(v) dv.$$

Now, $\beta(v) = vF(v)^{n-1} - \int_{V_L}^{v} F(t)^{n-1}dt \implies \frac{d\beta(v)}{dV_L} = F(V_L)^{n-1}$. Differentiating OBJ w.r.t. V_H we get:

$$\begin{aligned}
\frac{dOBJ}{dV_H} &= B_L f(V_L) - \beta(V_L) f(V_L) + \int_{V_L}^{1} \frac{d\beta(v)}{dV_L} f(v)dt \\
&= B_L f(V_L) - V_L F(V_L)^{n-1} f(V_L) + F(V_L)^{n-1}(1 - F(V_L)) \\
&= f(V_L)F(V_L)^{n-1}\left(\frac{B_L}{F(V_L)^{n-1}} - V_L + \frac{1 - F(V_L)}{f(V_L)}\right) \\
&= f(V_L)F(V_L)^{n-1}\left(\frac{B_L}{F(V_L)^{n-1}} - \psi(V_L)\right).
\end{aligned}$$

As $\frac{B_L}{F(V_L)^{n-1}}$ decreases with V_L, $\psi(V_L)$ increases with V_L, and $f(V_L)F(V_L)^{n-1}$ is non-negative, the root of $\frac{B_L}{F(V_L)^{n-1}} - \psi(V_L) = 0$ is the global maximum. Also, as the function is continuous, we can efficiently find a solution using a root finding algorithm such as the bisection method; let the solution be $\overline{V_L}$. The optimal V_L for this case will be $V_L = \max(V_{\text{low}}, \min(V_{\text{mid}}, \overline{V_L}))$.

2. $\beta(v) \geq B$ for $v \geq V_H \in [0, 1]$. We have the following equality:

$$\beta(V_H) = V_H\xi(V_H) - \int_{V_L}^{V_H} \xi(v)dv = V_H\eta(F(V_H)) - \int_{V_L}^{V_H} F(v)^{n-1}dv = B_H,$$

where $\eta(x) = \frac{1-x^n}{n(1-x)}$. Note that $\eta'(x) = \frac{1-x^n}{n(1-x)^2} - \frac{nx^{n-1}}{1-x} = \frac{1}{1-x}(\eta(x) - x^{n-1})$ and also that $\eta(x) \geq x^{n-1}$ and $\eta'(x) \geq 0$ for $x \in [0, 1]$.

For $u \geq v$, $\psi_u(v) = v - \frac{F(u)-F(v)}{f(v)}$. Observe that $\psi_u(v)$ is non-decreasing in v because $\frac{\psi_u(v)}{dv} = 2 + \frac{(F(u)-F(v))f'(v)}{f(v)^2}$ is obviously non-negative if $f'(v) \geq 0$, and if $f'(v) < 0$, then $\frac{\psi_u(v)}{dv} = 2 + \frac{(F(u)-F(v))f'(v)}{f(v)^2} \geq 2 + \frac{(1-F(v))f'(v)}{f(v)^2} = \frac{\psi(v)}{dv} \geq 0$.

As $\beta(V_H) = B_H$, we get $V_H \in [V_{\text{up}}, 1]$ and $V_L \in [V_{\text{mid}}, V_{\text{up}}]$ where V_{up} is the solution of the equation $B_H = V_{\text{up}}\eta(F(V_{\text{up}}))$. Differentiating $\beta(V_H) = B_H$

w.r.t. V_L we get:

$$\frac{d\beta(V_H)}{dV_L} = \frac{dB_H}{dV_L} = 0$$

$$\implies (V_H\eta'(F(V_H))f(V_H) + \eta(F(V_H)) - F(V_H)^{n-1})\frac{dV_H}{dV_L} + F(V_L)^{n-1} = 0$$

$$\implies (V_H\eta'(F(V_H))f(V_H) + (1 - F(V_H))\eta'(F(V_H)))\frac{dV_H}{dV_L} + F(V_L)^{n-1} = 0$$

$$\implies \frac{dV_H}{dV_L} = \frac{-F(V_L)^{n-1}}{\eta'(F(V_H))(V_H f(V_H) + (1 - F(V_H)))} \leq 0.$$

Given V_L and V_H, the objective can be written as:

$$OBJ = B_L F(V_L) + \int_{V_L}^{V_H} \beta(v)f(v)dv + B_H(1 - F(V_H)).$$

Differentiating OBJ w.r.t. V_L we get:

$$\begin{aligned}
\frac{dOBJ}{dV_L} &= B_L f(V_L) + \beta(V_H)f(V_H)\frac{dV_H}{dV_L} - \beta(V_L)f(V_L) \\
&\quad + \int_{V_L}^{V_H} F(V_L)^{n-1}f(v)dv - B_H f(V_H)\frac{dV_H}{dV_L} \\
&= B_L f(V_L) - V_L F(V_L)^{n-1}f(V_L) + \int_{V_L}^{V_H} F(V_L)^{n-1}f(v)dv \\
&\quad + (\beta(V_H) - B_H)f(V_H)\frac{dV_H}{dV_L} \\
&= B_L f(V_L) - V_L F(V_L)^{n-1}f(V_L) + \int_{V_L}^{V_H} F(V_L)^{n-1}f(v)dv \\
&= F(V_L)^{n-1}f(V_L)(\frac{B_L}{F(V_L)^{n-1}} - V_L + \frac{F(V_H) - F(V_L)}{f(V_L)}) \\
&= F(V_L)^{n-1}f(V_L)(\frac{B_L}{F(V_L)^{n-1}} - \psi_{V_H}(V_L)).
\end{aligned}$$

From the equation above, to find the solution of $\frac{dOBJ}{dV_L} = 0$, we need to solve for the values of V_L and V_H that satisfy $\frac{B_L}{F(V_L)^{n-1}} - \psi_{V_H}(V_L) = 0$ (and $\beta(V_H) = B_H$). As the F and ψ_{V_H} are continuous, we can efficiently find a solution using a root finding algorithm. Moreover, the pair of values for V_L and V_H that satisfies $\frac{B_L}{F(V_L)^{n-1}} - \psi_{V_H}(V_L) = 0$ is optimal because:

- The first term, $\frac{B_L}{F(V_L)^{n-1}}$, decreases with V_L.
- The second term, $\psi_{V_H}(V_L)$, has a derivative: $\frac{d\psi_{V_H}(V_L)}{dV_L} = \frac{\partial\psi_{V_H}(V_L)}{\partial V_H}\frac{dV_H}{dV_L} + \frac{\partial\psi_{V_H}(V_L)}{\partial V_L}$. As $\frac{\partial\psi_{V_H}(V_L)}{\partial V_H} = \frac{-f(V_H)}{f(V_L)} \leq 0$ and $\frac{dV_H}{dV_L} \leq 0$, we get $\frac{\partial\psi_{V_H}(V_L)}{\partial V_H}\frac{dV_H}{dV_L} \geq 0$. Also, $\frac{\partial\psi_{V_H}(V_L)}{\partial V_L} \geq 0$ as shown earlier. So, $\psi_{V_H}(V_L)$ is a non-decreasing function of V_L.

Let the values of V_L and V_H that satisfy $\frac{B_L}{F(V_L)^{n-1}} - \psi_{V_H}(V_L) = 0$ be $\overline{V_L}$ and $\overline{V_H}$, respectively. Overall, we have the optimal $V_L = \min(V_{\text{up}}, \max(V_{\text{mid}}, \overline{V_L}))$ and the optimal $V_H = \min(1, \max(V_{\text{up}}, \overline{V_H}))$.

One of the two cases, either $\beta(v)$ touches the upper threshold B_H or it does not, will give us the overall optimal solution.

Given the optimal expected allocation function $\xi(v)$, we can easily derive the optimal allocation function $\boldsymbol{x}(\boldsymbol{v})$, given in the theorem statement. $\square$

Proof Sketch of Lemma 5.15. The proof is very similar to Lemma 5.12. The main modifications are: first, we can check that in $[l(v), r(v)]$ if we flatten ξ, i.e., we set $\xi(y) = \frac{\int_{l(v)}^{r(v)} \xi(t) f(t) dt}{F(r(v)) - F(l(v))}$ for $y \in [l(v), r(v)]$, then we do not decrease the objective[8]; second, we account for the change in the objective for transferring area under ξ from $[l(u), r(u)]$ to $[l(v), r(v)]$, in aggregate, rather than from u to v as we did in Lemma 5.12. $\square$

[8]we do this for all points except for the points in $[l(V_L), r(V_L)]$ if $l(V_L) < V_L$, otherwise it might change $\beta^{-1}(B_L)$. Note that the statement of the lemma accommodates for this.

Chapter 6

Incentivizing a Target Group

6.1 Introduction

The dominant paradigm in contest design is to assume that the principal designs the prize structure so as to incentivize the agents to produce a higher total output. In this model, the designer values the contribution from each agent equally, i.e., the marginal output produced by any agent contributes the same marginal value to the designer's objective. In this chapter, we break away from this assumption: instead, we assume that the agents belong to two different groups and the designer may want to prioritize one group—the target group—over the other, non-target group.

Contests that give special importance to agents from a particular group are not uncommon in practice. Conferences give best student paper awards, in addition to best paper awards. Many competitions and hackathons are organized to elicit engagement from underrepresented groups: for example, hackathons as an intervention to get women interested in AI [71]. Contests are widely used for crowdsourcing, which is also an important source of training data for machine learning algorithms; in this context, eliciting input from disadvantaged groups is particularly important, as it helps the algorithms to learn decision-making rules that reflect opinions and preferences of such groups. Our work provides a better understanding of how to encourage contributors from such groups.

Like the previous chapter, we study an incomplete information (Bayesian) model. We assume that each agent is associated with an *ability*, which captures the amount of output they can produce per unit effort. In an incomplete information model for contests, it is generally assumed that the abilities of the agents are selected i.i.d. from a given distribution F. In our model, based on the assumption that the agents belong to one of the two groups, the abilities of the agents in the target and the non-target group are sampled from two distinct distributions F and G, respectively. Each agent belongs to the target group with a fixed probability μ.[1] The agents know the prize allocation scheme, their own ability, and the prior distributions of other agents' abilities. They act strategically to maximize their expected utility, where the utility of an agent is the prize they receive minus the effort they exert, and reach a Bayes–Nash equilibrium. The contest designer knows the prior distributions of the agents' abilities, and therefore can reason about the equilibrium behavior of the agents. She designs the prize allocation scheme so as to elicit equilibrium behavior that optimizes her own objective. We assume that the contest designer has a budget of 1 to use for the prizes.

The contest ranks the agents based on their outputs and awards the prizes based on the ranks. In our general model, we assume that contests have two sequences of prizes, one available to all agents, and the other available to the target group only. A conference sponsoring a best student paper award in addition to a best paper award is an example of such a contest. We also study two specific variants of this model. In the first variant, all prizes in the contest are equally accessible to all agents, and there are no group-specific prizes. One could come across such a design preference in situations where the contest designer does not want to discriminate among the agents, but still wants to incentivize an under-represented agent group. In this case, the contest does not discriminate while awarding the prizes, but the design of the prize structure (i.e., the value of the first prize, second prize, and so on) can incorporate the information about the ability distribution of the agents. This model captures an equal opportunity employer that treats all applicants equally, but has a preference to increase the diversity in the workplace. In the second variant, the contest has group-specific prizes only, with no prizes available to the entire population; examples include a hackathon for women in CS, or a hiring/scholarship contest with a strict group-specific quota.

We assume that the contest designer wants to maximize the expected output of the agents in the target group only. With minimal changes, the equilibrium and optimal contest characterization results (discussed below) for only group-specific or only general prizes can be extended to accommodate multiple groups (instead of just two) and objective functions that maximizes the weighted sum of the group outputs. These extensions are direct from the analysis and do not provide much added insight, so we do not discuss them in detail and remain focused on the two-group case. On

[1] We assume that μ is known based on historical data on the number of participants from the target group compared to the total number of participants. Similarly, the distributions F and G are also known from historical data.

the other hand, when there are both group-specific and general prizes, extension to multiple groups is beyond the scope of our current techniques.

6.1.1 Our Contribution

We initiate a theoretical study of contest design to elicit higher participation by agents from under-represented groups, and propose tractable models for doing so. We analyze the equilibrium behavior of the agents for three cases: *general prizes only*, where there is a sequence of rank-ordered prizes accessible to all agents, whether in the target or the non-target group; *group-specific prizes only*, where there are two different sequences of prizes, each accessible to one of the groups; *both general and target group-specific prizes*, where there are prizes accessible to both groups as well as prizes accessible to the target group only. The techniques used to analyze the equilibrium in these contests can be extended to related prize structures.

Based on these equilibrium characterizations, we study the properties of the contest that maximizes the total expected output of the agents from the target group. For the setting where all prizes are general, we prove that the optimal contest awards a prize of $1/\ell$ to the first ℓ agents and 0 to the remaining agents, where ℓ is a function of the relative frequency of target group agents, μ, and the ability distributions of the target group, F, and the non-target group, G. We give a closed-form formula for ℓ. We also study the effect of first-order and second-order stochastic dominance of F on G, and vice versa, on the optimal contest. For the setting where all prizes are group-specific, we show that it is optimal to award the full prize budget to the top-ranked agent, irrespective of the distributions F and G. This result matches similar results on maximizing total output [81]. We also compare general prizes and group-specific prizes, and show that either of these choices may be preferred depending upon the distributions F and G, and also upon the frequency of the target group agents in the population, μ, even if $F = G$.

For readability, certain proofs are provided at the end of the chapter in Section 6.6.

6.1.2 Related Work

Moldovanu and Sela [81] characterize the Bayes–Nash equilibrium for contests with rank-order allocation of prizes assuming incomplete information with i.i.d. types; Chawla and Hartline [28] prove the uniqueness of this equilibrium. Moldovanu and Sela [81] also show that awarding the entire prize to the top-ranked agent is optimal when the agents have (weakly) concave cost functions, but the optimal mechanism may have multiple prizes for convex cost functions. In contrast, we observe that even for linear cost functions, it may be beneficial to have multiple prizes to maximize the expected output from the target group.

The standard assumption in mechanism design and its sub-area of contest design is that the agents' types are sampled i.i.d. from some distribution. Our work extends the equilibrium analysis of Moldovanu and Sela [81] to an asymmetric model where the agents are from two groups with different distributions. Amann and Leininger [5] study an asymmetric model with two agents with types sampled independently from two different distributions; our analysis uses the idea of the function 'k' introduced by them, see Theorem 6.1. Characterizing equilibrium in contests with many agents with types sampled independently from different distributions has remained technically challenging; recently Olszewski and Siegel [87, 88] made progress by assuming very large numbers of agents and prizes, where an individual agent has infinitesimal effect on the equilibrium.

Bodoh-Creed and Hickman [21] model affirmative action in college admissions using contests. Their model has general non-linear utility and cost functions, but, like Olszewski and Siegel [87, 88], they assume that the numbers of agents and prizes are very large. They give first-order equilibrium conditions and show that two types of affirmative actions—(i) admissions preference schemes, where the outputs of agents from the target group are (artificially) amplified before ranking and prize allocation, and (ii) quotas, where there are separate pools of seats for the different groups (similar to our group-specific prizes only model, Section 6.3.3)—have the same sets of equilibria with identical actions by the agents and identical outcomes. Our model makes a stronger linearity assumption regarding the utility and the cost functions, which allows us to better characterize the equilibrium and study the optimal contest (prize) design problem.

There is a long line of literature on the effect of different types of affirmative action on contests. This literature generally assumes that there is only one prize, which either gets allocated to the top-ranked agent or gets allocated proportionally (Tullock [104] contests and their generalizations), and the contest designer introduces different kinds of interventions in the allocation of this prize. One such intervention is *favoritism*, where the objective is to maximize the total output by giving head-starts or handicaps to certain agents (see, e.g., [66, 51, 37]). A head-start adds a bonus to an agent's output, while a handicap decreases an agent's score by a fixed percentage. See [33] for a survey on contests and affirmative action.

Our work deviates from the widely studied objective of maximizing total output, by focusing on the output of a target group. Several papers have studied objectives other than the total output, such as maximum individual output [29], cumulative output from the top k agents [7, 53], total output of agents producing output in a given range [43].

There are also several papers that perform equilibrium analysis and optimal contest design in the complete information setting (e.g., [12, 11, 99]).

6.2 Model and Preliminaries

There are n agents. Any given agent belongs to the target (resp., non-target) group with probability μ (resp., $(1-\mu)$), independently of the other agents. Let $\boldsymbol{v} = (v_1, v_2, \ldots, v_n)$ be the ability profile of the agents, where v_i values are independent and identically distributed (i.i.d.) random variables from continuous and differentiable distributions F or G with support $[0, 1]$, depending upon whether agent i belongs to the target or the non-target group, respectively. Let f and g be the probability density functions (PDFs) of F and G. The n agents simultaneously produce the output profile $\boldsymbol{b} = (b_i)_{i \in [n]}$, so that the cost of agent i is b_i/v_i.

The contest awards two sequences of prizes $w_1 \geq w_2 \geq \ldots \geq w_n$ and $\omega_1 \geq \omega_2 \geq \ldots \geq \omega_n$, where the prize w_j is given to the j-th ranked agent overall (we call these prizes *general* prizes), and the prize ω_j is given to the j-th ranked agent among the agents in the target group (we call these *group-specific* prizes), with ties broken uniformly.[2] We assume that $\sum_j w_j + \sum_j \omega_j \leq 1$, i.e., the contest designer has a unit budget. Given a vector of outputs $\boldsymbol{b}$, let the allocation of general prizes be given by $\boldsymbol{x}(\boldsymbol{b}) = (x_{i,j}(\boldsymbol{b}))_{i,j \in [n]}$, where $x_{i,j} = 1$ if agent i is awarded the j-th prize and $x_{i,j} = 0$ otherwise. Similarly, let $\boldsymbol{y}(\boldsymbol{b}, \boldsymbol{t}) = (y_{i,j}(\boldsymbol{b}, \boldsymbol{t}))_{i,j \in [n]}$ be the allocation of target group prizes, where $\boldsymbol{t} = (t_i)_{i \in [n]}$ is the group label vector; $t_i \in \{T, N\}$, $t_i = T$ if the agent is in the target group, $t_i = N$ if not. We shall suppress the notation for $\boldsymbol{t}$ and write $\boldsymbol{y}(\boldsymbol{b})$ instead of $\boldsymbol{y}(\boldsymbol{b}, \boldsymbol{t})$.

[2] In our model, it does not matter how we break ties, so w.l.o.g., we can assume uniform tie breaking. In more detail, for a tie to happen, two players must produce exactly identical output. In our setting, this happens with zero probability, as there are no point masses in the probability distributions F and G, and the distributions of the output generated by the agents (that we derive as a result of our equilibrium analysis) also do not have any point masses.

If agent i is in the target group, her utility is given by

$$
\begin{aligned}
u_T(v_i, \boldsymbol{b}, \boldsymbol{t}) &= \sum_{j \in [n]} w_j x_{i,j}(\boldsymbol{b}) + \sum_{j \in [n]} \omega_j y_{i,j}(\boldsymbol{b}) - b_i / v_i \\
&\equiv v_i \left(\sum_{j \in [n]} w_j x_{i,j}(\boldsymbol{b}) + \sum_{j \in [n]} \omega_j y_{i,j}(\boldsymbol{b}) \right) - b_i, \qquad (6.1)
\end{aligned}
$$

where the equivalence above is true because scaling the utility function by a constant does not change the strategy of an agent. In other words, by scaling the utility of an agent by a positive constant, we change the utility of any action for that agent by the same factor, so the strategy of the agent does not change. The utility of an agent i in the non-target group is given by

$$
u_N(v_i, \boldsymbol{b}, \boldsymbol{t}) = v_i \sum_{j \in [n]} w_j x_{i,j}(\boldsymbol{b}) - b_i. \qquad (6.2)
$$

6.2.1 Mathematical Preliminaries

Let $p_j^H(v)$ denote the probability that a value $v \in [0,1]$ is the j-th highest among n i.i.d. samples from a distribution H, given by the expression

$$
p_j^H(v) = \binom{n-1}{j-1} H(v)^{n-j} (1 - H(v))^{j-1}.
$$

A key role in the Bayes–Nash equilibrium of rank-order allocation contests is played by the order statistics. Let $f_{n,j}^H$ be the PDF of the j-th highest order statistic out of n i.i.d. samples from H (with PDF h), given by the expression

$$
f_{n,j}^H(v) = \frac{n!}{(j-1)!(n-j)!} H(v)^{n-j} (1 - H(v))^{j-1} h(v). \qquad (6.3)
$$

We shall also frequently use the following two identities for order statistics:

$$
\begin{aligned}
H(v) f_{n-1,j}^H(v) &= \frac{n-j}{n} f_{n,j}^H(v); \\
(1 - H(v)) f_{n-1,j}^H(v) &= \frac{j}{n} f_{n,j+1}^H(v).
\end{aligned}
$$

Definition 6.1 ((Weak) First-Order Stochastic (FOS) Dominance)**.** A distribution F *FOS dominates* another distribution G, if for every x, F gives at least as high a probability of receiving at least x as does G: $1 - F(x) \geq 1 - G(x) \iff F(x) \leq G(x)$.

Definition 6.2 ((Weak) Second-Order Stochastic (SOS) Dominance)**.** A distribution F *SOS dominates* another distribution G, if for every x: $\int_{-\infty}^{x} [F(t) - G(t)] dt \leq 0$.

FOS dominance implies SOS dominance. Another sufficient condition for SOS dominance: F SOS dominates G if G is a mean-preserving spread of F.

Let us denote the inner product of two functions $\psi, \phi : [0,1] \to \mathbb{R}_+$ by

$$\langle \psi, \phi \rangle = \int_0^1 \psi(x)\phi(x)dx.$$

When appropriately normalized, $\langle \psi, \phi \rangle$ measures the similarity between the functions ψ and ϕ.

6.2.2 Equilibrium and Objective Function

We shall study the Bayes–Nash equilibrium of the agents, where the agents in the target group use a symmetric strategy $\alpha(v)$ (symmetric across the agents in the group), where $\alpha(v)$ is the output of an agent with ability v; similarly, the agents in the non-target group use a symmetric strategy $\beta(v)$. The objective of the contest designer is to maximize the expected total output generated by the agents in the target group:

$$\sum_{i \in [n]} \mathbb{E}_{v_i \sim F}[\alpha(v_i)] \, \mathbb{P}[t_i = T] = n \cdot \mu \cdot \mathbb{E}_{v_i \sim F}[\alpha(v_i)],$$

which is equivalent to maximizing $\mathbb{E}_{v \sim F}[\alpha(v)]$ because n and μ are (fixed) parameters given to the model.

6.3 Equilibrium Analysis

In this section, we characterize a symmetric Bayes–Nash equilibrium (symmetric across the agents in a given group). We first solve for the equilibrium of the contest with both general and target group-specific prizes, then we specialize it for the two special cases: (1) only general prizes; (2) only group-specific prizes.

6.3.1 Both Target Group and General Prizes

In the next theorem, we focus on contests that have a sequence of general prizes $(w_j)_{j \in [n]}$ and a sequence of target group-specific prizes $(\omega_j)_{j \in [n]}$. We will characterize a Bayes–Nash equilibrium where a player in the target group plays an action $\alpha(v)$ as a function of her ability $v \sim F$ and a player in the non-target group plays an action $\beta(v)$ as a function of her ability $v \sim G$. Note that all the players in the target group use the same strategy α and the players in the non-target group use the same strategy β.

Theorem 6.1. *A contest with general prizes $(w_j)_{j\in[n]}$ and target group-specific prizes $(\omega_j)_{j\in[n]}$ has a Bayes–Nash equilibrium where an agent with ability $v \in [0,1]$ uses strategy $\alpha(v)$ if she is in the target group and strategy $\beta(v)$ if she is in the non-target group, where $\alpha(v)$ and $\beta(v)$ are defined as:*

$$\beta(v) = \int_0^v (\mu f(k(y))k'(y) + (1-\mu)g(y))A(y)ydy,$$
$$\alpha(v) = \begin{cases} \beta(k^{-1}(v)), & 0 \le v \le k(1) \\ \beta(1) + \int_{k(1)}^v \mu f(y)C(y)ydy, & k(1) \le v \le 1 \end{cases},$$

where:

- *k is a function defined over $[0,1]$ and is the solution to the following ordinary differential equation, with boundary condition $k(0) = 0$,*

$$k'(v) = \frac{(1-\mu)g(v)(v-k(v))A(v)}{\mu f(k(v))(k(v)A(v) + k(v)B(v) - vA(v))}.$$

- $A(v) = \sum_{j\in[n]} w_j\psi_j(\mu F(k(v)) + (1-\mu)G(v))$.
- $B(v) = \sum_{j\in[n]} \omega_j\psi_j(\mu F(k(v)) + (1-\mu))$.
- $C(v) = \sum_{j\in[n]}(w_j + \omega_j)\psi_j(\mu F(v) + (1-\mu))$.
- $\psi_j(x) = \binom{n-1}{j-1}x^{n-j-1}(1-x)^{j-2}((n-j) - (n-1)x)$.

Notice that Theorem 6.1 does not provide a closed-form solution for $\alpha(v)$ and $\beta(v)$, it rather provides a characterization using the function $k(v)$ that relates $\alpha(v)$ and $\beta(v)$. Standard equilibrium analysis techniques using first-order equilibrium conditions will give a system of differential equations in $\alpha(v)$ and $\beta(v)$. Using the function $k(v) = \alpha^{-1}(\beta(v))$, we convert this system of differential equations in $\alpha(v)$ and $\beta(v)$ to a comparatively simpler single first-order explicit ordinary differential equation in only one dependent variable $k(v)$. After solving the ordinary differential equation for $k(v)$, we can use $k(v)$ to compute $\alpha(v)$ and $\beta(v)$ as described in the theorem statement. This technique of using the function $k(v)$ is motivated by the work of Amann and Leininger [5], where they use to it analyze two-agent asymmetric contests.

If the distributions F or G, or the prizes $(w_j)_{j\in[n]}$ or $(\omega_j)_{j\in[n]}$, are non-trivial, then it may be intractable to calculate the analytical solutions for $\alpha(v)$ and $\beta(v)$ derived in Theorem 6.1, but numerical solutions may be computed. Example 6.8 in Section 6.7 illustrates the use of Theorem 6.1 to compute $\alpha(v)$ and $\beta(v)$ for a particular instantiation of the problem.

Theorem 6.1 characterizes an equilibrium assuming that all players in a given group use the same strategy ($\alpha(v)$ or $\beta(v)$), and that $\alpha(v)$ and $\beta(v)$ are strictly increasing and (almost everywhere) differentiable functions of v. We believe that these assumptions are reasonable and the equilibrium is the most natural one for this model. It remains open to prove the uniqueness of this equilibrium or to characterize the set of all possible equilibria.

6.3.2 General Prizes Only

A contest designer may want to only award prizes that are equally accessible to all the agents. This model is particularly appealing for several real-life applications, because, although the design of the prize structure may incorporate distributional information about the ability of the agents in the two groups, the contest itself is completely unbiased.

In our model, this design choice corresponds to setting $\omega_j = 0$ for all $j \in [n]$. For this case, the expected utility for a target group agent and a non-target group agent for a given pair of ability v and output b is the same. We can specialize the equilibrium characterization in Theorem 6.1 to this case as follows:

Theorem 6.2. *A contest with only general prizes $(w_j)_{j\in[n]}$ has a Bayes–Nash equilibrium where an agent with ability $v \in [0,1]$ in the target or non-target group uses the strategy $\alpha(v)$, where $\alpha(v)$ is defined as:*

$$\alpha(v) = \sum_{j\in[n-1]} (w_j - w_{j+1}) \int_0^v y f_{n-1,j}^{\mu F+(1-\mu)G}(y)dy.$$

This equilibrium is unique, as follows from a result by Chawla and Hartline [28]. They prove that if a contest is anonymous (does not discriminate among the agents) and the abilities of the agents are sampled i.i.d. from a distribution, then the equilibrium is unique. With only general prizes, the contest is (i) anonymous because the general prizes are awarded without any bias towards players from any group, and (ii) effectively, the abilities of the agents are sampled i.i.d. from the distribution $\mu F(v) + (1-\mu)G(v)$. Also, given observation (ii), Theorem 6.2 can alternatively be derived using the characterization for the i.i.d. setting by Moldovanu and Sela [81]. Please see Section 5.3 for a detailed discussion about the equilibrium characterization for the i.i.d. setting.

6.3.3 Group-Specific Prizes Only

In this case, there are no general prizes accessible to agents from both groups: all the prizes are allocated based on strict quotas for the groups. This model is similar to the model of seat quotas for college admissions studied by Bodoh-Creed and Hickman [21]. Technically, in this case we assume that $w_j = 0$ for all $j \in [n]$. Here we focus on the target group; a similar result also holds for the non-target group, if there are any prizes reserved for them.

Theorem 6.3. *A contest with target only group-specific prizes, $(\omega_j)_{j\in[n]}$, has a Bayes–Nash equilibrium where an agent with ability $v \in [0,1]$ in the target group uses strategy $\alpha(v)$, where $\alpha(v)$ is defined as:*

$$\alpha(v) = \sum_{j\in[n-1]} (\omega_j - \omega_{j+1}) \int_0^v y f_{n-1,j}^{\mu F+(1-\mu)}(y)dy.$$

Chawla and Hartline [28]'s uniqueness result also applies to this equilibrium characterization, as it did for only general prizes (Section 6.3.2). We can transform an instance of our problem (without changing its set of equilibria) to satisfy these requirements: the abilities of the players are picked i.i.d. from the distribution $\mu F(v)+(1-\mu)$ and the contest awards prizes without discriminating among the players. Again, as this case can be seen as an i.i.d. model with distribution $\mu F(v)+(1-\mu)$, Theorem 6.3 can alternatively be derived using the characterization by Moldovanu and Sela [81] (discussed in Section 5.3).

6.4 Designing Prizes to Incentivize the Target Group

Based on the equilibrium characterizations in the previous section, in this section we investigate the properties of the optimal contest that maximizes the total output of the target group. We characterize the optimal contests for only general prizes and for only group-specific prizes. Characterization of the optimal contest for both general and group-specific prizes is a non-trivial open problem because of the absence of a closed-form equilibrium characterization for the case (the characterization in Theorem 6.1 involves an ordinary differential equation).

In Section 6.4.3, we compare the choice between only general and only group-specific prizes; we observe that either of them may be a better choice depending upon the situation. Even if the distributions of the target and the non-target group are the same, the relative frequency of target group agents, measured by μ, determines which prize allocation scheme is better. If μ is sufficiently high, then it is better to have only group-specific prizes, while the converse is true if μ is very low.

6.4.1 General Prizes Only

The strategy used by the agents in the target group is $\alpha(v)$, as derived in Theorem 6.2. It can be rewritten as

$$\alpha(v) = \sum_{j\in[n-1]} (w_j - w_{j+1}) \int_0^v y f_{n-1,j}^{\mu F+(1-\mu)G}(y)dy = \sum_{j\in[n-1]} \gamma_j \frac{1}{j} \int_0^v y f_{n-1,j}^{\mu F+(1-\mu)G}(y)dy, \tag{6.4}$$

where $\gamma_j = j(w_j - w_{j+1})$. Instead of optimizing over $(w_j)_{j\in[n]}$, we can optimize over $(\gamma_j)_{j\in[n-1]}$ with constraints $\gamma_j \geq 0$ and $\sum_j \gamma_j = 1$, to find the optimal contest that maximizes the target group's output.

Theorem 6.4. *The contest with only general prizes that maximizes the expected total output of the agents in the target group awards a prize of $1/k^*$ to the k^* top-ranked agents and a prize of 0 to the remaining $n - k^*$ agents, where k^* is defined as*

$$k^* \in \arg\max_{j\in[n-1]} \langle \psi, f_{n,j+1}^{\mu F+(1-\mu)G} \rangle$$

for $\psi(x) = \frac{x(1-F(x))}{1-\mu F(x)-(1-\mu)G(x)}$.

Let $H(x)$ denote $\mu F(x)+(1-\mu)G(x)$. The function $\psi(x) = x\frac{1-F(x)}{1-H(x)}$ in Theorem 6.4 is independent of the prize structure: it depends only on μ, F, and G, which are parameters given to the contest designer. The function $f_{n,j+1}^H(x)$ is the PDF of the $(j+1)$-th highest order statistic of the distribution H. The optimal $j = k^*$ maximizes the inner product between ψ and $f_{n,j+1}^H$, and therefore, the similarity between them. In other words, the optimal $j = k^*$ selects the order statistic that is as similar to ψ as possible.

Note that when we are maximizing the total output, then instead of an inner product of the order statistic with $\psi(x)$ in Theorem 6.4, we take an inner product with just x. As x is monotonically increasing, the order statistic that is most similar to x and that maximizes the inner product with x is the one with $j = 1$, as shown by Moldovanu and Sela [81]. This provides an argument in favor of allocating the entire prize budget to the first prize. In contrast, when we focus on the target group only, it may be optimal to distribute the prize among several top-ranked agents. This also means that we may lose a portion of the total output, and this loss can be $\Omega(n)$ as shown in Example 6.3. Also, by the *single-crossing* property, see Vojnovic [106] Chapter 3, if we flatten the prize structure then we increase the output of the agents with low ability and decrease the output of the agents with high ability.

We have partially omitted calculations from the examples in this section. These examples are presented in Section 6.7 with more details.

Example 6.3. Let $F(x) = 1 - (1 - x)^{n-1}$, $G(x) = ((n - 1)x - F(x))/(n - 2)$, and $\mu = 1/(n - 1)$. We get $H(x) = \mu F(x) + (1 - \mu)G(x) = x$, which is the uniform distribution over $[0, 1]$. Note that $\psi(x) = \frac{x(1-F(x))}{1-H(x)} = x(1 - x)^{n-2}$.

It can be checked that the order-statistic that maximizes the inner-product with $\psi(x)$ is $f^H_{n,n-1}(x) = n(n - 1)x(1 - x)^{n-1}$. So, the optimal value is $k^* = n - 2$. The total output generated for $k^* = n - 2$ is $\frac{2}{n(n+1)}$. On the other hand, we know that the expected total output is maximized by $j = 1$ and is equal to $\frac{(n-1)}{n(n+1)}$. The ratio between the two quantities is $(n - 1)/2 = \Omega(n)$.

In the above example, the distribution of the target group F was first-order stochastic (FOS) dominated by the distribution of the general population H. For this case, we observed that awarding prizes to more than one agent increases the expected output of the target group. F being FOS dominated by H means that the agents in the target group have lower ability than the general population. As flattening the prize structure increases the output of the lower ability agents and decreases the output of the higher ability agents, it makes sense that having a flatter prize structure incentivizes them.

The above result raises the question: What if the target group is equally able or stronger than the general population, i.e., if the distribution of the ability of an agent in the target group, F, FOS dominates the distribution of the ability of an agent in the non-target group, G; is it then optimal to award the prize to the top-ranked agent only? When $F = G = H$, we know that it is optimal to give prize to the top-ranked agent only, so when F FOS dominates H, one might expect this to be the case as well. However, the example below illustrates that this is not true.[3]

[3]In the preliminaries we assumed F to be continuous, strictly increasing, and differentiable in $[0, 1]$. In the following and the subsequent examples, the distribution F is continuous and weakly increasing but may not be strictly increasing and differentiable everywhere. If F is not strictly increasing, we can add a slight gradient to resolve the issue; on the other hand, we can smooth out the finite number of points where it is not differentiable. These changes will have a minimal effect on the objective value and our analysis holds. For ease of presentation, we shall not discuss these issues.

Example 6.4. Let $\mu = 1/8$, $F(x) = 1 - S(x)$, $G(x) = (x - \mu F(x))/(1 - \mu)$, where $S(x)$ is given by:

$$S(x) = \begin{cases} 1, & \text{if } x < 3/4 \\ 16(1-x)^2, & \text{if } 3/4 \leq x < 15/16 \\ 1 - x, & \text{if } x \geq 15/16 \end{cases}.$$

It can be verified that $F(x)$ and $G(x)$ are continuous cumulative distribution functions. The distribution of the general population is $H(x) = \mu F(x) + (1 - \mu)G(x) = x$. Also, $F(x) \leq H(x)$ for every x (and strict for some values of x), and therefore, F FOS dominates H. The optimal number of prizes k^* need not be 1; for example, for $n = 50$, $k^* = 11$.

Now we consider the case where F and H have the same mean but different variance, which implies second-order stochastic (SOS) dominance. The distribution with lower variance SOS dominates the one with higher variance, assuming both distributions have the same mean. The results are similar to FOS dominance: it may be optimal to have multiple prizes irrespective of whether F or H has a higher variance, as shown in the following examples. Note that FOS dominance implies SOS dominance, but not vice-versa.

Example 6.5. Let $\mu = 2/3$, $F(x) = 3x^2 - 2x^3$, and $G(x) = 3x - 6x^2 + 4x^3$. Observe that $H(x) = \mu F(x) + (1 - \mu)G(x) = x$, and the variance of H is $1/12$ while the variance of F is $1/20$, and both have mean $1/2$.

Solving for k^*, we get $k^* = \frac{5n}{6} - \frac{\sqrt{7n^2+30n+39}}{6} + \frac{1}{2}$. For example, for $n = 50$, $k^* = 19$.

Example 6.6. Let $\mu = 1/4$, $F(x) = 1 - S(x)$, and $G(x) = (x - \mu F(x))/(1 - \mu)$, where $S(x)$ is given by:

$$S(x) = \begin{cases} 1 - 48x/31, & \text{if } x < 31/96 \\ 1/2, & \text{if } 31/96 \leq x < 3/4 \\ 8(1-x)^2, & \text{if } 3/4 \leq x < 7/8 \\ 1 - x, & \text{if } x \geq 7/8 \end{cases}$$

It can be checked that $F(x)$ and $G(x)$ are valid and for the distribution of the general population we have $H(x) = \mu F(x) + (1 - \mu)G(x) = x$. It can also be checked that the mean of all the distributions is $1/2$, and the variance of H is $1/12 \approx 0.083$ while the variance of F is $6703/55296 \approx 0.121 > 1/12$. The optimal number of prizes k^* need not be 1; for example, for $n = 50$, $k^* = 11$.

6.4.2 Group-Specific Prizes Only

In this section, we study the optimal prize structure when the prizes are accessible to the target group only. This models situations when there are separate contests for the target and the non-target group (there can be a separate contest for only non-target group agents, this will not have any affect on the strategy of the target group agents). We see that the optimal contest allocates the entire prize budget to the first prize, i.e., gives a prize of 1 to the top-ranked agent in the group (Theorem 6.5). The proof of Theorem 6.5 extends the techniques used to provide a similar result for total output [81].

Theorem 6.5. *The contest with only group-specific prizes that maximizes the output of the target group awards a prize of* 1 *to the top-ranked agent and* 0 *to others.*

6.4.3 Comparing Group-Specific and General Prizes

In this section, we compare the choice among only group-specific prizes and only general prizes. This choice depends upon the distributions F and G, and the parameter μ that captures the relative frequency of target-group agents. Even if we factor out the effect of the distributions F and G, then also the choice among group-specific and general prizes depends upon μ. If μ is high enough, then it is better to have group-specific prizes. On the other hand, if μ is quite low, then having only group-specific prizes leads to very low competition, and therefore, a low output is produced by the agents, and using general prizes is better than group-specific prizes. Example 6.7 below illustrates this. It is an open direction to further study the choice between group-specific and general prizes for general classes of distributions F and G.

Example 6.7. (Repeated with detailed calculations in Example 6.13) Let $F(x) = G(x) = x$. As shown in Theorem 6.5, it is optimal to allocate the entire budget to the first prize for group-specific prizes. For $F = G$, a similar result holds for general prizes as well.

We compare the following three cases:

(A) Only general prizes of value 1; $w_1 = 1$, $w_j = 0$ for $j > 1$, $\omega_j = 0$ for $j \in [n]$.

(B) Only target group-specific prizes of value 1; $w_j = 0$ for $j \in [n]$, $\omega_1 = 1$, $\omega_j = 0$ for $j > 1$.

(C) Only target group-specific prizes of value μ; $w_j = 0$ for $j \in [n]$, $\omega_1 = \mu$, $\omega_j = 0$ for $j > 1$.

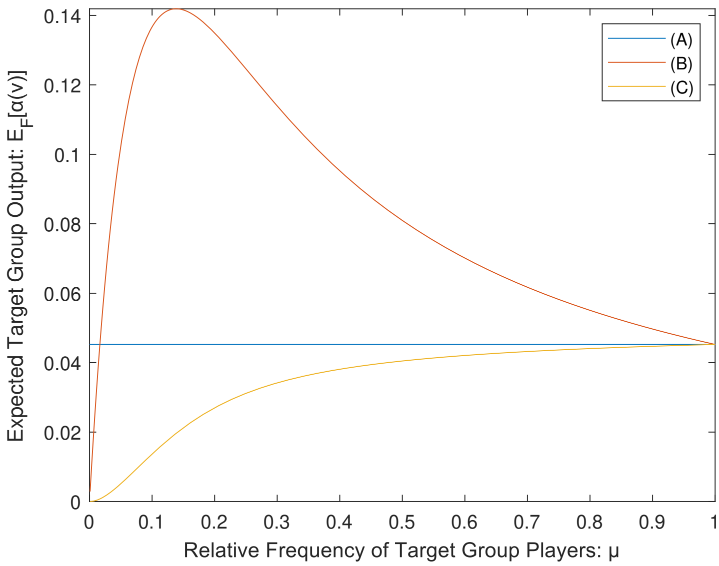

Figure 6.1: Expected Target Group Output vs Relative Frequency of Target Group Players. (A) Only general prizes of value 1. (B) Only target group-specific prizes of value 1. (C) Only target group-specific prizes of value μ.

As μ is the fraction of target group agents in the population, the contest organizer may want to have a contest for the target group agents with a prize of μ. This distributes the prize of 1 equally between the target group and the non-target group. Case (C) above captures this scenario.

Let $\alpha_A(v)$, $\alpha_B(v)$, and $\alpha_C(v)$ be the output generated by a player with ability v for the three cases (A), (B), and (C), respectively.

The expected output for case (A) general prizes is:

$$\mathbb{E}_{v\sim F}[\alpha_A(v)] = \frac{n-1}{n(n+1)}.$$

The expected output for case (B) group-specific prizes of value 1 is:

$$\mathbb{E}_{v\sim F}[\alpha_B(v)] = \int_0^1 v(1-F(v))f_{n-1,1}^{1-\mu+\mu F}(v)dv$$
$$= \frac{1}{n(n+1)\mu^2}((n-1)-(1-\mu)(n+1)+(1-\mu)^n(2(1-\mu)+(n+1)\mu)).$$

Although the solution above may seem difficult to interpret, but one can observe that as $\mu \to 0$, it goes to 0, but as $\mu \to 1$, it reaches $(n-1)/(n(n+1))$ from above.

We can observe that the expected output for case (C) group-specific prizes of value

μ is equal to μ times the expected output for case (B):

$$\begin{aligned}\mathbb{E}_{v\sim F}[\alpha_C(v)] &= \mu\,\mathbb{E}_{v\sim F}[\alpha_B(v)]\\ &= \frac{1}{n(n+1)\mu}((n-1)-(1-\mu)(n+1)+(1-\mu)^n(2(1-\mu)+(n+1)\mu)).\end{aligned}$$

The derivative of the above expression is non-negative in μ and so it is increasing, for $\mu \in [0,1]$. For $\mu \to 0$, it goes to 0, and for $\mu \to 1$, it goes to $(n-1)/(n(n+1))$. So, $\mathbb{E}_{v\sim F}[\alpha_C(v)] \leq (n-1)/(n(n+1))$ for $\mu \in [0,1]$.

Figure 6.1 plots the expected output for the three cases for $n = 20$.

We observe that case (A) always dominates case (C), so a general prize of 1 is better than a group-specific prize of μ, although in both cases the amount of prize in the contest per target group agent is the same. This happens because of higher competition in case (A): for general prizes, the target group agents produce higher output to compete with the non-target group agents.

Comparing case (A) general prize of 1 and (B) group-specific prize of 1, we see that if μ is very small, then case (A) dominates case (B), and the opposite is true if μ is sufficiently large. This happens because of too little competition in case (B) for very small μ.

6.5 Conclusion

We studied rank-order allocation of prizes that aim to maximize the output of a target group. We provide equilibrium characterization for these contests and study the properties of the optimal contest. For unbiased contests (i.e., with general prizes only), although it is preferable to award the prize to the top-performing player if the target group players are *similar* to the overall population, if the target group players are very *different*, then we may want to flatten out the prize structure. This does not only mean that we flatten out the prizes if the target group players are *weaker* than others; we also demonstrate that if the target group players are *stronger* but quite different from others, it may still be a good idea to flatten out the prizes. However, for contests with fixed prize quotas (i.e., with group-specific prizes only), it may not be useful to flatten the prizes. Comparing unbiased contests and fixed quota contests, the size of the target group in proportion to the size of the population plays an important role.

An important open problem is to understand the properties of the equilibrium and the optimal contest when there are both group-specific and general prizes. We characterize the equilibrium for this case using an ordinary differential equation, but there is further scope to understand the properties of this equilibrium. On the other hand, understanding the optimal contest for this case is largely open; we only characterize the optimal contest for specific sub-cases. One way to make this tractable may be to assume large numbers of agents and prizes; such assumptions allow for a stronger and more tractable equilibrium characterization by making the influence of an individual agent infinitesimal [87]. Our equilibrium and optimal contest characterization results for only group-specific or only general prizes can be extended with minimal changes to accommodate multiple groups (instead of just two) and objective functions that maximizes the weighted sum of the group outputs. These extensions are direct from the analysis. On the other hand, when there are both group-specific and general prizes, extension to multiple groups is beyond the scope of our current techniques, and is an important direction of future research. Another direction is to better compare the choice between only group-specific and only general prizes by providing more general conditions on μ, F, and G. Our work focused on contests with a rank-order allocation of prizes with incomplete information and linear cost functions. Instead of rank-order allocation of prizes, we can study a proportional allocation of prizes (Tullock contests) and its generalizations; one motivation for proportional allocation is the randomness and unpredictability of outcome in the real world. Similarly, we may relax the assumption of linear cost functions, or we can study complete information models instead of incomplete information.

Ethical Statement There may be both short-term and long-term implications for designing contests specifically to incentivize a target group. For example, even if the contest itself is unbiased (Section 6.4.1), optimizing the expected total output of the players in the target group may mean that we decrease the expected total output of the entire population, and in particular decrease the expected total output of the non-target group. Moreover, such effects can be stronger if the target group forms a smaller fraction of the participants. Our analysis and examples try to capture such effects. For any given contest, this means that we are optimizing for the target group with a cost to society. But for a longer time frame, the answer is less obvious, and we believe that it is out of the scope of our work. This question concerns whether practices such as affirmative action have an overall positive or negative impact on the society in the long run. Our work addresses the following question: if we do want to use affirmative action in contests to incentivize the target group (with particular design choices, such as the prize allocation scheme being unbiased), how do we do it most effectively, and how does it affect the output of all the players in the given contest.

6.6 Omitted Proofs

6.6.1 From Section 6.3

Proof of Theorem 6.1. Let us assume that $\alpha : [0, 1] \to \mathbb{R}_+$ and $\beta : [0, 1] \to \mathbb{R}_+$ are strictly increasing and (almost everywhere) differentiable functions and are the strategies used by the agents in the target group and the non-target group, respectively. We also assume that $\alpha(0) = \beta(0) = 0$, i.e., an agent with no ability doesn't produce any output, and that $\alpha(v) > 0$ and $\beta(v) > 0$ for $v > 0$, i.e., an agent with a non-zero ability produces a non-zero output. Let $\overline{\alpha} = \alpha(1)$ and $\overline{\beta} = \beta(1)$; we shall observe later in the proof that $\alpha(x) \geq \beta(x)$ for all x, and therefore, $\overline{\alpha} \geq \overline{\beta}$. Let the distribution of the output produced by an agent in the target group be denoted by $\widehat{F}$; $\widehat{F}$ and the corresponding PDF $\widehat{f}$ are given by:

$$\widehat{F}(x) = F(\alpha^{-1}(x)); \quad \widehat{f}(x) = f(\alpha^{-1}(x))\frac{d\alpha^{-1}(x)}{dx},$$

where $x \in [0, \overline{\alpha}]$. Similarly, the distribution and the PDF of the output produced by a non-target group agent are:

$$\widehat{G}(x) = G(\beta^{-1}(x)); \quad \widehat{g}(x) = g(\beta^{-1}(x))\frac{d\beta^{-1}(x)}{dx},$$

where $x \in [0, \overline{\beta}]$. Any given agent belongs to the target group with a probability μ, so the overall distribution of the output produced by an agent is given by the following mixture distribution:

$$\widehat{H}(x) = \mu\widehat{F}(x) + (1-\mu)\widehat{G}(x) = \begin{cases} \mu F(\alpha^{-1}(x)) + (1-\mu)G(\beta^{-1}(x)), & 0 \leq x \leq \overline{\beta}, \\ \mu F(\alpha^{-1}(x)) + (1-\mu), & \overline{\beta} \leq x \leq \overline{\alpha}. \end{cases}$$

For the general prizes $(w_j)_{j\in[n]}$, which are accessible to all agents, the probability that an agent who produces an output of x receives the j-th prize is given by $p_j^{\widehat{H}}(x)$, which is equal to

$$p_j^{\widehat{H}}(x) = \binom{n-1}{j-1}\widehat{H}(x)^{n-j}(1-\widehat{H}(x))^{j-1}.$$

The target group-specific prizes $(\omega_j)_{j\in[n]}$ are allocated to only the agents in the target group. For these prizes, we can effectively assume that all the agents in the non-target group produce < 0 output, and then find the probability that an agent in the target group has the j-th highest output. This output distribution, denoted by $\tilde{H}$, is given by:

$$\tilde{H}(x) = \mu\widehat{F}(x) + (1-\mu) = \mu F(\alpha^{-1}(x)) + (1-\mu),$$

and the probability that the x is j-th highest value is given by $p_j^{\tilde{H}}(x)$. Note that $\tilde{H}$ is not continuous and (its density function) has a point-mass of $(1-\mu)$ at 0.

First-order conditions. We now derive the first-order equilibrium conditions. Let us focus on a particular agent i. Let us assume that the agent has an ability v and produces an output x. All the other agents are playing with strategy α or β depending upon their group. If agent i is in the target group, his expected utility is:

$$u_T(v,x) = v\left(\sum_{j\in[n]} w_j p_j^{\widehat{H}}(x) + \sum_{j\in[n]} \omega_j p_j^{\tilde{H}}(x)\right) - x, \tag{6.5}$$

while, if the agent is not in the target group, then his expected utility is:

$$u_N(v,x) = v\sum_{j\in[n]} w_j p_j^{\widehat{H}}(x) - x. \tag{6.6}$$

For a target group agent, the first-order condition implies that the derivative of $u_T(v,x)$ with respect to x, at $x = \alpha(v)$, is: 0 if $x > 0$, and ≤ 0 if $x = 0$. For $x > 0$, we have:

$$\left.\frac{du_T(v,x)}{dx}\right|_{v=\alpha^{-1}(x)} = 0 \implies \alpha^{-1}(x)\left(\sum_{j\in[n]} w_j \frac{dp_j^{\widehat{H}}(x)}{dx} + \sum_{j\in[n]} \omega_j \frac{dp_j^{\tilde{H}}(x)}{dx}\right) = 1. \tag{6.7}$$

Similarly, for non-target group agent we have:

$$\left.\frac{du_N(v,x)}{dx}\right|_{v=\beta^{-1}(x)} = 0 \implies \beta^{-1}(x)\sum_{j\in[n]} w_j \frac{dp_j^{\widehat{H}}(x)}{dx} = 1. \tag{6.8}$$

From the two equations above, we can observe that for any x, $\alpha^{-1}(x) \leq \beta^{-1}(x)$, because the prizes w_j's and ω_j's and the derivatives of $p^{\widehat{H}}(x)$ and $p^{\tilde{H}}(x)$ are non-negative. $\alpha^{-1}(x) \leq \beta^{-1}(x)$ implies $\alpha(v) \geq \beta(v)$ for any given v, so an agent in the target group produces more output compared to an agent in the non-target group with the same ability. We also get that $\alpha^{-1}(\beta(v)) \leq \beta^{-1}(\beta(v)) = v$.

Second-order conditions. Note that equation (6.7) is valid for all values of $v > 0$ and $x = \alpha(v)$. Let $\xi(x)$ denote

$$\xi(x) = \left(\sum_{j\in[n]} w_j \frac{dp_j^{\widehat{H}}(x)}{dx} + \sum_{j\in[n]} \omega_j \frac{dp_j^{\tilde{H}}(x)}{dx}\right)$$

Select any values v_1, v, v_2 such that $v_1 < v < v_2$. Evaluating equation (6.7) at v_1 and $x_1 = \alpha(v_1)$, we get

$$\left.\frac{du_T(v_1, x)}{dx}\right|_{x=x_1} = v_1\xi(x_1) - 1 = 0.$$

Note that $\xi(x)$ non-negative for every x (can be deduced from the formula for $\xi(x)$, also done later in the proof). As $v > v_1$, replacing v_1 by v in the equation above, we have

$$v\xi(x_1) - 1 > v_1\xi(x_1) - 1 = 0 = v\xi(x) - 1.$$

So, $v(\xi(x_1) - \xi(x)) > 0$ if $x_1 < x = \alpha(v)$. Similarly, let $x_2 = \alpha(v_2)$, and as $v < v_2$, we get

$$v\xi(x_2) - 1 < v_2\xi(x_2) - 1 = 0 = v\xi(x) - 1.$$

So, $v(\xi(x_2) - \xi(x)) < 0$ if $x_2 > x = \alpha(v)$. These two equations give us the second-order optimality condition. So, $x = \alpha(v)$ maximizes the utility of a target group agent compared to any other $x > 0$.

Now, let us focus our attention to $x = 0$. At $x = 0$, there is a discontinuity in the win probability, the probability of win is only right-continuous, because the non-target group players are not eligible for the target group-specific prizes, which happens with non-zero probability. But, right-continuity is enough to ensure that for any player with $v > 0$, $x = \alpha(v)$ (satisfying the first-order conditions at x) is a better strategy for agent v compared to any other $x \geq 0$. And, for a player with $v = 0$, $x = 0$ is trivially an optimal strategy.

Similar arguments can be applied for non-target group agents to show that $x = \beta(v)$ is an equilibrium if it satisfies the first-order conditions.

Equilibrium strategies. We now solve for the equilibrium strategies using the first-order conditions derived earlier. We separate into two cases, depending upon the value of x, the output.

Case 1: $x \leq \overline{\beta}$**.**

Let us define a function $k(v)$ over the domain $[0, 1]$ as $k(v) = \alpha^{-1}(\beta(v))$. Note that $k(\beta^{-1}(x)) = \alpha^{-1}(x)$, and $k(0) = 0$ as per our assumption that $\alpha(0) = \beta(0) = 0$ and $\alpha(v) > 0$ and $\beta(v) > 0$ for $v > 0$. For $x \leq \overline{\beta}$, replacing $\alpha^{-1}(x)$ by $k(\beta^{-1}(x))$ in $\widehat{H}(x)$, we get:

$$\widehat{H}(x) = \mu F(\alpha^{-1}(x)) + (1-\mu)G(\beta^{-1}(x)) = \mu F(k(\beta^{-1}(x))) + (1-\mu)G(\beta^{-1}(x)),$$

which shows that $\widehat{H}(x)$ can be written as function of k and β, without directly using α. The same is true also for functions based on $\widehat{H}$, like $p_j^{\widehat{H}}(x)$ and $\frac{dp_j^{\widehat{H}}(x)}{dx}$. Differentiating

$\widehat{H}(x)$ w.r.t. x we get:

$$\frac{d\widehat{H}(x)}{dx} = (\mu f(k(\beta^{-1}(x)))k'(\beta^{-1}(x)) + (1-\mu)g(\beta^{-1}(x)))\frac{d\beta^{-1}(x)}{dx}.$$

Now, we do a change of variable, replacing x by $\beta(v)$, we get:

$$\begin{aligned}\widehat{H}(\beta(v)) &= \mu F(k(v)) + (1-\mu)G(v);\\ \frac{d\widehat{H}(\beta(v))}{d\beta(v)} &= \frac{d\widehat{H}(\beta(v))}{dv}\frac{dv}{d\beta(v)} = (\mu f(k(v))k'(v) + (1-\mu)g(v))\frac{dv}{d\beta(v)}.\end{aligned}$$

Similarly, we can write (by differentiating $\tilde{H}$ w.r.t. x, for $x > 0$):

$$\begin{aligned}\tilde{H}(\beta(v)) &= (1-\mu) + \mu F(k(v));\\ \frac{d\tilde{H}(\beta(v))}{d\beta(v)} &= \frac{d\tilde{H}(\beta(v))}{dv}\frac{dv}{d\beta(v)} = \mu f(k(v))k'(v)\frac{dv}{d\beta(v)}.\end{aligned}$$

Observe from the equations above that $\widehat{H}(\beta(v))$ and $\tilde{H}(\beta(v))$ are functions of only v and $k(v)$, and $\frac{d\widehat{H}(\beta(v))}{dv}$ and $\frac{d\tilde{H}(\beta(v))}{dv}$ are functions of only v, $k(v)$, and $k'(v)$, and do not directly dependent upon $\beta(v)$.

Expanding $\frac{dp_j^{\widehat{H}}(x)}{dx}$ and changing variable x to $\beta(v)$, we get:

$$\begin{aligned}\frac{dp_j^{\widehat{H}}(x)}{dx} =& \binom{n-1}{j-1}\widehat{H}(x)^{n-j-1}(1-\widehat{H}(x))^{j-2}((n-j)(1-\widehat{H}(x)) - (j-1)\widehat{H}(x))\frac{d\widehat{H}(x)}{dx}\\ =& A_j(v)(\mu f(k(v))k'(v) + (1-\mu)g(v))\frac{dv}{d\beta(v)},\end{aligned}$$

where $A_j(v) = \psi_j(\widehat{H}(\beta(v)))$, where $\psi_j(x) = \binom{n-1}{j-1}x^{n-j-1}(1-x)^{j-2}((n-j)-(n-1)x)$. $A_j(v)$ is a function of v and $k(v)$. In a similar manner, we get:

$$\frac{dp_j^{\tilde{H}}(x)}{dx} = B_j(v)\mu f(k(v))k'(v)\frac{dv}{d\beta(v)},$$

where $B_j(v) = \psi_j(\tilde{H}(\beta(v)))$.

Replacing $\frac{dp_j^{\widehat{H}}(x)}{dx}$ and $\frac{dp_j^{\tilde{H}}(x)}{dx}$ in the first-order condition, equation (6.7), and changing x to $\beta(v)$, we get:

$$\alpha^{-1}(x)\left(\sum_{j\in[n]} w_j \frac{dp_j^{\widehat{H}}(x)}{dx} + \sum_{j\in[n]} \omega_j \frac{dp_j^{\tilde{H}}(x)}{dx}\right) = 1$$

$$\iff \alpha^{-1}(\beta(v))\left(\sum_{j\in[n]} w_j A_j(v)(\mu f(k(v))k'(v) + (1-\mu)g(v))\right.$$

$$\left. + \sum_{j\in[n]} \omega_j B_j(v)\mu f(k(v))k'(v)\right) = \frac{d\beta(v)}{dv}$$

$$\iff k(v)k'(v)\mu f(k(v))\left(\sum_{j\in[n]} w_j A_j(v) + \sum_{j\in[n]} \omega_j B_j(v)\right)$$

$$+ k(v)(1-\mu)g(v)\sum_{j\in[n]} w_j A_j(v) = \frac{d\beta(v)}{dv}.$$

Let $A(v) = \sum_{j\in[n]} w_j A_j(v)$ and $B(v) = \sum_{j\in[n]} \omega_j B_j(v)$; from equation above

$$k(v)k'(v)\mu f(k(v))(A(v) + B(v)) + k(v)(1-\mu)g(v)A(v) = \frac{d\beta(v)}{dv}. \tag{6.9}$$

In a similar manner, from first-order condition (6.8):

$$vk'(v)\mu f(k(v))A(v) + v(1-\mu)g(v)A(v) = \frac{d\beta(v)}{dv}. \tag{6.10}$$

Combining these two equations, we get:

$$k'(v) = \frac{(1-\mu)g(v)(v-k(v))A(v)}{\mu f(k(v))(k(v)A(v) + k(v)B(v) - vA(v))}. \tag{6.11}$$

The expression above is a explicit first-order ordinary differential equation and can be calculated analytically based on the distributions F and G. As per our assumptions on F and G, and with boundary condition for $k(0) = 0$, there exists a unique solution to the equation above. After calculating the solution for $k(v)$, we can calculate $\beta(v)$ by using equation (6.10):

$$\beta(v) = \int_0^v (\mu k'(y)f(k(y)) + (1-\mu)g(y))A(y)y dy. \tag{6.12}$$

Let us denote the maximum value that k attains by $\overline{k} = k(1)$. Going back to the definition of k, $k(v) = \alpha^{-1}(\beta(v)) \implies k^{-1}(v) = \beta^{-1}(\alpha(v))$ for $v \in [0, \overline{k}]$. So, after calculating k and β as shown above, we can calculate α as:

$$\alpha(v) = \beta(k^{-1}(v)), \text{ for } v \in [0, \overline{k}]. \tag{6.13}$$

Note that at $\overline{k}$, $\alpha(\overline{k}) = \beta(k^{-1}(\overline{k})) = \beta(1) = \overline{\beta}$, which is exactly the range we were working with: $0 \leq x \leq \overline{\beta}$.

Case 2: $\overline{\beta} \leq x \leq \overline{\alpha}$.

For $x \geq \overline{\beta}$, we have

$$\widehat{H}(x) = \mu F(\alpha^{-1}(x)) + (1-\mu) = \tilde{H}(x).$$

As $\widehat{H}(x) = \tilde{H}(x)$ for $x \geq \overline{\beta}$, we also have $p_j^{\widehat{H}}(x) = p_j^{\tilde{H}}(x)$. As we did earlier, we do a change of variable from x to $\alpha(v)$ to get

$$\widehat{H}(\alpha(v)) = \mu F(v) + (1-\mu); \ \frac{d\widehat{H}(\alpha(v))}{d\alpha(v)} = \mu f(v)\frac{dv}{d\alpha(v)}.$$

Differentiating $p_j^{\widehat{H}}(x)$ w.r.t. x and changing variable to $\alpha(v)$, we get

$$\begin{aligned}\frac{dp_j^{\widehat{H}}(x)}{dx} =& \binom{n-1}{j-1}\widehat{H}(x)^{n-j-1}(1-\widehat{H}(x))^{j-2}((n-j)(1-\widehat{H}(x)) - (j-1)\widehat{H}(x))\frac{d\widehat{H}(x)}{dx} \\ =& C_j(v)\mu f(v)\frac{dv}{d\alpha(v)},\end{aligned}$$

where $C_j(v) = \psi_j(\tilde{H}(\alpha(v)))$.

From the first order conditions, we get (note that $\widehat{H}(x) = \tilde{H}(x)$)

$$\begin{aligned}\alpha^{-1}(x)\left(\sum_{j\in[n]} w_j \frac{dp_j^{\widehat{H}}(x)}{dx} + \sum_{j\in[n]} \omega_j \frac{dp_j^{\tilde{H}}(x)}{dx}\right) = 1 \\ \iff v\sum_{j\in[n]}(w_j+\omega_j)C_j(v)\mu f(v) = \frac{d\alpha(v)}{dv}.\end{aligned}$$

Using this differential equation, along with the boundary condition $\alpha(\overline{k}) = \overline{\beta}$ that we derived earlier, we get

$$\alpha(v) = \overline{\beta} + \int_{\overline{k}}^{v} y\sum_{j\in[n]}(w_j+\omega_j)C_j(y)\mu f(y)dy. \tag{6.14}$$

□

Proof of Theorem 6.2. This result is a specific instantiation of Theorem 6.1. Note that the utility functions are the same for both target group and non-target group agents, $u_T(v,x) = u_N(v,x)$ as given in equations (6.5) and (6.6). With reference to Theorem 6.1 and its proof, we have $\alpha(v) = \beta(v)$ and $k(v) = \alpha^{-1}(\beta(v)) = v$, and therefore

$$\alpha(v) = \int_0^v y\sum_{j\in[n]} w_j\psi_j(\mu F(y) + (1-\mu)G(y))(\mu f(y) + (1-\mu)g(y))dy,$$

where $\psi_j(x) = \binom{n-1}{j-1} x^{n-j-1}(1-x)^{j-2}((n-j)-(n-1)x)$, as before.

Let $H(v) = \mu F(v) + (1-\mu)G(v)$. We have $\alpha(v) = \int_0^v y \sum_{j\in[n]} w_j \psi_j(H(y))h(y)dy$. Rewriting $\alpha(v)$ using the PDF for the order statistic of H, $f^H_{n,j}$ (formula given in equation (6.3)), we get

$$\alpha(v) = \sum_{j\in[n-1]} (w_j - w_{j+1}) \int_0^v y f^H_{n-1,j}(y)dy,$$

as required. □

Proof of Theorem 6.3. Substituting $w_j = 0$ for every $j \in [n]$ in Theorem 6.1, we get $\beta(v) = 0$, $k(v) = 0$, and $C(v) = \sum_{j\in[n]} \omega_j \psi_j(\mu F(v) + (1-\mu))$ for $v \in [0,1]$, and therefore,

$$\alpha(v) = \int_0^v \mu f(y)C(y)ydy.$$

Let $H(v) = \mu F(v) + (1-\mu)$. We have $\alpha(v) = \int_0^v y \sum_{j\in[n]} \omega_j \psi_j(H(y))h(y)dy$. Rewriting $\alpha(v)$ using the PDF for the order statistic of H, $f^H_{n,j}$, we have

$$\alpha(v) = \sum_{j\in[n-1]} (\omega_j - \omega_{j+1}) \int_0^v y f^H_{n-1,j}(y)dy,$$

as required. □

6.6.2 From Section 6.4

Proof of Theorem 6.4. The expected output of an agent in the target group is $\mathbb{E}_{v\sim F}[\alpha(v)]$, and maximizing this quantity maximizes the total output produced by all agents in the target group because μ and n are constants. Let $H(v) = \mu F(v) + (1-\mu)G(v)$.

$$\begin{aligned}
\mathbb{E}_{v\sim F}[\alpha(v)] = \int_0^1 \alpha(v)dF(v) &= \int_0^1 \left(\sum_{j\in[n-1]} \gamma_j \frac{1}{j} \int_0^v y f^H_{n-1,j}(y)dy \right) dF(v) \\
&= \sum_{j\in[n-1]} \gamma_j \frac{1}{j} \int_0^1 \left(\int_y^1 dF(v) \right) y f^H_{n-1,j}(y)dy \\
&= \sum_{j\in[n-1]} \gamma_j \frac{1}{j} \int_0^1 y(1-F(y)) f^H_{n-1,j}(y)dy.
\end{aligned}$$

As $\gamma_j \geq 0$ and $\sum_j \gamma_j = 1$, from the above expression we can observe that the optimal solution sets $\gamma_k > 0$ only if $\left(\frac{1}{j}\int_0^1 y(1-F(y))f^H_{n-1,j}(y)dy\right)$ is maximized by $j = k$.

Moreover, there is always an optimal solution where $\gamma_{k^*} = 1$ for some $k^* \in [n-1]$ and $\gamma_j = 0$ for $j \neq k^*$, where k^* is given by:

$$\begin{aligned}
k^* &\in \arg\max_{j\in[n-1]} \frac{1}{j}\int_0^1 y(1-F(y))f^H_{n-1,j}(y)dy \\
&= \arg\max_{j\in[n-1]} \frac{1}{n}\int_0^1 y\frac{1-F(y)}{1-H(y)}f^H_{n,j+1}(y)dy \\
&= \arg\max_{j\in[n-1]} \langle \psi, f^H_{n,j+1}\rangle,
\end{aligned}$$

where $\psi(x) = x\frac{1-F(x)}{1-H(x)}$. □

Proof of Theorem 6.5. From the equilibrium characterization in Theorem 6.3, we have

$$\alpha(v) = \sum_{j\in[n-1]} (\omega_j - \omega_{j+1}) \int_0^v y f^{\mu F+(1-\mu)}_{n-1,j}(y)dy = \sum_{j\in[n-1]} \gamma_j \frac{1}{j}\int_0^v y f^{\mu F+(1-\mu)}_{n-1,j}(y)dy,$$

where $\gamma_j = j(\omega_j - \omega_{j+1}) \geq 0$.

Let $H(x) = (1-\mu) + \mu F(x)$. We have $h(x) = \mu f(x)$ for $x > 0$ and $(1 - H(x)) = \mu(1-F(x))$. Let

$$H_{n,j}(x) = \sum_{\ell=0}^{j-1} \binom{n}{\ell} H(x)^{n-\ell}(1-H(x))^\ell$$

be the distribution of the j-th highest order statistic of $H(x)$, and let $h_{n,j}(x)$ be its PDF, which is well-defined for $x > 0$, but at $x = 0$, $H_{n,j}(0)$ has a non-zero point mass. Note that $f^{\mu F+(1-\mu)}_{n,j}(x) = h_{n,j}(x)$.

The objective of the designer is

$$\begin{aligned}
\mathbb{E}_{v\sim F}[\alpha(v)] &= \int_0^1 \left(\sum_{j\in[n-1]} \gamma_j \frac{1}{j}\int_0^v y h_{n-1,j}(y)dy \right) dF(v) \\
&= \sum_{j\in[n-1]} \gamma_j \frac{1}{j}\int_0^1 y(1-F(y))h_{n-1,j}(y)dy \\
&= \sum_{j\in[n-1]} \gamma_j \frac{1}{\mu j}\int_0^1 y(1-H(y))h_{n-1,j}(y)dy \\
&= \sum_{j\in[n-1]} \gamma_j \frac{1}{\mu n}\int_0^1 y h_{n,j+1}(y)dy.
\end{aligned}$$

As $1/(\mu n)$ is a constant, to maximize the objective, the designer sets $\gamma_j > 0$ only for values of j that maximize $\int_0^1 y h_{n,j+1}(y)dy$. From the definition of $H_{n,k}$, we have

$$1 = H_{n,j+1}(1) = H_{n,j+1}(0) + \int_0^1 h_{n,j+1}(y)dy \implies \int_0^1 h_{n,j+1}(y)dy = 1 - H_{n,j+1}(0).$$

We also know that for any x and $j < n$,

$$H_{n,j}(x) = \sum_{\ell=0}^{j-1} \binom{n}{\ell} H(x)^{n-\ell}(1-H(x))^{\ell} \leq \sum_{\ell=0}^{j} \binom{n}{\ell} H(x)^{n-\ell}(1-H(x))^{\ell} = H_{n,j+1}(x).$$

Using this, we get

$$\begin{aligned} H_{n,1+1}(x) \leq H_{n,j+1}(x) &\implies 1 - H_{n,2}(0) \geq 1 - H_{n,j+1}(0) \\ &\implies \int_0^1 h_{n,2}(y)dy \geq \int_0^1 h_{n,j+1}(y)dy, \quad (6.15) \end{aligned}$$

for any $j \in [n-1]$. Observe that $h_{n,2}(y)$ is either *single-crossing* with respect to $h_{n,j+1}(y)$, i.e., there is a point $x \in [0,1]$ s.t. $h_{n,2}(y) \leq h_{n,j+1}(y)$ for $y \leq x$ and $h_{n,2}(y) > f^H_{n,j+1}(y)$ for $y > x$, or $h_{n,2}(y)$ is always greater than $h_{n,j+1}(y)$, because

$$\frac{h_{n,2}(y)}{h_{n,j+1}(y)} = \frac{\binom{n-1}{1} H(y)^{n-2}(1-H(y))^1}{\binom{n-1}{j} H(y)^{n-j-1}(1-H(y))^j} = \frac{(n-j-1)!j!}{(n-2)!(1-H(y))^{j-1}} \left(\frac{H(y)}{1-H(y)}\right)^{j-1},$$

$\frac{H(y)}{1-H(y)}$ is strictly increasing and goes from $\frac{H(0)}{1-H(0)}$ to ∞ as y goes to 1, and the coefficient of $\left(\frac{H(y)}{1-H(y)}\right)^{j-1}$ is a constant. Using: (i) the identity function y is an increasing function, (ii) $h_{n,2}(y)$ is single-crossing w.r.t. $h_{n,j+1}(y)$, and (iii) equation (6.15), we have

$$\int_0^1 y h_{n,2}(y)dy \geq \int_0^1 y h_{n,j+1}(y)dy.$$

So, it is optimal to only have a first prize. □

6.7 Additional Examples

6.7.1 Equilibrium Analysis

The following example illustrates the use of Theorem 6.1 to compute $\alpha(v)$ and $\beta(v)$.

Example 6.8. Let $\mu = 1/2$, $F(v) = v^s$, and $G(v) = v^t$, for $s, t > 0$. Let $w_1 = 1/2$, $w_j = 0$ for $j > 1$, $\omega_1 = 1/2$, and $\omega_j = 0$ for $j > 1$.

We get $\mu F(k(v)) + (1-\mu)G(v) = (k(v)^s + v^t)/2$, $\mu F(k(v)) + (1-\mu) = (k(v)^s + 1)/2$, and $\mu F(v) + (1-\mu) = (v^s+1)/2$. Further, we get

$$\begin{aligned} A(v) &= \sum_{j\in[n]} w_j \psi_j(\mu F(k(v)) + (1-\mu)G(v)) = w_1\psi_1(\mu F(k(v)) + (1-\mu)G(v)) \\ &= \frac{1}{2}(n-1)\left(\frac{k(v)^s + v^t}{2}\right)^{n-2}. \end{aligned}$$

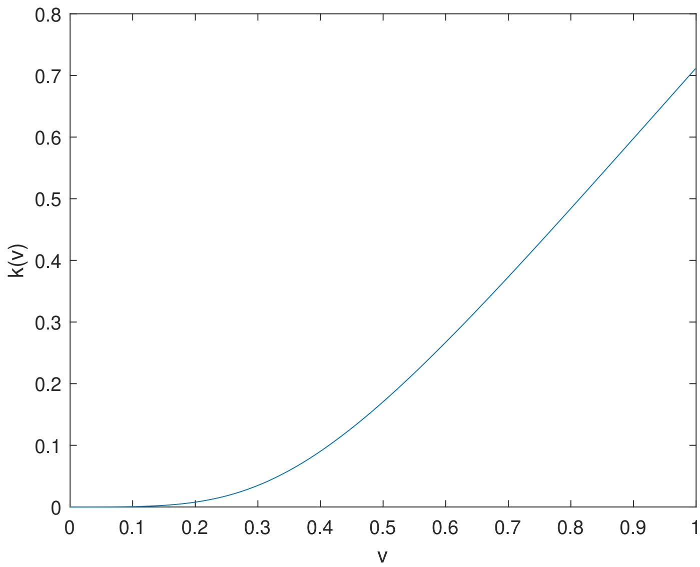

Figure 6.2: $k(v)$ vs v (Example 6.8).

Similarly, we get $B(v) = \frac{1}{2}(n-1)\left(\frac{k(v)^s+1}{2}\right)^{n-2}$ and $C(v) = (n-1)\left(\frac{v^s+1}{2}\right)^{n-2}$. Putting it all together, we get

$$k'(v) = \frac{tv^{t-1}(v-k(v))(k(v)^s+v^t)^{n-2}}{sk(v)^{s-1}(k(v)(k(v)^s+1)^{n-2}-(v-k(v))(k(v)^s+v^t)^{n-2})}.$$

Numerical solution for $k(v)$ for $n = 5$, $s = 1/2$, and $t = 1$ is given in Figure 6.2. Once we have $k(v)$, we compute $\beta(v)$ and $\alpha(v)$ using

$$\beta(v) = \int_0^v \frac{n-1}{2}\frac{sk(y)^{s-1}k'(y)+ty^{t-1}}{2}\left(\frac{k(y)^s+y^t}{2}\right)^{n-2} ydy,$$

$$\alpha(v) = \begin{cases} \beta(k^{-1}(v)), & 0 \le v \le k(1), \\ \beta(1) + \int_{k(1)}^v (n-1)\frac{sy^{s-1}}{2}\left(\frac{y^s+1}{2}\right)^{n-2} ydy, & k(1) \le v \le 1, \end{cases}$$

given in Figure 6.3. Further, we can use it to calculate the expected output of the players, for a player in the target group it is $\mathbb{E}_{v\sim F}[\alpha(v)] \approx 0.103$ and for a player in the non-target group player it is $\mathbb{E}_{v\sim G}[\beta(v)] \approx 0.052$.

6.7.2 Further Details on the Examples from Section 6.4.1

The examples from Section 6.4.1 have been repeated here with more details.

Example 6.9 (Example 6.3). Let $F(x) = 1-(1-x)^{n-1}$, $G(x) = ((n-1)x - F(x))/(n-2)$, and $\mu = 1/(n-1)$. We get $H(x) = \mu F(x) + (1-\mu)G(x) = x$, which

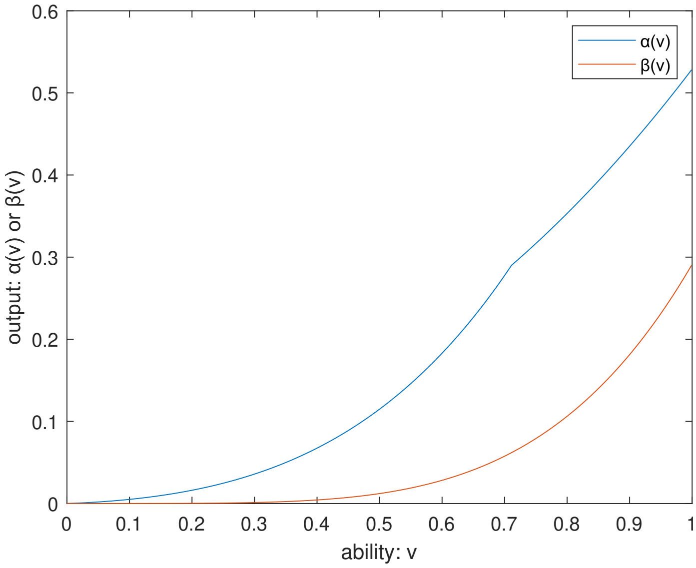

Figure 6.3: $\alpha(v)$ and $\beta(v)$ vs v (Example 6.8).

is the uniform distribution over $[0, 1]$. Note that

$$\psi(x) = \frac{x(1 - F(x))}{1 - H(x)} = \frac{x(1 - x)^{n-1}}{(1 - x)} = x(1 - x)^{n-2}.$$

It can be checked that the order-statistic that maximizes the inner-product with $\psi(x)$ is $f^H_{n,n-1}(x) = n(n - 1)x(1 - x)^{n-1}$. So, the optimal value is $k^* = n - 2$. The total output generated for $k^* = n - 2$ is

$$\frac{1}{n}\int_0^1 x f^H_{n,n-1}(x)dx = \frac{2}{n(n+1)}\int_0^1 f^H_{n+1,n-1}(x)dx = \frac{2}{n(n+1)}.$$

On the other hand, we know that the expected total output is maximized by $j = 1$ and is equal to

$$\frac{1}{n}\int_0^1 x f^H_{n,2}(x)dx = \frac{(n-1)}{n(n+1)}\int_0^1 f^H_{n+1,2}(x)dx = \frac{(n-1)}{n(n+1)}.$$

The ratio between the two quantities is $(n - 1)/2 = \Omega(n)$.

Example 6.10 (Example 6.4)**.** Let $\mu = 1/8$, $F(x) = 1 - S(x)$, $G(x) = (x - \mu F(x))/(1 - \mu)$, where $S(x)$ is given by:

$$S(x) = \begin{cases} 1, & \text{if } x < 3/4 \\ 16(1-x)^2, & \text{if } 3/4 \leq x < 15/16 \\ 1 - x, & \text{if } x \geq 15/16 \end{cases}.$$

It can be verified that $F(x)$ and $G(x)$ are continuous cumulative distribution functions. The distribution of the general population is $H(x) = \mu F(x) + (1-\mu)G(x) = x$. Also, $F(x) \leq H(x)$ for every x (and strict for some values of x), and therefore, F FOS dominates H. The objective value is:

$$\begin{aligned}
&\frac{1}{j}\left(\int_0^{\frac{3}{4}} x f^H_{n-1,j}(x)dx + \int_{\frac{3}{4}}^{\frac{15}{16}} 16x(1-x)^2 f^H_{n-1,j}(x)dx + \int_{\frac{15}{16}}^{1} x(1-x) f^H_{n-1,j}(x)dx\right)\\
&= \frac{n-j}{j \cdot n} F^H_{n,j}\left(\frac{3}{4}\right) + \frac{16(n-j)(j+1)}{n(n+1)(n+2)}(F^H_{n+2,j+2}(\frac{15}{16}) - F^H_{n+2,j+2}(\frac{3}{4}))\\
&\qquad + \frac{n-j}{n(n+1)}\left(1 - F^H_{n+1,j+1}\left(\frac{15}{16}\right)\right),
\end{aligned}$$

where $F^H_{n,j}$ is the CDF corresponding to the PDF $f^H_{n,j}$. Note that $F^H_{n,j}$ is equal to the Beta distribution with parameters $n+1-j$ and j. The optimal value for j may be derived with an involved calculation, but for our purpose, we just want to illustrate that the optimal j is not 1; it can be checked that for $n = 50$, the optimal value for j is 11 and the optimal objective value is approximately 0.0498.

Example 6.11 (Example 6.5)**.** Let $\mu = 2/3$, $F(x) = 3x^2 - 2x^3$, and $G(x) = 3x - 6x^2 + 4x^3$. Observe that $H(x) = \mu F(x) + (1-\mu)G(x) = x$, and the variance of H is $1/12$ while the variance of F is $1/20$, and both have mean $1/2$. The objective value is given by:

$$\begin{aligned}
&\frac{1}{j}\int_0^1 x(1 - 3x^2 + 2x^3) f^H_{n-1,j}(x)dx\\
&\qquad\qquad = \frac{n-j}{j \cdot n}\left(1 - \frac{(n+1-j)(n+2-j)}{(n+1)(n+2)}\left(3 - 2 \cdot \frac{n+3-j}{n+3}\right)\right).
\end{aligned}$$

Solving for j, we get $j = \frac{5n}{6} - \frac{\sqrt{7n^2+30n+39}}{6} + \frac{1}{2}$; as j is integer, either the floor or the ceiling of this number is optimal. For example, for $n = 50$, the optimal value is 19.

Example 6.12 (Example 6.6)**.** Let $\mu = 1/4$, $F(x) = 1 - S(x)$, and $G(x) = (x - \mu F(x))/(1-\mu)$, where $S(x)$ is given by:

$$S(x) = \begin{cases} 1 - 48x/31, & \text{if } x < 31/96 \\ 1/2, & \text{if } 31/96 \leq x < 3/4 \\ 8(1-x)^2, & \text{if } 3/4 \leq x < 7/8 \\ 1 - x, & \text{if } x \geq 7/8 \end{cases}$$

It can be checked that $F(x)$ and $G(x)$ are valid and for the distribution of the general population we have $H(x) = \mu F(x) + (1-\mu)G(x) = x$. It can also be checked that the mean of all the distributions is $1/2$, and the variance of H is $1/12 \approx 0.083$ while the variance of F is $6703/55296 \approx 0.121 > 1/12$. The objective value is:

$$\frac{1}{j}\left(\int_0^{31/96} x\left(1-\frac{48}{31}x\right) f_{n-1,j}^H(x)dx + \int_{31/96}^{3/4} x\cdot\frac{1}{2}\cdot f_{n-1,j}^H(x)dx\right.$$
$$\left. + \int_{3/4}^{7/8} x\cdot 8\cdot(1-x)^2 f_{n-1,j}^H(x)dx + \int_{7/8}^{1} x(1-x) f_{n-1,j}^H(x)dx\right)$$
$$= \frac{n-j}{2\cdot j\cdot n}\left(F_{n,j}^H(31/96) + F_{n,j}^H(3/4)\right) - \frac{48(n-j)(n+1-j)}{31\cdot j\cdot n\cdot(n+1)} F_{n+1,j}^H(31/96)$$
$$+ \frac{8(n-j)(j+1)}{n(n+1)(n+2)}(F_{n+2,j+2}^H(7/8) - F_{n+2,j+2}^H(3/4)) + \frac{n-j}{n(n+1)}\left(1 - F_{n+1,j+1}^H(7/8)\right).$$

Note that $F_{n,j}^H$ is the Beta distribution with parameters $n+1-j$ and j. For $n = 50$, the optimal value for j is 11 and the optimal objective value is approximately 0.0249.

6.7.3 Further Details on the Example from Section 6.4.3

Example 6.13 (Example 6.7)**.** Let $F(x) = G(x) = x$. As shown in Theorem 6.5, it is optimal to allocate the entire budget to the first prize for group-specific prizes. For $F = G$, a similar result holds for general prizes as well.

We compare the following three cases:

(A) Only general prizes of value 1; $w_1 = 1$, $w_j = 0$ for $j > 1$, $\omega_j = 0$ for $j \in [n]$.

(B) Only target group-specific prizes of value 1; $w_j = 0$ for $j \in [n]$, $\omega_1 = 1$, $\omega_j = 0$ for $j > 1$.

(C) Only target group-specific prizes of value μ; $w_j = 0$ for $j \in [n]$, $\omega_1 = \mu$, $\omega_j = 0$ for $j > 1$.

As μ is the fraction of target group agents in the population, the contest organizer may want to have a contest for the target group agents with a prize of μ. This distributes the prize of 1 equally between the target group and the non-target group. Case (C) above captures this scenario.

Let $\alpha_A(v)$, $\alpha_B(v)$, and $\alpha_C(v)$ be the output generated by a player with ability v for the three cases (A), (B), and (C), respectively.

The expected output for case (A) general prizes is:

$$\mathbb{E}_{v\sim F}[\alpha_A(v)] = \int_0^1 v(1-F(v)) f_{n-1,1}^F(v)dv = (n-1)\int_0^1 (1-v)v^{n-1}dv = \frac{n-1}{n(n+1)}.$$

The expected output for case (B) group-specific prizes of value 1 is:

$$\mathbb{E}_{v\sim F}[\alpha_B(v)] = \int_0^1 v(1-F(v))f_{n-1,1}^{1-\mu+\mu F}(v)dv$$
$$= (n-1)\int_0^1 v(1-v)(1-\mu+\mu v)^{n-2}\mu dv.$$

Let $H(x) = 1-\mu+\mu F(x) = 1-\mu+\mu x$. We have $(1-H(x)) = \mu(1-x)$ and $h(x) = \mu$. Let us write $f_{n,j}^H(x)$ as $h_{n,j}(x)$ and $F_{n,j}^H(x)$ as $H_{n,j}(x)$. Now,

$$v = \frac{1}{\mu}\mu v = \frac{1}{\mu}((1-\mu+\mu v)-(1-\mu)) = \frac{1}{\mu}(H(v)-(1-\mu)).$$

Incorporating this into the expected output formula, we get

$$\begin{aligned}
\mathbb{E}_{v\sim F}[\alpha_B(v)] &= (n-1)\int_0^1 v(1-v)(1-\mu+\mu v)^{n-2}\mu dv \\
&= \frac{(n-1)}{\mu^2}\int_0^1 (H(v)-(1-\mu))(1-H(v))H(v)^{n-2}h(v)dv \\
&= \frac{1}{\mu^2}\int_0^1 (H(v)-(1-\mu))(1-H(v))h_{n-1,1}(v)dv \\
&= \frac{1}{n\mu^2}\int_0^1 (H(v)-(1-\mu))h_{n,2}(v)dv \\
&= \frac{1}{n\mu^2}\left(\int_0^1 H(v)h_{n,2}(v)dv - (1-\mu)\int_0^1 h_{n,2}(v)dv\right) \\
&= \frac{1}{n\mu^2}\left(\frac{n-1}{n+1}\int_0^1 h_{n+1,2}(v)dv - (1-\mu)(1-H_{n,2}(0))\right) \\
&= \frac{1}{n\mu^2}\left(\frac{n-1}{n+1}(1-H_{n+1,2}(0)) - (1-\mu)(1-H_{n,2}(0))\right).
\end{aligned}$$

Now, $H_{n,2}(0) = H(0)^n + n(1-H(0))H(0)^{n-1} = (1-\mu)^n + n\mu(1-\mu)^{n-1} = (1-\mu)^{n-1}(1-\mu+n\mu)$. Similarly, $H_{n+1,2}(0) = (1-\mu)^n(1-\mu+(n+1)\mu)$. Replacing this, we get

$$\begin{aligned}
\mathbb{E}_{v\sim F}[\alpha_B(v)] = \frac{1}{n\mu^2}\Bigg(&\frac{n-1}{n+1}(1-(1-\mu)^n(1-\mu+(n+1)\mu)) \\
&- (1-\mu)(1-(1-\mu)^{n-1}(1-\mu+n\mu))\Bigg)
\end{aligned}$$
$$= \frac{1}{n(n+1)\mu^2}((n-1)-(1-\mu)(n+1)+(1-\mu)^n(2(1-\mu)+(n+1)\mu)).$$

Although the solution above may seem difficult to interpret, but one can observe that as $\mu \to 0$, it goes to 0, but as $\mu \to 1$, it reaches $(n-1)/(n(n+1))$ from above.

We can observe that the expected output for case (C) group-specific prizes of value μ is equal to μ times the expected output for case (B):

$$\begin{aligned}\mathbb{E}_{v\sim F}[\alpha_C(v)] &= \mu\,\mathbb{E}_{v\sim F}[\alpha_B(v)] \\ &= \frac{1}{n(n+1)\mu}((n-1)-(1-\mu)(n+1)+(1-\mu)^n(2(1-\mu)+(n+1)\mu)).\end{aligned}$$

The derivative of the above expression is non-negative in μ and so it is increasing, for $\mu \in [0,1]$. For $\mu \to 0$, it goes to 0, and for $\mu \to 1$, it goes to $(n-1)/(n(n+1))$. So, $\mathbb{E}_{v\sim F}[\alpha_C(v)] \leq (n-1)/(n(n+1))$ for $\mu \in [0,1]$.

Figure 6.1 plots the expected output for the three cases for $n = 20$.

We observe that case (A) always dominates case (C), so a general prize of 1 is better than a group-specific prize of μ, although in both cases the amount of prize in the contest per target group agent is the same. This happens because of higher competition in case (A): for general prizes, the target group agents produce higher output to compete with the non-target group agents.

Comparing case (A) general prize of 1 and (B) group-specific prize of 1, we see that if μ is very small, then case (A) dominates case (B), and the opposite is true if μ is sufficiently large. This happens because of too little competition in case (B) for very small μ.

Part III

Simultaneous Contests

Chapter 7

Equal-Sharing Allocation of Prizes

7.1 Introduction

In some contests, the prizes (rewards) may have monetary value, which incentivizes the players to make the costly investments. In other scenarios, the prizes may be associated with reputation and social status. For example, many online forums and websites depend upon user-generated content to provide value to their customers, and award badges, which do not have any monetary value but provide social reputation (e.g., StackOverflow and Quora); see, e.g., [62].

In the presence of multiple simultaneous contests, each player may explicitly select one or more contests and invest efforts so as to win the associated prizes. Moreover, sometimes contest participation is implicit: players engage in various activities, and each contest awards prizes to some of the players based on their performance in a specific subset of activities.

Consider, for instance, the setting where several social media platforms or news websites compete to attract customers. The potential customers are not homogeneous: e.g., some may be interested in politics, while others focus on sports or technology. It is therefore natural to model this setting as a set of simultaneous contests, with each individual contest corresponding to a group of customers with similar preferences. The platforms can take actions that make them more attractive to potential customers. Indeed, some of the actions, such as improving the interface, or increasing the update frequency, may impact the platform's performance with respect to several customer groups. That is, we can think of platforms as engaging in several *activities*, with their performance in each contest depending on the effort they invest in these activities. Different customer groups may value different mixtures of activities in different ways: e.g., while consumers of financial and sports news care about frequent updates, those who read the gossip column are happy with daily or even weekly updates. Thus, by increasing her investment in an activity, a player may improve her performance in several—but not all—contests.

In this chapter, we study a formal model that can capture scenarios of this type. In our model, there are several players and several simultaneous contests, as well as a set of activities. Each player selects their effort level for each activity (and may incur a cost for doing so, or face budget constraints), and each contest j has its own success function, which specifies combinations of effort for each activity that are sufficient to succeed in j. In addition to that, we assume that each contest allocates identical prizes to all agents that meet its criteria[1]: e.g., if, to win contest j, it suffices to produce 2 units along activity ℓ, then the player who produces 3 units along ℓ and the player who produces 30 units along ℓ receive the same prize from j; however, the value of the prize in contest j may depend on the number of players who meet the criteria of j. We refer to this setting as *multi-activity games*.

From the perspective of the player, we distinguish between two models: (1) the *cost model*, which has cost functions for producing output, and (2) the *budget model*, which has generalized budget constraints (feasible subsets of the space). The cost functions and the budget constraints may be different across players, capturing the fact that some players may be able to perform better in some activities compared to others.

Our model is very general: we impose very mild and natural constraints on the contests' success functions and prize allocation rules. In particular, we assume that the total value of the prizes allocated to the winners of each contest is a monotonically non-decreasing and concave function of the number of winners; these assumptions capture the idea that the value of an award decreases as it gets awarded to a larger number of players (see, e.g., [62]). A specific instantiation of this model is a contest that has a *fixed total prize* of V, and every winner gets a share of V/k if there are k winners.

Besides our social media platform example, several other situations can be modeled using contests with equal sharing allocation of prizes, where players compete by engaging in activities valued by multiple contests. For instance, funding agencies are generally not able to *perfectly discriminate* among the applicants to select the most deserving ones. They might score candidates based on attributes such as strength of

[1] The motivation for equal sharing allocation is somewhat similar to that of proportional allocation (e.g., in Tullock contests [104]), with equal sharing becoming relevant in situations where the contests use explicit rules (or criteria) to decide on an allocation.

proposal, publication history, and management experience, and allocate their budget to one of the eligible candidates or distribute it among them. Because of this uniform distribution of the funding, from the perspective of an expected utility maximizing applicant, this situation can be approximated using a contest with an equal sharing allocation with a fixed total prize.

7.1.1 Our Results

We observe that, under mild assumptions, multi-activity games are *exact potential games* [83] (see Definition 7.4). This guarantees the existence of a pure-strategy Nash equilibrium (PNE). Moreover, an approximate PNE can be computed in pseudo-polynomial time. In fact, a sequence of ϵ-better/best[2] response moves converges to an ϵ-PNE, with the number of steps proportional to $1/\epsilon$ (and pseudo-polynomial in other parameters). For the budget model, for the natural definition of social welfare in these games, we observe that the price of anarchy (PoA) is at most 2 and the price of stability (PoS) can be close to 2, so both these bounds are tight.[3] However, for the cost model, for meaningful definitions of social welfare, the PoS can be infinite.

We then study the computation complexity of finding an equilibrium in these games, which is the main focus of this chapter. This portion of the chapter concentrates on a specific instantiation of the model, where the contests' criteria and the players' budget constraints are all linear; let us call this model the *linear model*. (Here, we discuss the results for the budget model. With similar assumptions, similar results hold for the cost model as well.) We show that, for the linear model, it is NP-hard for a player to best respond to the strategies of other players. This hardness result holds even for a game with only one player; in other words, the hardness is due to the optimization problem that a player faces while playing the game. We also prove that there exists no polynomial-time approximation scheme unless P=NP. (Here, NP-hardness for best-response directly implies NP-hardness for equilibrium computation.) On the positive side, we obtain fixed-parameter tractability results: we show that best response is polynomial-time computable if either the number of contests or the number of activities is a constant.

The NP-hardness result for best-response motivates us to further restrict our model: we assume that a player can produce output only along a small (polynomial) number of portfolios of activities. Mathematically, a portfolio corresponds to

[2]An ϵ-better response move increases the utility of a player by at least ϵ.

[3]The price of anarchy (stability) is the ratio between the social welfare of the optimal solution and the social welfare of an equilibrium solution, in the worst case over instances of the problem, and in the worst (best) case over corresponding equilibrium solutions.

a direction in the activity space along which a player can produce output. This restricted model captures salient features of the original model: e.g., it maintains the property that the contests can have overlapping criteria.

With a simple transformation of the activity space, the portfolio model can be converted to an equivalent model where a player produces output along a single activity only, i.e., only along the axes, where an activity in the new model corresponds to a portfolio in the old model. In our discussion, we call this new model the *single-activity model* to differentiate it from the original model, which we call the *multi-activity model*.

The positive results for the multi-activity model automatically carry over to the single-activity model. Additionally, it is computationally easy for the players to best respond in the single-activity model. However, we get a different hardness result. Even for the linear model, it is PLS-complete to compute a PNE and CLS-complete (CLS=PPAD∩PLS) to compute a mixed-strategy Nash equilibrium (MNE). These hardness results, particularly the CLS-hardness result, are interesting because single-activity games form a strict, structured and well-motivated subclass of explicit congestion games (a contest awards a $1/k$ fraction of a fixed prize to each winner if there are k winners, but a congestion game can have a cost that is an arbitrary function of the number of winners), yet finding an MNE in these games has the same computational complexity as finding an MNE of explicit congestion games [9] (and finding a fixed-point of gradient descent [49]). We also prove some fixed-parameter tractability results with respect to the number of players and the number of contests.

The rest of this chapter is organized as follows. After summarizing the related work (Section 7.1.2), in Section 7.2 we introduce the general multi-activity model. We also prove the existence of PNE, the pseudo-polynomial convergence of ϵ-best-response dynamics to ϵ-PNE, and present our PoA results. In Section 7.3 we establish the hardness of best-response in linear multi-activity models. Section 7.4 focuses on the single-activity model and presents our results on PLS-completeness and CLS-completeness.

Omitted proofs and additional results are provided in Section 7.6.

7.1.2 Related Work

The model of simultaneous contests with equal sharing allocation of prizes has been studied before in the literature [76, 106]. At a high level, our contribution is to (i) generalize the model and extend the positive results and (ii) study the complexity of computing equilibria.

The linear budget model with a fixed total prize has previously been studied by May et al. [76], to model situations such as the social media platform example discussed earlier. Their theoretical results are similar to our positive results: (i) they prove existence of PNE by showing that the game is an exact potential game; (ii) they establish a PoA bound of 2. For (i), existence of PNE, we give a simpler proof by explicitly constructing the potential function that lifts these results to our general model. Our proof also makes it transparent that the PNE exists because of the equal sharing property of the contests (the *congestion* property), and the other restrictions of the model of May et al. [76]—the linear budget constraints, the linear criteria of contests, and the fixed total prize—are not necessary for the result. Moreover, the proof clarifies that using the budget model is also non-essential, as the result holds for the cost model as well. For (ii), the PoA bound, May et al. [76] prove the result from first principles. In contrast, we use the result of Vetta [105] for submodular social welfare functions to derive the same result in our—more general—setting. In summary, we extend the positive results of [76] to a general model (and we also study computational complexity, which was not considered by May et al. [76]). May et al. [76] perform an empirical study and show that the real-life behavior of social media curators resembles the predictions of the model. Bilo et al. [17] consider a model similar to the single-activity model of our work and show inclusion in the class PLS (but no hardness results).

Models of simultaneous contests that do not have activities (i.e., the players directly produce outputs for each contest, or, equivalently, there is a one-to-one mapping between the activities and the contests) have been extensively studied. Cost-based models where the prizes are awarded based on the players' ranks have been considered by a number of authors [6, 7, 10, 39], including empirical work [7, 39, 70, 111]. Colonel Blotto games, where the players have budget constraints and the prize is awarded to the highest-output player for each contest, were proposed by Borel [23], and have received a significant amount of attention in the literature (e.g., [22, 103, 55, 19, 20, 13, 98, 92, 57, 1, 92, 60, 109]). Simultaneous contests with proportional allocation have been studied by, e.g., [50, 112, 91].

Two related papers are by Birmpas et al. [18] and Elkind et al. [43] (part of Chapter 5). Both these papers do not have activities, i.e., the players produce output directly for the contests, which makes their models a bit different (simpler) than ours, but they add complexity along other dimensions. Therefore, their results are not directly comparable to ours. Birmpas et al. [18] have both budgets and costs in the same model, and they give a constant factor PoA bound by augmenting players' budgets when computing the equilibrium welfare (but not when computing the optimal welfare). Elkind et al. [43] consider a model with only one contest and in the case of incomplete information. Their focus is on mechanism design, and for one of the objectives studied in the paper, they prove that the optimal contest distributes its prize equally to all players who produce output above some threshold, similar to the contests in our paper.

The complexity class PLS (Polynomial Local Search) and the concepts of PLS-hardness and PLS-completeness were introduced by Johnson et al. [64]. PLS consists of discrete local optimization problems whose solutions are easy to verify (the cost of a given solution can be computed in polynomial time and its local neighborhood can be searched in polynomial time). Similar to NP-hard problems, PLS-hard problems are believed to be not solvable in polynomial time. Several natural problems, such as finding a locally optimal solution of MAX-CUT, were shown to be PLS-complete by Schaffer [94]. The problem of finding a PNE in explicit congestion games (which always have a PNE) is also PLS-complete [48], from which it follows that better or best response dynamics take an exponential time to converge in the worst case [64].

The class PPAD (Polynomial Parity Arguments on Directed graph) was introduced by Papadimitriou [89]. Like PLS problems, PPAD problems always have solutions. For PPAD, the existence of a solution is based on a parity argument: In a directed graph where each vertex has at most one predecessor and one successor, if there exists a source vertex (i.e., a vertex with no predecessor), then there exists some other degree-1 vertex. One of the most well-known results in algorithmic game theory is that the problem of finding a mixed-strategy Nash equilibrium (MNE) is PPAD-complete [36, 30].

Recent work has determined the complexity of computing an MNE of an explicit congestion game. The class PPAD∩PLS represents problems that can be solved both by an algorithm that solves PPAD problems and by an algorithm that solves PLS problems. Finding an MNE of an explicit congestion game is in PPAD∩PLS: indeed, this problem can be solved either by finding a PNE (which is also an MNE) using an algorithm for PLS problems or by computing an MNE using an algorithm for PPAD problems. Recently, Fearnley et al. [49] proved that finding a fixed-point of a smooth 2-dimensional function $f : [0,1]^2 \to \mathbb{R}$ using gradient descent is complete for the class PPAD∩PLS; based on this result, Babichenko and Rubinstein [9] proved that finding an MNE of an explicit congestion game is also (PPAD∩PLS)-complete. PLS-complete, PPAD-complete, and (PPAD∩PLS)-complete problems are considered to be hard problems with no known polynomial-time algorithms.

7.2 General Model, Pure Nash Equilibrium, and Price of Anarchy

In this section, we formally define the general multi-activity model, prove the existence of pure Nash equilibria, and show that the price of anarchy (PoA) is 2 for the budget variant of the model and $+\infty$ for the cost variant.

We consider a set of n players, $N = [n]$,[4] who simultaneously produce output along k activities, $K = [k]$. There are m contests, $M = [m]$, which award prizes to the players based on their outputs. The contests may have different prizes and may value the activities differently.

We study two models: in one, the players have output production costs; and in the other, the players have generalized output budgets. Player $i \in N$ chooses an output vector $\boldsymbol{b}_i = (b_{i,\ell})_{\ell \in K} \in \mathbb{R}^k_{\geq 0}$.

- **Cost.** Player i incurs a cost of $c_i(\boldsymbol{b}_i)$ for producing $\boldsymbol{b}_i$, where $c_i : \mathbb{R}^k_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a non-decreasing cost function with $c_i(\mathbf{0}) = 0$ (normalized).

- **Budget.** The player does not incur any cost for the output vector $\boldsymbol{b}_i$, but $\boldsymbol{b}_i$ is restricted to be in a set $\mathcal{B}_i \subseteq \mathbb{R}^k_{\geq 0}$. We assume that $\mathbf{0} \in \mathcal{B}_i$, i.e., players are always allowed to not participate.[5]

Technically, the budget model is a special case of the cost model, and all the positive results for the cost model automatically carry over to the budget model. We make the distinction because when we study the restricted linear models in Section 7.3, the linear cost model and the linear budget model will have different formulations.

Let $\boldsymbol{b} = (\boldsymbol{b}_i)_{i \in N} = (b_{i,\ell})_{i \in N, \ell \in K}$. Each contest $j \in M$ is associated with a pair of functions $f_j : \mathbb{R}^k_{\geq 0} \to \mathbb{R}_{\geq 0}$ and $v_j : \mathbb{Z}_{\geq 0} \to \mathbb{R}_{\geq 0}$. The function f_j, which is an increasing function such that $f_j(\mathbf{0}) = 0$, determines the set of winners of contest j: we say that player i *wins* contest j if $f_j(\boldsymbol{b}_i) \geq 1$ and set $N_j(\boldsymbol{b}) = \{i \in N \mid f_j(\boldsymbol{b}_i) \geq 1\}$. Let $n_j(\boldsymbol{b}) = |N_j(\boldsymbol{b})|$. The function v_j determines how the prizes are allocated: each player in $N_j(\boldsymbol{b})$ receives a prize of $v_j(n_j(\boldsymbol{b}))$. The total prize allocated by contest j is then $n_j(\boldsymbol{b}) \cdot v_j(n_j(\boldsymbol{b}))$. We make the following assumptions about the function v_j, which are necessary for our price of anarchy bounds.

1. $v_j(\ell)$ is a non-increasing function of ℓ: $v_j(\ell) \geq v_j(\ell + 1)$.

[4]Let $[\ell] = \{1, 2, \ldots, \ell\}$ for any positive integer $\ell \in \mathbb{Z}_{>0}$.

[5]The budget model has an alternative interpretation—each player selects a subset of activities among a feasible set of subset of activities for that player—as discussed in Section 7.2.3.

2. $\ell \cdot v_j(\ell)$ is a non-decreasing function of ℓ: $\ell \cdot v_j(\ell) \leq (\ell+1) \cdot v_j(\ell+1)$.

3. $\ell \cdot v_j(\ell)$ is a weakly concave function of ℓ, i.e., $(\ell+1) \cdot v_j(\ell+1) - \ell \cdot v_j(\ell)$ is a non-increasing function of ℓ: $(\ell+2) \cdot v_j(\ell+2) - 2 \cdot (\ell+1) \cdot v_j((\ell+1) + \ell \cdot v_j(\ell) \leq 0$. This condition says that the rate of increase in the total prize allocated by a contest j weakly decreases as the number of winners increases.

Some examples of functions v_j that satisfy these conditions are:

- $x \cdot v_j(x) = x$ or $v_j(x) = 1$. Here, the total value of the prize scales linearly with the number of winners, i.e., the prize awarded to a winner does not change as the number of winners increases.

- $x \cdot v_j(x) = 1$ or $v_j(x) = 1/x$. Here, the total value of the prize remains constant, i.e., the prize awarded to a winner decreases and is equal to the inverse of the number of winners.

- $x \cdot v_j(x) = \sqrt{x}$ or $v_j(x) = 1/\sqrt{x}$. This sits between the previous two examples. Here, the total value of the prize increases, but the prize awarded to a winner decreases as the number of winners increases.

The utility of a player i in the cost and the budget model is, respectively,

$$u_i^C(\boldsymbol{b}) = \sum_{j \in M} v_j(n_j(\boldsymbol{b})) \cdot \mathbb{1}_{\{i \in N_j(\boldsymbol{b})\}} - c_i(\boldsymbol{b}_i), \quad u_i^B(\boldsymbol{b}) = \sum_{j \in M} v_j(n_j(\boldsymbol{b})) \cdot \mathbb{1}_{\{i \in N_j(\boldsymbol{b})\}}$$

where $\mathbb{1}_{\{\dots\}}$ is the indicator function; we omit the superscripts B and C if they are clear from the context.

We will be interested in stable outcomes of multi-activity games, as captured by their Nash equilibria.

Definition 7.1 (Pure-Strategy Nash Equilibrium (PNE))**.** A pure strategy profile $\boldsymbol{b} = (\boldsymbol{b}_i, \boldsymbol{b}_{-i})$ is a *pure Nash equilibrium (PNE)* if for every $i \in N$ and every action $\boldsymbol{b}'_i$ of player i we have $u_i(\boldsymbol{b}) \geq u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})$.

Definition 7.2 (Mixed-Strategy Nash Equilibrium (MNE))**.** A mixed strategy profile $\mu = \times_{\ell \in [n]} \mu_\ell$ is a *mixed Nash equilibrium (MNE)* if for every $i \in N$ and every distribution over actions μ'_i of player i we have
$\mathbb{E}_{\boldsymbol{b} \sim \mu}[u_i(\boldsymbol{b})] \geq \mathbb{E}_{\boldsymbol{b}'_i \sim \mu'_i, \boldsymbol{b}_{-i} \sim \mu_{-i}}[u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})]$, where $\mu_{-i} = \times_{\ell \neq i} \mu_\ell$.

We now give definitions of the price of anarchy and the price of stability, which will be used to study the efficiency of equilibria in our models.

Definition 7.3 (Price of Anarchy (PoA) and Price of Stability (PoS))**.** Let $\mathcal{G}$ denote a class of games and let $\mathcal{I} \in \mathcal{G}$ denote a particular instance of the game. For a given instance of a game $\mathcal{I}$, let $Action(\mathcal{I})$ denote the set of action profiles and $Eq(\mathcal{I})$ denote the set of all MNE. Let $sw(\boldsymbol{b})$ denote the social welfare for an action profile $\boldsymbol{b} \in Action(\mathcal{I})$. The *price of anarchy* and the *price of stability* for class $\mathcal{G}$ are defined as, respectively,

$$PoA = \max_{\mathcal{I}\in\mathcal{G}} \frac{\max_{\boldsymbol{b}\in Action(\mathcal{I})} sw(\boldsymbol{b})}{\min_{\mu\in Eq(\mathcal{I})} \mathbb{E}_{\boldsymbol{b}\sim\mu}[sw(\boldsymbol{b})]}; \quad PoS = \max_{\mathcal{I}\in\mathcal{G}} \frac{\max_{\boldsymbol{b}\in Action(\mathcal{I})} sw(\boldsymbol{b})}{\max_{\mu\in Eq(\mathcal{I})} \mathbb{E}_{\boldsymbol{b}\sim\mu}[sw(\boldsymbol{b})]}.$$

I.e., the PoA (PoS) is the ratio between the optimal social welfare and the equilibrium social welfare in the worst (best) case over possible equilibria and in the worst case over instances of the game.

7.2.1 Existence of Pure-Strategy Nash Equilibrium

We start by showing that multi-activity games are exact potential games. We then use the classic result of Monderer and Shapley [83] to conclude that multi-activity games always have pure Nash equilibria.

Definition 7.4 (Exact Potential Games)**.** [83] A normal form game is an *exact potential game* if there exists a potential function ϕ such that for any player i with utility function u_i, any two strategies $\boldsymbol{b}_i$ and $\boldsymbol{b}'_i$ of player i, and any strategy profile $\boldsymbol{b}_{-i}$ of the other players it holds that

$$u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) - u_i(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) = \phi(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) - \phi(\boldsymbol{b}_i, \boldsymbol{b}_{-i}).$$

Theorem 7.1. *[83] Exact potential games always have a pure-strategy Nash equilibrium. Indeed, every pure strategy profile that maximizes ϕ is a PNE.*

We define the potential functions for the multi-activity budget and cost games as, respectively,

$$\phi^C(\boldsymbol{b}) = \sum_{j\in M} \sum_{\ell\in[n_j(\boldsymbol{b})]} v_j(\ell) - \sum_{i\in N} c_i(\boldsymbol{b}_i), \qquad \phi^B(\boldsymbol{b}) = \sum_{j\in M} \sum_{\ell\in[n_j(\boldsymbol{b})]} v_j(\ell); \tag{7.1}$$

we omit the superscripts B and C if they are clear from the context. We use these potential functions to prove the existence of PNE in multi-activity games. In Section 7.4, we shall also use them to study the complexity of computing equilibria in these games.

Theorem 7.2. *A multi-activity budget/cost game is an exact potential game, and hence has a pure-strategy Nash equilibrium.*

We note that the crucial property required for the proof of Theorem 7.2 is the equal sharing property; some of the other assumptions made in our model, e.g., that the cost functions $c_i(b_i)$ are non-decreasing, $\ell \cdot v_j(\ell)$ are non-decreasing and weakly concave, etc., are not essential for the proof. We also note that for a restricted version of the model the same result was proved by [76, 106], but we believe that our proof using the potential function is simpler.

7.2.2 Approximate Pure-Strategy Nash Equilibrium using Better-Response

The characterization in Theorem 7.2 is very useful for finding approximate equilibria of multi-activity games.

Definition 7.5 (ϵ-Pure-Strategy Nash Equilibrium)**.** A pure strategy profile $\boldsymbol{b} = (\boldsymbol{b}_i, \boldsymbol{b}_{-i})$ is an *ϵ-PNE* if for every $i \in N$ and every action $\boldsymbol{b}'_i$ of player i we have $u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) \leq u_i(\boldsymbol{b}) - \epsilon$.

Definition 7.6 (ϵ-Better-Response)**.** For a pure strategy profile $\boldsymbol{b} = (\boldsymbol{b}_i, \boldsymbol{b}_{-i})$, a player $i \in N$, and an action $\boldsymbol{b}'_i$ of player i, the move from $\boldsymbol{b}_i$ to $\boldsymbol{b}'_i$ is an *ϵ-better-response move* if $u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) > u_i(\boldsymbol{b}) + \epsilon$.

From the definitions above, it is immediate that a pure strategy profile is an ϵ-PNE if and only if it does not admit any ϵ-better-response moves. Now, as multi-activity games are exact potential games, each ϵ-better-response increases the potential by at least ϵ. As the potential function is bounded from above, a sequence of ϵ-better-response moves necessarily terminates, and the resulting profile is an ϵ-PNE.

Corollary 7.3. *In multi-activity games, any sequence of ϵ-better-response moves arrives to an ϵ-PNE in at most $n \cdot m \cdot (\max_j v_j(1))/\epsilon$ steps.*

Corollary 7.3 provides a pseudo-polynomial bound in the number of steps required for convergence to an approximate PNE, which can be meaningful in situations where the prizes (i.e., the values $v_j(1)$) are not very large. It also implies that an ϵ-PNE can be computed in pseudo-polynomial time if we can compute ϵ-better-responses efficiently.

7.2.3 Social Welfare, Price of Anarchy, and Price of Stability

To start, let us briefly discuss what would be a natural definition(s) of social welfare for multi-activity games.

Consider first the budget model. From the perspective of the players, social welfare naturally corresponds to the total prize allocated. This definition is also natural from the perspective of the contests. Indeed, consider the motivating example involving the social media curators (players) and subscribers (contests). The welfare of the curators corresponds to the total subscribers' attention they receive, which is equal to the total prize allocated. On the other hand, the welfare of the subscribers corresponds to the compatible curators who serve them, which again corresponds to the total prize allocated. For this definition of social welfare, we prove an upper bound of 2 on the price of anarchy (PoA) and a lower bound of $2 - o(1)$ on the price of stability (PoS), so both of these bounds are tight.

For the cost model, the formulation of social welfare can be similarly motivated, with or without subtracting the cost. If one assumes that the cost to the players is a sunk cost, then it is reasonable to subtract it from the social welfare. On the other hand, if one assumes that the cost gets transferred to the contest organizers (or the society), then it should not be subtracted. In any case, for both of these definitions, the PoS (and therefore, PoA) turns out to be infinite.

7.2.3.1 Budget Model

For the budget model, the social welfare is equal to the total prize that gets allocated:

$$sw(\boldsymbol{b}) = \sum_{i \in N} u_i(\boldsymbol{b}) = \sum_{i \in N} \sum_{j \in M} v_j(n_j(\boldsymbol{b})) \cdot \mathbb{1}_{\{i \in N_j(\boldsymbol{b})\}} = \sum_{j \in M} v_j(n_j(\boldsymbol{b})) \cdot n_j(\boldsymbol{b}). \tag{7.2}$$

To prove the upper bound of 2 on PoA, we shall use the result of Vetta [105] for *submodular* social welfare functions. This result states that the PoA can be upper-bounded by 2 if the following conditions are satisfied:

1. the utility of the players and the social welfare are measured in the same units,
2. the total utility of the players is at most the social welfare,
3. the social welfare function is non-decreasing and submodular, and
4. the private utility of a player is at least as much as the *Vickrey utility* (see Definition 7.8).

By definition, our model satisfies the first two requirements on this list. To show that it also satisfies the last two requirements, we reinterpret it using a different notation.

In the budget game, an action by a player i effectively corresponds to a subset of the contests that i wins. The output vector produced by the player to win a given set of contests is not important as long as the set of contests that she wins remains the same.[6] Hence, we will represent the action taken by player i as a subset of elements from a set of size m, as follows: Let $\mathcal{G}$ denote a set of $n \cdot m$ elements partitioned into n size-m sets $\mathcal{G}^1, \mathcal{G}^2, \ldots, \mathcal{G}^n$, where $\mathcal{G}^i = \{g_1^i, g_2^i, \ldots, g_m^i\}$. We represent the set of feasible actions for player i by $\mathcal{A}_i \subseteq 2^{\mathcal{G}^i}$: an action $A_i \in \mathcal{A}_i$ contains g_j^i if and only if player i satisfies the criteria for contest $j \in M$. Note that $\emptyset \in \mathcal{A}_i$, because we assume that each player is allowed to not participate and produce a $\mathbf{0}$ output (which does not win her any contests). An action profile $A = (A_i)_{i \in N} \in \times_{i \in N} \mathcal{A}_i$ can be equivalently written as $\cup_{i \in N} A_i \in \cup_{i \in N} \mathcal{A}_i$ because the sets $\mathcal{A}_i$ are disjoint as per our notation.

As before, let $N_j(A) = \{i \mid g_j^i \in A\}$ denote the players who win contest j under action profile A, and let $n_j(A) = |N_j(A)|$. The utility functions and the social welfare function can be rewritten using the new set notation as follows:

$$u_i(A) = \sum_{j \in M} v_j(n_j(A)) \cdot \mathbb{1}_{\{g_j^i \in A\}}; \qquad sw(A) = \sum_{j \in M} v_j(n_j(A)) \cdot n_j(A).$$

Let us formally define submodular set functions.

Definition 7.7 (Submodular Functions)**.** A function $f : 2^\Omega \to \mathbb{R}$ is *submodular* if for every pair of subsets $S, T \subseteq \Omega$ such that $S \subseteq T$ and every $x \in \Omega \setminus T$ we have

$$f(S \cup \{x\}) - f(S) \geq f(T \cup \{x\}) - f(T).$$

The next two lemmas prove that our model satisfies the conditions required to use the result of Vetta [105].

Lemma 7.4. *For the budget model, if the total prize allocated by a contest is a weakly concave function of the number of winners of the contest, then the social welfare function is submodular.*

Definition 7.8 (Vickrey Utility)**.** The *Vickrey utility* of player i at an action profile A is the loss incurred by other players due to i's participation, i.e.,

$$u_i^{\text{Vickrey}}(A) = \sum_{q \neq i} u_q(\emptyset, A_{-i}) - \sum_{q \neq i} u_q(A_i, A_{-i}),$$

[6]Not important for PoA analysis, but very important for computational complexity: the budget constraint for a player i boils down to selecting a subset of contests in some feasible set of subsets, say $\mathcal{A}_i \in 2^M$. The budget constraint is a concise way of representing this $\mathcal{A}_i$, and $\mathcal{A}_i$ may be of size exponential in the representation.

where $\emptyset$ is the action to not participate (or, equivalently, produce 0 output and not win any contests).

Lemma 7.5. *For the budget model, if the total prize allocated by a contest is a non-decreasing function of the number of winners of the contest, then for every profile A and every player i, the utility of i is at least as large as her Vickrey utility, i.e., $u_i(A_i, A_{-i}) \geq u_i^{Vickrey}(A)$.*

Theorem 7.6. *The social welfare in any mixed-strategy Nash equilibrium of a multi-activity budget game is at least $1/2$ of the optimum social welfare.*

The following result complements the upper bound of 2 for PoA by giving a lower bound of 2 for PoS, which is based upon an example in [106] and ensuring that the equilibrium is unique.

Theorem 7.7. *There are instances of multi-activity budget games where the social welfare in every mixed-strategy Nash equilibrium approaches $1/2$ of the optimum social welfare as the number of players grows.*

7.2.3.2 Cost Model

As discussed before, we consider two definitions of social welfare for the cost model. The first definition does not subtract the costs from social welfare and corresponds to the definition given in (7.2). The second definition subtracts the costs from the social welfare and is equal to

$$\overline{sw}(\boldsymbol{b}) = \sum_{i \in N} (u_i(\boldsymbol{b}) - c_i(\boldsymbol{b})) = \sum_{j \in M} v_j(n_j(\boldsymbol{b})) \cdot n_j(\boldsymbol{b}) - \sum_{i \in N} c_i(\boldsymbol{b}). \tag{7.3}$$

Next, we prove that the PoS can be unbounded for both definitions of social welfare in the cost model.

Theorem 7.8. *There are instances of multi-activity cost games where the social welfare in every mixed-strategy Nash equilibrium can be arbitrarily low compared to the optimum social welfare. This holds even if there are at most two players, two activities, and two contests.*

7.3 Multi-Activity Games: Hardness of Best-Response

In this section, we focus on a restricted model: each contest uses a linear criterion, and the budget constraint or the cost function of each player are also linear. We call this model the *linear multi-activity model*. Formally, for an output profile $\boldsymbol{b} = (\boldsymbol{b}_i)_{i \in N} = (b_{i,\ell})_{i \in N, \ell \in K}$, the winners of contest j are $N_j(\boldsymbol{b}) = \{i \mid \sum_{\ell \in K} w_{j,\ell} b_{i,\ell} \geq 1\}$, where $w_{j,\ell} \in \mathbb{R}_{\geq 0}$ is a non-negative weight that contest j has for activity ℓ. Similarly, the linear budget constraint of a player i is of the form $\sum_{\ell \in K} \beta_{i,\ell} b_{i,\ell} \leq 1$, where $\beta_{i,\ell} \in \mathbb{R}_{>0}$. Likewise, a linear cost function for player i is of the form $\sum_{\ell \in K} c_{i,\ell} b_{i,\ell}$, where $c_{i,\ell} \in \mathbb{R}_{>0}$.

We also impose another constraint: we assume that for each contest its total prize is fixed. That is, each contest $j \in M$ is associated with a total prize V_j, and if there are ℓ winners, each winner gets a prize V_j/ℓ.

We study the computational complexity of best-response in this linear multi-activity model. Observe that it suffices to consider this problem for $n = 1$. Indeed, consider a player $i \in N$. If there are ℓ winners other than i for a given contest j, then i gets a prize of $v_j(\ell + 1) = V_j/(\ell + 1)$ from this contest if she satisfies this contest's criteria, and 0 otherwise. By scaling the values of all contests appropriately, we reduce i's optimization problem to one where i is the only player in the game.

In the next theorem, we prove that finding a best-response exactly or approximately is NP-hard.

Theorem 7.9. *In the linear multi-activity model, both cost and budget, a player cannot approximate a best response beyond a constant factor in polynomial time unless* P = NP.

We provide two separate reductions for the cost and budget models. For the cost model, we reduce from 4-Regular-SetCover, which is the SetCover problem where every set has a size of exactly 4. For the budget model, we reduce from MAX-2-SAT-3, which is the MAX-SAT problem with a 2-CNF formula where each variable occurs at most 3 times. Both 4-Regular-SetCover and MAX-2-SAT-3 problems cannot be approximated better than a constant factor unless $P = NP$, and our reductions preserve these hardness of approximations.

We note that in a single-player case finding a best response is equivalent to finding a PNE. We obtain the following corollary.

Corollary 7.10. *In the linear multi-activity model, both cost and budget, the problem of computing an exact or an approximate PNE is* NP*-hard.*

Theorem 7.9 proves that the problem of finding a best response in the linear multi-activity model does not admit a polynomial-time approximation scheme. For restricted versions of the model, a constant factor approximation can be found. For example, for the budget model, if the contests have $\{0, 1\}$ weights for the activities, and the player has a budget that she can distribute across any of the activities (the hard instance constructed in Theorem 7.9 for the budget model satisfies these conditions), then the problem becomes a submodular maximization problem with a polynomial-time constant-factor approximation algorithm. However, we feel that a constant-factor approximation result is of limited usefulness in the context of computing a best response or a Nash equilibrium.

Next, we study the fixed-parameter tractability of the problem of computing a best response. There are three natural parameters of the model: the number of players n, the number of contests m, and the number of activities k. We have already shown that the problem is NP-hard even with only one player, $n = 1$. On the positive side, we show that the problem becomes tractable if either the number of contests m or the number of activities k is a constant.

Theorem 7.11. *In the linear multi-activity model, both cost and budget, a player can compute a best response in polynomial time if either the number of contests or the number of activities is bounded by a constant.*

7.4 Single-Activity Games: CLS-Completeness

In this section, we focus our attention on the single-activity model, and show that, for both cost and budget models, it is PLS-complete to find a pure Nash equilibrium and (PPAD∩PLS)-complete to find a mixed Nash equilibrium (MNE). Note that the set of MNE of a game is a super-set of the set of PNE, so finding an MNE is at least as easy as finding a PNE.

In a single-activity game, for every player i there is at most one activity ℓ for which the output $b_{i,\ell}$ may be strictly positive; for every other activity $\ell' \neq \ell$ it holds that $b_{i,\ell'} = 0$. Additionally, we assume that each contest i has a fixed total prize, which it distributes equally among all winners.

Theorem 7.12. *In the linear single-activity model, both cost and budget, it is* $(\mathrm{PPAD} \cap \mathrm{PLS})$*-complete to find a mixed-strategy Nash equilibrium.*

A recent paper by Fearnley et al. [49] proved the interesting result that finding a fixed-point of a 2-dimensional smooth function ($f : [0,1]^2 \to \mathbb{R}$) given by a circuit of addition and multiplication operators using gradient descent, 2D-GD-FixedPoint, is complete for the class PPAD∩PLS. Based on this result, Babichenko and Rubinstein [9] proved that the problem of computing an MNE of an explicit congestion game, ExpCong, is also complete for PPAD∩PLS. Babichenko and Rubinstein [9] do this by first reducing 2D-GD-FixedPoint to identical interest 5-polytensor games, 5-Polytensor, and then 5-Polytensor to ExpCong. In our proof for Theorem 7.12, we reduce 5-Polytensor to the single-activity game, proving the required result.

Before moving to the proof, let us define identical interest polytensor games, κ-Polytensor. Polytensor games are a generalization of the better known polymatrix games, Polymatrix; specifically, Polymatrix = 2-Polytensor.

Definition 7.9 (Polymatrix: Identical Interest Polymatrix Game)**.** There are n players. Player i chooses from a finite set of actions A_i. The utility of player i for action profile $(a_i, \boldsymbol{a}_{-i}) = \boldsymbol{a} = (a_j)_{j\in[n]} \in \times_{j\in[n]} A_j$ is given by $u_i(\boldsymbol{a}) = \sum_{j\in[n], j\neq i} u_{i,j}(a_i, a_j)$, where $u_{i,j} : A_i \times A_j \to \mathbb{R}_{\geq 0}$. The players have identical interest, i.e., $u_{i,j} = u_{j,i}$.

The definition of κ-Polytensor games is similar to that of Polymatrix games: instead of $u_{i,j}$ we have u_S, where $S \subseteq [n], |S| = \kappa$.

Definition 7.10 (κ-Polytensor: Identical Interest κ-Polytensor Game)**.** There are n players. Player i chooses from a finite set of actions A_i. The utility of player i for action profile $\boldsymbol{a} = (a_j)_{j\in[n]} \in \times_{j\in[n]} A_j$ is given by $u_i(\boldsymbol{a}) = \sum_{S\subseteq[n], i\in S, |S|=\kappa} u_S(\boldsymbol{a}_S)$, where $\boldsymbol{a}_S = (a_j)_{j\in S}$ is the action profile of the players in S and $u_S : \times_{j\in S} A_j \to \mathbb{R}_{\geq 0}$.

When the number of actions for each player, $|A_i|$, is bounded by m, note that the representation size of a κ-Polytensor game is $O(n^\kappa m^\kappa)$. In particular, if κ is a constant and $m = poly(n)$, then κ-Polytensor games admit a succinct representation.

Proof Sketch for Theorem 7.12. Below, we provide a reduction from 3-Polytensor to the budget game for a cleaner presentation of the main steps (unlike 5-Polytensor, 3-Polytensor is not known to be (PPAD∩PLS)-complete). Similar steps with more calculations apply to 5-Polytensor, for both cost and budget models (we provide this argument in Section 7.6).

Take an arbitrary instance of 3-POLYTENSOR with n players; we shall use the same notation as in Definition 7.10. We construct a single-activity game with n players, $\sum_{i\in[n]} |A_i|$ activities, and a polynomial number of contests to be defined later.

The $\sum_{i\in[n]} |A_i|$ activities have a one-to-one association with the actions of the players. The activities are partitioned into n subsets, so that the i-th subset has size $|A_i|$ and is associated with player i; we identify these activities with the set A_i. Player i has a budget of 1 that they can use to produce output along any activity from A_i, but they have 0 budget for the activities in A_j for $j \neq i$. Effectively, as we are in a single-activity model, player i selects an activity from the activities in A_i and produces an output of 1 along it. Note that the players have disjoint sets of activities for which they can produce outputs.

All the contests we construct are associated with exactly three players and at most three activities. We shall denote a contest by $\mathcal{C}_{i,j,k}(A)$, where (i) $S = \{i, j, k\}$ are the three distinct players whose utility function in the polytensor game, u_S, will be used to specify the prize of contest $\mathcal{C}_S(A)$; (ii) the contest $\mathcal{C}_S(A)$ awards its prize to any player who produces an output of at least 1 along the activities in A; (iii) the activities in A are from $A_i \cup A_j \cup A_k$ with $|A| \leq 3$ and $|A \cap A_\ell| \leq 1$ for $\ell \in S$. We shall call a contest $\mathcal{C}_S(A)$ a Type-ℓ contest if $|A| = \ell$.

Let us focus on a fixed set of three players $S = \{i, j, k\}$. We create contests to exactly replicate the utility that these players get from u_S. If we can do this, then, by repeating the same process for every triple of players, we will replicate the entire 3-POLYTENSOR game. The utility that player i gets from u_S is $u_S(a_i, a_j, a_k)$, where a_i, a_j, and a_k are the actions of the three players. We have the following contests:

Type-3 Contests. Let us add a contest $\mathcal{C}_S(a_i, a_j, a_k)$ with prize $v_S(a_i, a_j, a_k)$ for every $(a_i, a_j, a_k) \in A_i \times A_j \times A_k$. Later, we shall specify the $v_S(a_i, a_j, a_k)$ values based on $u_S(a_i, a_j, a_k)$ values. Contest $\mathcal{C}_S(a_i, a_j, a_k)$ distributes a prize of $v_S(a_i, a_j, a_k)$ to players who produce output along the activities a_i, a_j, or a_k.

Say, the players i, j, k select the actions a_i^*, a_j^*, a_k^*. The total prize that player i gets from the contests we added is:

$$\frac{v_S(a_i^*, a_j^*, a_k^*)}{3} + \sum_{a_j \neq a_j^*} \frac{v_S(a_i^*, a_j, a_k^*)}{2} + \sum_{a_k \neq a_k^*} \frac{v_S(a_i^*, a_j^*, a_k)}{2} + \sum_{\substack{a_j \neq a_j^* \\ a_k \neq a_k^*}} v_S(a_i^*, a_j, a_k) \quad (7.4)$$

In expression (7.4), the first term is for the prize that i shares with j and k, the second term is for the prizes that i shares with k, but not with j, the third term is for the prizes that i shares with j, but not with k, and the fourth term is for the prizes that i does not share with j or k.

In expression (7.4), the first term $\frac{1}{3}v_S(a_i^*, a_j^*, a_k^*)$ resembles the utility that the players obtain in the polytensor game, $u_S(a_i^*, a_j^*, a_k^*)$. If we were to set $v_S(a_i^*, a_j^*, a_k^*) = 3u_S(a_i^*, a_j^*, a_k^*)$, then it would be exactly equal to it. However, we also need to take care of the additional terms in expression (7.4). Hence, we will add Type-1 and Type-2 contests to cancel these terms.

The expression in (7.4) can be rewritten as $\sum_{\substack{a_j \in A_j \\ a_k \in A_k}} v_S(a_i^*, a_j, a_k) - \sum_{a_j \neq a_j^*} \frac{v_S(a_i^*, a_j, a_k^*)}{2} - \sum_{a_k \neq a_k^*} \frac{v_S(a_i^*, a_j^*, a_k)}{2} - \frac{2v_S(a_i^*, a_j^*, a_k^*)}{3}$.

Type-1 Contests. Let us add a contest $\mathcal{C}_S(a_i')$ with prize

$$\sum_{a_i \neq a_i', a_j \in A_j, a_k \in A_k} v_S(a_i, a_j, a_k)$$

for every $a_i' \in A_i$. This contest $\mathcal{C}_S(a_i')$ awards its prize to any player who produces output along activity a_i' (effectively, it awards the prize to player i if they produce output along a_i', because no other player can produce output along a_i'). Similarly, we add the contests $\mathcal{C}_S(a_j')$ and $\mathcal{C}_S(a_k')$ for $a_j' \in A_j$ and $a_k' \in A_k$, respectively. The total prize that player i gets from Type-1 and Type-3 contests is $\sum_{a_i, a_j, a_k} v_S(a_i, a_j, a_k) - \sum_{a_j \neq a_j^*} \frac{1}{2} v_S(a_i^*, a_j, a_k^*) - \sum_{a_k \neq a_k^*} \frac{1}{2} v_S(a_i^*, a_j^*, a_k) - \frac{2}{3} v_S(a_i^*, a_j^*, a_k^*)$. As $\sum_{a_i, a_j, a_k} v_S(a_i, a_j, a_k)$ does not depend upon the action a_i^* selected by i, the utility of player i is effectively

$$-\sum_{a_j \neq a_j^*} \frac{1}{2} v_S(a_i^*, a_j, a_k^*) - \sum_{a_k \neq a_k^*} \frac{1}{2} v_S(a_i^*, a_j^*, a_k) - \frac{2}{3} v_S(a_i^*, a_j^*, a_k^*).$$

Type-2 Contests. Let us add a contest $\mathcal{C}_S(a_i', a_j')$ with prize

$$\sum_{a_k \in A_k} \frac{1}{2} v_S(a_i', a_j', a_k)$$

for every $a_i' \in A_i$ and $a_j' \in A_j$. This contest $\mathcal{C}_S(a_i', a_j')$ awards its prize to players who produce output along activity a_i' or a_j'. In a similar manner, we add contests corresponding to the actions of the other 5 possible combinations of players among the three players i, j, k, e.g., $\mathcal{C}_S(a_i', a_k')$ for $a_i' \in A_i$ and $a_k' \in A_k$, and so on. The net utility that player i gets from Type-1, Type-2 and Type-3 contests is $\frac{1}{3} v_S(a_i^*, a_j^*, a_k^*)$. We set $v_S(a_i, a_j, a_k) = 3u_S(a_i, a_j, a_k)$ for every $(a_i, a_j, a_k) \in A_i \times A_j \times A_k$, and we are done. $\square$

In Theorem 7.12, we analyzed the complexity of computing an MNE. As we have shown earlier, single-activity cost and budget games always have a PNE, and therefore, it is relevant to know the complexity of computing a PNE. In the next theorem, we prove that computing a PNE is PLS-complete, and this is true even if all the players are identical. In the proof, we reduce MAX-CUT to the problem of finding a PNE in a particular class of single-activity games with identical players. For this class of single-activity games, finding an MNE is easy, which highlights that the class of single-activity games where PNE is hard (PLS-complete) to compute is strictly larger than the class of single-activity games where MNE is hard (PPAD∩PLS-complete) to compute.

Theorem 7.13. *In the single-activity models, both cost and budget, it is* PLS*-complete to find a pure-strategy Nash equilibrium. The result holds even if all players are identical.*

A direct corollary of this PLS-completeness result is that better/best-response dynamics takes an exponential number of steps to converge for some instances of the problem [64].

Regarding fixed-parameter tractability (FPT), in Section 7.6 we show that a PNE is efficiently computable for both cost and budget single-activity models if the number of players is a constant, and for the budget model if the number of contests is a constant. Providing FPT results for other cases remains open.

7.5 Conclusion

In this chapter, we studied a model of simultaneous contests and analyzed the existence, efficiency, and computational complexity of equilibria in these contests. Given the real-life relevance of the three-level model, player–activity–contest, it will be interesting to study it for prize allocation rules other than the equal-sharing allocation, such as rank-based allocation, proportional allocation, etc. For these contests, one may investigate the properties of the equilibria and their computational complexity. We also believe that there is much to explore regarding the computational complexity of simultaneous contests, in general.

7.6 Omitted Proofs and Additional Results

7.6.1 From Section 7.2

Proof of Theorem 7.2. We shall prove the result for the cost game. The budget game can be reduced to the cost game by endowing it with cost function $c(\cdot)$ that is always equal to 0 (the presence of the budget constraint does not affect the proof).

Let player i make a move by changing its strategy from $\boldsymbol{b}_i$ to $\boldsymbol{b}'_i$, and let $\boldsymbol{b}_{-i}$ be the pure strategy profile of the other players. Let $S \subseteq M$ and $S' \subseteq M$ be the contests that player i wins for strategy profiles $\boldsymbol{b} = (\boldsymbol{b}_i, \boldsymbol{b}_{-i})$ and $(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})$, respectively. We have

$$n_j(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) = \begin{cases} n_j(\boldsymbol{b}) - 1 & \text{if } j \in S \setminus S' \\ n_j(\boldsymbol{b}) + 1 & \text{if } j \in S' \setminus S \\ n_j(\boldsymbol{b}) & \text{in all other cases.} \end{cases}$$

The change in the potential $\phi(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) - \phi(\boldsymbol{b})$ can be expressed as

$$\begin{aligned} &\sum_{j \in M} \left(\sum_{\ell \in [n_j(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})]} v_j(\ell) - \sum_{\ell \in [n_j(\boldsymbol{b})]} v_j(\ell) \right) - (c_i(\boldsymbol{b}'_i) - c_i(\boldsymbol{b}_i)) \\ &\qquad\qquad + \sum_{q \in N \setminus \{i\}} \underbrace{(c_q(\boldsymbol{b}_q) - c_q(\boldsymbol{b}_q))}_{=0}, \\ &= \underbrace{\left(\sum_{j \in S' \setminus S} v_j(n_j(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})) - c_i(\boldsymbol{b}'_i) \right)}_{u_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})} - \underbrace{\left(\sum_{j \in S \setminus S'} v_j(n_j(\boldsymbol{b})) - c_i(\boldsymbol{b}_i) \right)}_{u_i(\boldsymbol{b})}, \end{aligned}$$

which is exactly equal to the change in the utility of the player i. □

Proof of Corollary 7.3. The potential function is bounded from above by $\phi(\boldsymbol{b}) \leq \sum_{j \in M} \sum_{\ell \in [n_j(\boldsymbol{b})]} v_j(\ell) \leq \sum_{j \in M} \sum_{\ell \in [n_j(\boldsymbol{b})]} v_j(1) \leq \sum_{j \in M} n \cdot v_j(1) \leq n \cdot m \cdot (\max_j v_j(1))$. As the game is an exact potential game, every ϵ-better-response move increases the potential by more than ϵ. So, in $n \cdot m \cdot (\max_j v_j(1))/\epsilon$ steps the sequence should converge to an ϵ-PNE. □

Proof of Lemma 7.4. Let $S, T \subseteq \mathcal{G}$, $S \subseteq T$, and $x \in \mathcal{G} \setminus T$. We need to show that

$$sw(S \cup \{x\}) - sw(S) \geq sw(T \cup \{x\}) - sw(T).$$

Let $x = g^i_j$ for some player i and contest j. As $S \subseteq T$, $n_j(S) \leq n_j(T)$. Now, as the total prize allocated by contest j (which is equal to $\ell \cdot v_j(\ell)$ if there are ℓ winners) is a concave function, we get

$$\begin{aligned} &(n_j(S) + 1) \cdot v_j(n_j(S) + 1) - n_j(S) \cdot v_j(n_j(S)) \\ &\qquad\qquad \geq (n_j(T) + 1) \cdot v_j(n_j(T) + 1) - n_j(T) \cdot v_j(n_j(T)). \end{aligned}$$

As $n_j(S \cup \{x\}) = n_j(S) + 1$ and $n_j(T \cup \{x\}) = n_j(T) + 1$, and $n_{j'}(S \cup \{x\}) = n_{j'}(S)$ and $n_{j'}(T \cup \{x\}) = n_{j'}(T)$ for $j' \neq j$, we get the required result. □

Proof of Lemma 7.5. By definition of u_i,

$$u_i(A_i, A_{-i}) = \sum_{j:g_j^i \in A_i} v_j(n_j(A_i, A_{-i})) = \sum_{j:g_j^i \in A_i} v_j(n_j(\emptyset, A_{-i}) + 1). \tag{7.5}$$

By changing her action from $\emptyset$ to A_i, player i starts winning contests in $J = \{j \mid g_j^i \in A_i\}$. Therefore, she affects the prize allocation of players who were already receiving prizes from contests in J. In particular, for any $j \in J$, the $n_j(\emptyset, A_{-i})$ players who were receiving a prize of $v_j(n_j(\emptyset, A_{-i}))$ now receive a prize of $v_j(n_j(\emptyset, A_{-i}) + 1)$. So,

$$\begin{aligned}&\sum_{\iota \neq i} u_\iota(\emptyset, A_{-i}) - \sum_{\iota \neq i} u_\iota(A_i, A_{-i}) \\ &\qquad = \sum_{j|g_j^i \in A_i} n_j(\emptyset, A_{-i}) \cdot (v_j(n_j(\emptyset, A_{-i})) - v_j(n_j(\emptyset, A_{-i}) + 1)). \end{aligned} \tag{7.6}$$

As the total prize allocated by a contest is a non-decreasing function of the number of winners, we have

$$\begin{aligned}&(n_j(\emptyset, A_{-i}) + 1) \cdot v_j(n_j(\emptyset, A_{-i}) + 1) \geq n_j(\emptyset, A_{-i}) \cdot v_j(n_j(\emptyset, A_{-i})) \\ &\implies v_j(n_j(\emptyset, A_{-i}) + 1) \geq n_j(\emptyset, A_{-i}) \cdot (v_j(n_j(\emptyset, A_{-i})) - v_j(n_j(\emptyset, A_{-i}) + 1)) \\ &\implies \sum_{j:g_j^i \in A_i} v_j(n_j(\emptyset, A_{-i}) + 1) \\ &\qquad\qquad \geq \sum_{j:g_j^i \in A_i} n_j(\emptyset, A_{-i}) \cdot (v_j(n_j(\emptyset, A_{-i})) - v_j(n_j(\emptyset, A_{-i}) + 1)).\end{aligned}$$

Combining this inequality with (7.5) and (7.6), we get the required result. □

Proof of Theorem 7.6. The social welfare is a non-decreasing function, as the total prize allocated by a contest weakly increases as the number of winners increases. This observation, together with Lemmas 7.4 and 7.5, shows that our setting satisfies the necessary conditions of Theorem 5 of Vetta [105]. This gives us the required bound on the price of anarchy. □

Proof of Theorem 7.7. We describe an instance of the game, adapted from Example 5.37 of Vojnovic [106], that proves the required result.

Let there be n players, n activities, and n contests, i.e., $m = k = n$. Each player has a budget of 1 that he can spend on any of the activities. Contest 1 has a prize of $n + 1$ and the remaining $n - 1$ contests have a prize of 1 each. Contest i gives (a share of) its prize to any player whose output is at least 1 for activity i.

This game has a unique equilibrium where every player produces an output of 1 for activity 1. This happens because even if all players play action 1, each player gets a share of the prize from contest 1, which is equal to $(n+1)/n > 1$. On the other hand, by spending their budget on any other activity they can get a prize of at most 1.

Now, the social welfare for this equilibrium is $n+1$. On the other hand, if each player selects a different activity then the social welfare would be $2n$. So, we get a price of stability equal to $\frac{2n}{n+1} \to 2$ as $n \to \infty$. □

Proof of Theorem 7.8. We analyze the two variants of the model separately.

Subtracting Costs from Social Welfare. We will construct an instance with two players, two activities, and two contests, i.e., $n = k = m = 2$. Contest j distributes a prize of $2+\epsilon$ to players who produce an output of 1 along activity j and does not care about the output along activity $(3-j)$. Each player has a linear cost function with a slope of 1 for each activity.

This game has a unique equilibrium where both players produce an output of 1 for both activities. Because, irrespective of what the other player is doing, a player would get a prize of strictly greater than 1 by producing an output of 1 (with cost 1) for each activity. The total welfare for this equilibrium is $2*(2+\epsilon) - 2*(1+1) = 2\epsilon$.

Consider a socially optimal situation where player i produces output along activity i: the total prize allocation is still $2*(2+\epsilon)$, but the total cost is 2, so the total welfare is $2+2\epsilon$. Therefore, the price of anarchy is $\frac{2+2\epsilon}{2\epsilon}$, which tends to $+\infty$ as $\epsilon \to 0$.

Not Subtracting Costs from Social Welfare. In this definition of social welfare, too, a simple example shows that the price of anarchy is unbounded. Consider an instance with one player, two activities, and two contests. The two activities have a one-to-one correspondence with the contests, i.e., contest j cares only about activity j. The player can win the first contest, which has a prize of ϵ, with a cost of $\epsilon/2$, and she can win the second contest, which has a prize of 1, with a cost of $1+\epsilon/2$. In equilibrium (unique), the player only wins the first contest. The PoS is $1/\epsilon$, which tends to $+\infty$ as $\epsilon \to 0$. Here, the social welfare and the utility of the player are not aligned, as the player cares about costs whereas the social welfare does not. Hence a bad PoS is expected. □

7.6.2 From Section 7.3

Proof of Theorem 7.9. The following two lemmas prove the result separately for the cost and the budget models.

Lemma 7.14. *In the linear multi-activity cost model, there is no polynomial-time approximation scheme for best-response, unless* P = NP.

Proof of Lemma 7.14. We provide a reduction from 4-Regular-SetCover, which is NP-hard to approximate by a factor of $2 - \epsilon$ for any $\epsilon > 0$ [59].

4-Regular-SetCover: An instance of this problem is given by a universe $\mathcal{U} = \{1, 2, \ldots, \hat{m}\}$, a set $\mathcal{S} = \{S_1, S_2, \ldots, S_{\hat{k}}\}$, where each S_j is a subset of $\mathcal{U}$ of size 4 ($|S_j| = 4$). The objective is to find a set cover of $\mathcal{U}$ of minimum possible size, i.e., find a collection of sets $\mathcal{S}' \subseteq \mathcal{S}$ with minimum possible $|\mathcal{S}'|$ such that $\cup_{S_j \in \mathcal{S}'} = \mathcal{U}$. We can assume without loss of generality that $\cup_{S_j \in \mathcal{S}} = \mathcal{U}$ (as otherwise we obviously cannot have a set cover).

We create an instance of the multi-activity cost model with 1 player, $m = \hat{k} + \hat{m}$ contests, and a set of $k = \hat{k}$ activities K. We split the contests into two groups of size $\hat{k}$ and $\hat{m}$, respectively: $M_1 = \{1, \ldots, \hat{k}\}$ and $M_2 = \{\hat{k} + 1, \ldots, \hat{k} + \hat{m}\}$.

- Each contest in M_1 has a prize of value $(\hat{m} + 1)$. Let the weights for these contests be defined as:
$$w_{j,\ell} = \begin{cases} 1, & \text{for } \ell = j \\ 0, & \text{otherwise} \end{cases}$$
where $j \in M_1$ and $\ell \in K$. That is, a contest $j \in M_1$ awards a prize of $(\hat{m} + 1)$ to the player as long as she produces an output of at least 1 along activity j.

- Each contest in M_2 has a prize of value $1 + 1/(\hat{m} + 1)$. To receive the prize of contest $(j + \hat{k}) \in M_2$, the player needs to produce a total output of 1 along activities $\{\ell \mid j \in S_\ell\}$. Formally, the weights are:
$$w_{j,\ell} = \begin{cases} 1, & \text{if } (j - \hat{k}) \in S_\ell \\ 0, & \text{otherwise} \end{cases}$$
where $j \in M_2$ and $\ell \in K$.

- The player has a marginal cost of $(\hat{m} + 2)$ for producing output along any of the activities, i.e., $c_{1,\ell} = \hat{m} + 2$ for $\ell \in K$.

Observe that the contests in M_1 are much more valuable to the player than the ones in M_2: each contest in M_1 has a prize of $(\hat{m} + 1)$, which is more than the total prize of all contests in M_2 which is equal to $\hat{m} + \hat{m}/(\hat{m} + 1)$. Note that it is sub-optimal for the player to produce an output that is strictly more than 1 along any activity: this would contribute to her cost without affecting the prizes she gets. Also, it is sub-optimal for a player to produce an output that is strictly between 0 and 1 for any activity. Indeed, let $K' = \{\ell \in K \mid 0 < b_{1,\ell} < 1\}$, and set $\beta = \sum_{\ell \in K} b_{1,\ell}$. Suppose for the sake of contradiction that $K' \neq \emptyset$. The activities in K' do not help the player to win contests in M_1. If $\beta < 1$, they do not help the player to win contests in M_2 either, so the player can improve her utility by setting $b_{1,\ell} = 0$ for each $\ell \in K'$. If $\beta \geq 1$, these efforts may contribute towards winning contests in M_2, whose total prize value is $\hat{m} + \hat{m}/(\hat{m}+1)$, but then they incur a cost of at least $\beta \cdot (\hat{m}+2) \geq \hat{m}+2 > \hat{m} + \hat{m}/(\hat{m}+1)$, so again the player can improve her utility by setting $b_{1,\ell} = 0$ for each $\ell \in K'$. Hence, when playing a best response, the player produces binary outputs (either 0 or 1) along each activity. In other words, she selects a subset of activities from K and produces an output of 1 for each.

Without loss of generality, we can also assume that the player will select an action (subset of activities) that wins her all the contests in M_2. For contradiction, say the player selects an action $A \subset K$ such that there is a contest $(j+\hat{k}) \in M_2$ that the player does not win. As per our assumption that every element of the universe $\mathcal{U}$ in the 4-Regular-SetCover instance is covered by at least one set in $\mathcal{S}$, so there is an $\ell \in [\hat{k}]$ such that $j \in S_\ell$. $\ell \notin A$ because the player does not win contest $j+\hat{k}$. Now, if the player selects action $A' = \{\ell\} \cup A$, i.e., produces an output of 1 along activity ℓ in addition to the ones in A, then her prize increases by $(\hat{m}+1)$ because of contest $\ell \in M_1$ and by $1 + 1/(\hat{m}+1)$ because of contest $(j+\hat{k}) \in M_2$, while her cost increases by $(\hat{m}+2)$. So, her utility for playing A' is strictly more than A.

In summary, given any action by a player, we can can convert this action to another action where the player produces an output of 1 for a subset of activities $A \subseteq K$, produces an output of 0 for activities in $K \setminus A$, and wins all the contests in M_2. And, this conversion increases the player's utility and can be done in polynomial time. Further, this new action directly corresponds to a set cover for the 4-Regular-SetCover instance.

We claim that the player cannot compute an action $A \subseteq K$ in polynomial time that gives her a utility of at least $(3+\epsilon)/4$ of the optimal utility, for any $\epsilon > 0$. As per our discussion above, w.l.o.g., A is a set cover, i.e., $\cup_{\ell \in A} S_\ell = \mathcal{U}$. Let $A^* \subseteq [\hat{k}]$ be an optimal set cover. The utility of the player for action A is $\hat{m} + \frac{\hat{m}}{\hat{m}+1} - |A|$ while the utility for action A^* is $\hat{m} + \frac{\hat{m}}{\hat{m}+1} - |A^*|$. If the utility for playing A is at least $(3+\epsilon)/4$ of the optimal utility, then

$$\hat{m} + \frac{\hat{m}}{\hat{m}+1} - |A| \geq \frac{3+\epsilon}{4}\left(\hat{m} + \frac{\hat{m}}{\hat{m}+1} - |A^*|\right)$$
$$\implies \frac{1-\epsilon}{4}\left(\hat{m} + \frac{\hat{m}}{\hat{m}+1}\right) + \frac{3+\epsilon}{4}|A^*| \geq |A|. \quad (7.7)$$

Now, as $|S| = 4$ for every $S \in \mathcal{S}$, therefore

$$|A^*| \geq \frac{|\mathcal{U}|}{4} = \frac{\hat{m}}{4} \implies 4|A^*| \geq \hat{m} \implies 5|A^*| \geq \hat{m} + 1 \geq \hat{m} + \frac{\hat{m}}{\hat{m}+1}.$$

Putting this in (7.7) gives us

$$\frac{1-\epsilon}{4} 5|A^*| + \frac{3+\epsilon}{4}|A^*| \geq |A| \implies |A| \leq (2-\epsilon)|A^*|.$$

So, A is a $2 - \epsilon$ approximation of the optimal set cover A^* for the arbitrary 4-Regular-SetCover instance, which is not possible unless $P = NP$ [59].

□

Lemma 7.15. *In the linear multi-activity budget model, there is no polynomial-time approximation scheme for best-response, unless* P = NP.

Proof of Lemma 7.15. We prove this result by providing a reduction from MAX-2-SAT-3. In an instance of the MAX-2-SAT-3 problem, we are given a 2-CNF formula where each variable occurs at most 3 times, and the objective is to find an assignment of the variables to satisfy as many clauses of the 2-CNF formula as possible. It is known that it is NP-hard to approximate this problem beyond a constant factor [14].

Consider an instance of MAX-2-SAT-3 with $\hat{n}$ variables $x_1, x_2, \ldots, x_{\hat{n}}$. Let there be $\hat{m}$ clauses in the 2-CNF formula; w.l.o.g. assume that $\hat{m} \geq \hat{n}$. We construct an instance of the multi-activity budget game with one player, $k = 2\hat{n}$ activities, and $m = \hat{m} + 4\hat{n}$ contests. The player has a budget of $\hat{n}$ that she can allocate across any of the $2\hat{n}$ activities. Each contest awards a prize of 1.

- Activities: Corresponding to each variable in the 2-CNF formula, we have two activities in the game. For variable x_i, we have activity i with output b_i, corresponding to the literal x_i, and we have activity $\hat{n} + i$ with output $b_{\hat{n}+i}$, corresponding to the literal $\neg x_i$. We sometimes denote $b_{\hat{n}+i}$ by b'_i for a cleaner presentation.

- Contests: The $m = \hat{m} + 4\hat{n}$ contests are of two types:

    1. Contests M_1. There are $\hat{m}$ contests of this type, and these contests correspond to the clauses in the 2-CNF formula. For a clause $x_i \wedge x_j$, we have a contest with condition $b_i + b_j \geq 1$. Similarly, for a clause $x_i \wedge \neg x_j$, there is a contest with condition $b_i + b'_j \geq 1$, and so on for the other types of clauses. A contest in M_1 can be identified by the associated pair of activities $(i, j) \in [2\hat{n}] \times [2\hat{n}]$.

2. Contests M_2. There are $4\hat{n}$ contests of this type, and they correspond to the $2\hat{n}$ possible literals in the 2-CNF formula. For each $i \in [\hat{n}]$, there are two contests with condition $b_i \geq 0$ and two with condition $b'_i \geq 0$.

As each contest has a prize of 1, the objective of the player is to satisfy as many contests as possible.

Let us assume that for linear multi-activity budget games, we have a polynomial-time algorithm that approximates the best response by a factor $\gamma \in (0, 1]$. Let $\boldsymbol{b}$ be a solution for the instance of the problem we just constructed, found in polynomial time using this approximation algorithm.

From this approximate solution $\boldsymbol{b}$, we construct a $\{0, 1\}$-integral approximate solution with at least as good an approximation ratio as $\boldsymbol{b}$, as described below:

1. If there is an $i \in [2\hat{n}]$ such that $b_i > 1$, set $b_i = 1$. Notice that this transformation does not affect the number of satisfied contests. In the subsequent transformations of $\boldsymbol{b}$, we will not touch these b_i values (where $b_i = 1$) and the contests that contain them. So, the contests where they show up will remain satisfied. Let these contests be denoted by $M^{(1)}$.

2. Let $M_1^{(2)}$ denote the contests in $M_1 \setminus M^{(1)}$ that are satisfied by b. Do the following updates to $\boldsymbol{b}$ and $M_1^{(2)}$ until $M_1^{(2)} = \emptyset$:

   - Take an arbitrary contest (i, j) in $M_1^{(2)}$. We know that the following conditions hold: $b_i + b_j \geq 1$, $b_i < 1$, and $b_j < 1$.
   - Set $b_i = 1$ and $b_j = 0$. As the value of b_i increases, the contests that contain b_i and were satisfied before this update remain satisfied after this update. On the other hand, as every variable may occur in at most three clauses, b_j may occur in at most two more contests in $M_1^{(2)}$ in addition to the current contest (i, j). These contests become unsatisfied as a result of this transformation, but they are compensated by the two contests in M_2 that correspond to b_i, which were unsatisfied before this update, but become satisfied after this update.
   - Remove all contests from $M_1^{(2)}$ that correspond to the activities i or j, i.e., set $M_1^{(2)} = M_1^{(2)} \setminus \{(\ell, t) \mid \{\ell, t\} \cap \{i, j\} \neq \emptyset\}$.

   As the size of $M_1^{(2)}$ decreases after each update without decreasing the total number of satisfied contests, the process terminates with a solution as good as before.

3. As a result of the previous step, there is no contest in M_1 that is satisfied but has an output strictly less than 1 for both of its activities. In other words, for any contest $(i, j) \in M_1$, $b_i + b_j \geq 1 \implies \max(b_i, b_j) = 1$. Moreover, none of the contests in M_2 can be satisfied by an activity i with $b_i < 1$. So, we can safely set $b_i = 0$ for all i with $b_i < 1$, without decreasing the number of satisfied contests.

Thus, we can assume without loss of generality that the solution for the instance of the linear multi-activity budget game is $\{0, 1\}$-integral. We can also assume that $\sum_{i \in 2\hat{n}} b_i = \hat{n}$, because setting additional b_i values to 1 does not decrease the objective.

As mentioned before, MAX-2-SAT-3 is NP-hard to approximate beyond a constant factor, even when the number of clauses is of the same order as the number of variables, i.e., $\hat{m} = \Theta(\hat{n})$, as proved by Berman and Karpinski [14], given in the theorem below.

Theorem 7.16. *[14] For any* $0 < \epsilon < 1/2$*, it is* NP*-hard to decide whether an instance of* MAX-2-SAT-3 *with* $\hat{m} = \mu_1 \hat{n}$ *clauses has a truth assignment that satisfies* $(\mu_2 + 1 - \epsilon)\hat{n}$ *clauses, or every truth assignment satisfies at most* $(\mu_2 + \epsilon)\hat{n}$ *clauses, where* μ_1 *and* μ_2 *are two positive constants.*[7]

Theorem 7.16 proves that MAX-2-SAT-3 cannot be approximated better than a factor $(\mu_2 + 1)/\mu_2 - \epsilon$. Now, if an instance of the MAX-2-SAT-3 problem has a truth assignment that satisfies $\gamma\hat{n}$ clauses, then the constructed instance of the linear budget game has a best response that satisfies $(\gamma + 4)\hat{n}$ contests. So, for any $0 < \epsilon < 1/2$, it is NP-hard to decide whether an instance of the linear budget game with $m = \hat{m} + 4\hat{n} = (\mu_1 + 4)\hat{n}$ contests has a best response that satisfies $(\mu_2 + 5 - \epsilon)\hat{n}$ contests, or it can be at most $(\mu_2 + 4 + \epsilon)\hat{n}$. This gives us an approximation ratio of $(\mu_2 + 5)/(\mu_2 + 4) - \epsilon$.

□

□

Proof of Theorem 7.11. We prove the result first for the case when the number of activities is a constant and then for the case when the number of players is a constant.

Constant number of activities (k constant). Notice that in the linear multi-activity model, the criterion for a given contest j is a linear constraint of the form $\sum_{\ell \in K} w_{j,\ell} b_{i,\ell} \geq 1$, where $w_{j,\ell} \in \mathbb{R}_{\geq 0}$ is a non-negative weight that contest j has for

[7] $\mu_1 = 2016$ and $\mu_2 = 2011$. Berman and Karpinski [14] call the problem 3-OCC-MAX-2SAT instead of MAX-2-SAT-3.

activity ℓ and $\boldsymbol{b}_i = (b_{i,\ell})_{\ell \in K}$ is the output vector for a player. This linear constraint divides $\mathbb{R}^k$ into two half-spaces. The set of half-spaces in $\mathbb{R}^k$ has a VC-dimension of $k+1$ [65], and therefore, the set of constraints (one for each contests) also has a VC-dimension of at most $k+1$. Let $\mathcal{S}$ be the following set of subsets of contests:

$$\mathcal{S} = \{S \subseteq M \mid \exists \boldsymbol{b}_i \in \mathbb{R}^k_{\geq 0} \text{ s.t. } \sum_{\ell \in K} w_{j,\ell} b_{i,\ell} \geq 1, \forall j \in S, \\ \text{and} \sum_{\ell \in K} w_{j,\ell} b_{i,\ell} < 1, \forall j \in S\}.$$

In other words, for every set $S \in \mathcal{S}$, there exists an action $\boldsymbol{b}_i$ which satisfies the criteria for all the contests in S and none of the contests not in S. As the set of criteria corresponding to the contests has a VC-dimension of $k+1$, the set $\mathcal{S}$ is of size $O(m^{k+1})$ and the elements of $\mathcal{S}$ can be enumerated in time polynomial in m when k is fixed.[8] The elements of $\mathcal{S}$ can be computed inductively, see the proof of Lemma 3.1 of Kearns and Vazirani [65].

Now, let us pick a set $S \in \mathcal{S}$. For the cost model, we can efficiently solve a linear program where the objective is to minimize the (linear) cost function of the player and the constraints are the ones corresponding to the contests in set S. Similarly, for the budget model, we need to check the feasibility of the set of constraints for contests in S plus the budget constraint of the player, which can also be done efficiently. Finally, we compile the optimal solution for each $S \in \mathcal{S}$ to find the optimal solution overall.

Constant number of contests (m constant). The proof is similar to the one for the case when k is a constant. There are 2^m possible subsets of contests, we can efficiently enumerate over them because m is a constant. For each subset of contests $S \subseteq M$, we repeat the same process (solve a linear programming problem) as the previous case when k was a constant. □

7.6.3 From Section 7.4

Proof of Theorem 7.12. For now, let us focus on the single-activity budget model. At the end, we extend the result to the cost model.

In the proof sketch given in Section 7.4, we reduced a 3-Polytensor game to a single-activity budget game to emphasize the main ideas and give a cleaner presentation, but 3-Polytensor games have not yet been proven to be (PPAD∩PLS)-complete. Rather, 5-Polytensor games are (PPAD∩PLS)-complete [9], so we start from an instance of 5-Polytensor.

[8] See Lemma 3.1 and 3.2 of Chapter 3 of the book by Kearns and Vazirani [65].

Take an arbitrary instance of 5-Polytensor with n players; we shall use the same notation as Definition 7.10. We construct a single-activity game with n players, $\sum_{i\in[n]} |A_i|$ activities, and a polynomial number of contests to be defined later.

The $\sum_{i\in[n]} |A_i|$ activities have a one-to-one association with the actions of the players. The activities are partitioned into n subsets; the i-th subset has size $|A_i|$ and is associated with player i; we identify these activities by A_i. Player i has a budget of 1 that they can use to produce output along any activity from A_i, but they have 0 budget for the activities in A_j for $j \neq i$. Effectively, as we are in a single-activity model, player i selects an activity from the activities in A_i and produces an output of 1 along it. Note that the players have disjoint sets of activities for which they can produce outputs.

All the contests we construct are associated with exactly five players and at most five activities. We shall denote a contest by $\mathcal{C}_S(A)$, where

- S is the set of five distinct players based on whose joint utility function in the polytensor game, u_S, we shall specify the prize of contest $\mathcal{C}_S(A)$;
- the contest $\mathcal{C}_S(A)$ awards its prize to any player who produces an output of at least 1 along the activities in A;
- the activities in A are from $\cup_{i\in S} A_i$ with $|A| \leq 5$ and $|A \cap A_i| \leq 1$ for $i \in S$.

We shall call a contest $\mathcal{C}_S(A)$ a Type-ℓ contest if $|A| = \ell$.

We create the contests $(\mathcal{C}_S(A))_A$ to exactly replicate the utility that the five players in S get from u_S. If we can do this, then repeating the same process for every set of five players, we will replicate the entire 5-Polytensor game. The utility that player $i \in S$ gets from u_S is $u_S(\boldsymbol{a}_S)$, where the $\boldsymbol{a}_S$ are the actions of the five players. We have the following contests:

Type-5 Contests. Let us add a contest named $\mathcal{C}_S(\boldsymbol{a}_S)$ with prize $v_S(\boldsymbol{a}_S)$ for every $\boldsymbol{a}_S \in \times_{i\in S} A_i$. We shall specify the $v_S(\boldsymbol{a}_S)$ values based on the $u_S(\boldsymbol{a}_S)$ values, later. Contest $\mathcal{C}_S(\boldsymbol{a}_S)$ distributes its prize to players who produce output along the activities $\boldsymbol{a}_S$.

Say the players S play the actions $\boldsymbol{a}^*_S = (a^*_j)_{j\in S}$. Let $A^-_j = A_j \setminus \{a^*_j\}$; and for $T \subseteq S$, let $A_T = \times_{j\in T} A_j$ and $A^-_T = \times_{j\in T} A^-_j$. The total prize that player $i \in S$ gets from the Type-5 contests is:

$$\sum_{\ell\in[5]} \sum_{T\subseteq S, i\in T, |T|=\ell} \sum_{\boldsymbol{a}_{S\setminus T}\in A^-_{S\setminus T}} \frac{v_S(\boldsymbol{a}^*_T, \boldsymbol{a}_{S\setminus T})}{\ell}.$$

In the above formula, the outer summation with ℓ is for the number of players that player i has to share their prize with, including i; the middle summation is for the ℓ players, T; the inner summation is for the actions of the players in $S \setminus T$. The same formula can be rewritten as

$$\sum_{\boldsymbol{a}_{S\setminus\{i\}} \in A_{S\setminus\{i\}}} v_S(a_i^*, \boldsymbol{a}_{S\setminus\{i\}}) - \sum_{2\le\ell\le5} \sum_{T\subseteq S, i\in T, |T|=\ell} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \frac{\ell-1}{\ell} v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}).$$

Type-1 Contests. Let us add a contest named $\mathcal{C}_S(a_i')$ with prize $\sum_{a_i\neq a_i', \boldsymbol{a}_{S\setminus\{i\}} \in A_{S\setminus\{i\}}} v_S(a_i, \boldsymbol{a}_{S\setminus\{i\}})$ for every $i \in S$ and $a_i' \in A_i$. This contest $\mathcal{C}_S(a_i')$ awards its prize to any player who produces output along activity a_i'. The total prize that player i gets from Type-1 and Type-5 contests is

$$\sum_{\boldsymbol{a}_S \in A_S} v_S(\boldsymbol{a}_S) - \sum_{2\le\ell\le5} \sum_{T\subseteq S, i\in T, |T|=\ell} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \frac{\ell-1}{\ell} v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}).$$

As $\sum_{\boldsymbol{a}_S \in A_S} v_S(\boldsymbol{a}_S)$ does not depend upon the action a_i^* selected by i, the utility of player i is effectively

$$-\sum_{2\le\ell\le5} \sum_{T\subseteq S, i\in T, |T|=\ell} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \frac{\ell-1}{\ell} v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}).$$

Type-2 Contests. Let us add a contest named $\mathcal{C}_S(a_i', a_j')$ with prize $\sum_{\boldsymbol{a}_{S\setminus\{i,j\}} \in A_{S\setminus\{i,j\}}} \frac{v_S(a_i', a_j', \boldsymbol{a}_{S\setminus\{i,j\}})}{2}$ for every $i, j \in S, i \neq j$ and $a_i' \in A_i$ and $a_j' \in A_j$. This contest $\mathcal{C}_S(a_i', a_j')$ awards its prize to any player who produces output along activity a_i' or a_j'. The net utility of player i gets from Type-5, Type-1, and Type-2 contests is

$$\begin{aligned}
&\sum_{3\le\ell\le5} \sum_{T\subseteq S, i\in T, |T|=\ell} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \left(\frac{1}{2} \frac{4.\binom{3}{\ell-2}}{\binom{4}{\ell-1}} - \frac{\ell-1}{\ell} \right) v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}) \\
&\qquad = \sum_{T\subseteq S, i\in T, |T|=3} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \frac{1}{3} v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}) \\
&\qquad\qquad + \sum_{T\subseteq S, i\in T, |T|=4} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \frac{3}{4} v_S(\boldsymbol{a}_T^*, \boldsymbol{a}_{S\setminus T}) + \frac{6}{5} v_S(\boldsymbol{a}_S^*).
\end{aligned}$$

Type-3 Contests. Let us add a contest named $\mathcal{C}_S(\boldsymbol{a}_T')$ with prize $\sum_{\boldsymbol{a}_T\neq\boldsymbol{a}_T', \boldsymbol{a}_{S\setminus T} \in A_{S\setminus T}} \frac{v_S(\boldsymbol{a}_T, \boldsymbol{a}_{S\setminus T})}{3}$ (effectively, the prize is $-\sum_{\boldsymbol{a}_{S\setminus T} \in A_{S\setminus T}} \frac{v_S(\boldsymbol{a}_T', \boldsymbol{a}_{S\setminus T})}{3}$) for every $T \subset S$ with $|T| = 3$ and for every $\boldsymbol{a}_T' \in A_T$. This contest $\mathcal{C}_S(\boldsymbol{a}_T')$ awards its prize to any player who produces output along activities $\boldsymbol{a}'_T$. The net utility of player i gets from Type-5, Type-1, Type-2, and Type-3 contests is

$$\sum_{T \subseteq S, i \in T, |T|=3} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \left(\frac{1}{3} - \frac{1}{3}\right) v_S(\boldsymbol{a}^*_T, \boldsymbol{a}_{S\setminus T})$$
$$+ \sum_{T \subseteq S, i \in T, |T|=4} \sum_{\boldsymbol{a}_{S\setminus T} \in A^-_{S\setminus T}} \left(\frac{3}{4} - 1\right) v_S(\boldsymbol{a}^*_T, \boldsymbol{a}_{S\setminus T}) + \left(\frac{6}{5} - 2\right) v_S(\boldsymbol{a}^*_S)$$
$$= -\sum_{j \neq i} \sum_{a_j \in A^-_j} \frac{1}{4} v_S(\boldsymbol{a}^*_{S\setminus\{j\}}, a_j) - \frac{4}{5} v_S(\boldsymbol{a}^*_S).$$

Type-4 Contests. Let us add a contest named $\mathcal{C}_S(\boldsymbol{a}'_T)$ with prize $\sum_{\boldsymbol{a}_{S\setminus T} \in A_{S\setminus T}} \frac{v_S(\boldsymbol{a}'_T, \boldsymbol{a}_{S\setminus T})}{4}$ for every $T \subset S$ with $|T| = 4$ and for every $\boldsymbol{a}'_T \in A_T$. This contest $\mathcal{C}_S(\boldsymbol{a}'_T)$ awards its prize to any player who produces output along activities $\boldsymbol{a}'_T$. The net utility of player i gets from all the five types of contests is

$$\frac{1}{5} v_S(\boldsymbol{a}^*_S).$$

We set $v_S(\boldsymbol{a}_S) = 5u_S(\boldsymbol{a}_S)$ for every $\boldsymbol{a}_S \in A_S$, and we are done.

Cost. It is easy to adapt the above proof to the single-activity cost model. All the players, contests, and activities are the same, we just need to replace the budgets with costs. In the proof above, we used the budgets to make sure that player i can produce output only along activities in A_i (note that we are in a single-activity model so they will produce output along exactly one activity). We can achieve that same for the cost model: for player i, set a very low cost for the activities in A_i and very high cost for the activities in A_j, $j \neq i$, such that player i always wants to produce output along one of the activities in A_i and never along the activities not in A_i.

□

Proof of Theorem 7.13. We prove this result by reducing the problem of finding a locally optimal max-cut solution in a graph, which is known to be a PLS-complete problem [64], to our problem. Let (V, E) be an arbitrary graph and let $w_e \in \mathbb{Z}_{\geq 0}$ be the weight of edge $e \in E$. A cut $(S, \overline{S})$ (where $\overline{S} = V \setminus S$) is locally optimal if one cannot improve the max-cut objective $\sum_{(u,v) \in E, u \in S, v \in \overline{S}} w_{u,v}$ by moving a vertex u from its current partition S (or $\overline{S}$) to the other partition $\overline{S}$ (or S).

Budget. Let us first focus on the single-activity budget model, we shall slightly tweak the same construction for the single-activity cost model.

Construction. We construct an instance of the budget game corresponding to the given graph (V, E) as follows:

1. Let there be $|V|$ players.

2. Let there be $2|V|$ activities, which we identify by α_v and β_v for $v \in V$.

3. Let each player have a budget of 1 that they can spend on any of the activities. As we are in a single-activity setup, in a pure strategy, a player would select an activity and produce an output of 1 along that activity. Note that the strategy space is symmetric across the players.

4. Let there be $|V| + 2|E|$ contests.

   - We identify the first $|V|$ contests using the vertices. The contest that corresponds to $v \in V$ equally distributes a large prize of $4\sum_{e \in E} w_e + 2$ to the players who produce a total output of at least 1 along the two activities α_v and β_v corresponding to v, i.e., a player i receives a share of the prize if $b_{i,\alpha_v} + b_{i,\beta_v} \geq 1$. We call these contests *vertex* contests.
   - We identify the other $2|E|$ contests using the edges, two contests per edge. For an edge $e = (u, v) \in E$, one of the contests distributes a prize of w_e to the players who produce an output of at least 1 along the two activities α_u and α_v, and the other contest distributes a prize of same value w_e to the players who produce an output of at least 1 along β_u and β_v. We call these contests *edge* contests.

Analysis. Observe that the total prize awarded by the $2|E|$ edge contests is $2(\sum_{e \in E} w_e)$, while the prize awarded by a single vertex contest is $4\sum_{e \in E} w_e + 2$, which is much larger. As a player can produce output along only one activity, and as each vertex contest v values a unique disjoint pair of activities (α_v, β_v), no player can win prizes from more than one vertex contest. Also, as the number of players is equal to the number of the vertex contests, each player can get a complete share of one of the vertex contests. In a PNE, each player selects a distinct pair (α_v, β_v) and produces output along one of the activities in the pair, because a player can move from a shared prize to an unshared prize increasing their utility by at least $2\sum_{e \in E} w_e + 1$ from the vertex contests, which dominates maximum total loss of $2\sum_{e \in E} w_e$ from the $2|E|$ edge contests. Let us denote the player who produces output along the pair of activities (α_v, β_v) as i_v.

Let $S = \{v \in V \mid b_{i_v,\alpha_v} = 1\}$ and $\overline{S} = V \setminus S = \{v \in V \mid b_{i_v,\beta_v} = 1\}$. For an edge (u, v), if both players i_u and i_v are in S or $\overline{S}$, then they share the prize of $w_{u,v}$, on the other hand if one of them is in S and the other in $\overline{S}$ then they each get $w_{u,v}$. The value of the potential function (Equation (7.1)) ϕ can be written using S and $\overline{S}$ as

$$\begin{aligned}\phi(S,\overline{S}) &= \left(4\sum_{e\in E} w_e + 2\right)|V| + \left(\frac{3}{2}\sum_{e\in E} w_e + \frac{1}{2}\sum_{(u,v)\in E, u\in S, v\in\overline{S}} w_{u,v}\right)\\ &= \left(4|V| + \frac{3}{2}\right)\sum_{e\in E} w_e + 2|V| + \frac{1}{2}\sum_{(u,v)\in E, u\in S, v\in\overline{S}} w_{u,v},\end{aligned}$$

where the contribution of $|V|(4\sum_{e\in E} w_e + 2)$ is from the $|V|$ vertex contests and a contribution of $3w_{u,v}/2$ (if both u and v are in S or $\overline{S}$) or $2w_{u,v}$ (if one of u or v is in S and the other in $\overline{S}$) is from the two contests corresponding to the edge (u,v). As we have argued before, a player i_v will never change their associated vertex v in a better response move. A move by the player i_v of shifting from α_v to β_v, or vice-versa, corresponds to moving the vertex v across the cut $(S,\overline{S})$ in the max-cut problem. The value of the potential is effectively equal to the objective of the max-cut problem: we have multiplied the max-cut objective with a positive constant factor of $1/2$ and added a constant term. So, we have a one-one correspondence between local search in max-cut and dynamics in the single-activity budget game.

Cost. Note that for the budget game constructed above, the budget constraint was not very crucial in the analysis. No player had an incentive to produce an output of strictly more than 1 along the activity they selected, because all the contests already awarded the prizes as soon as the players produced an output of 1. In the cost model, we remove the budget constraint and set $c_{i,\ell} = 4\sum_{e\in E} w_e + 2$ for every player i and activity ℓ. The same analysis used in the budget model holds here too, the value of the potential function (Equation (7.1)) here is:

$$\begin{aligned}\phi(S,\overline{S}) &= \left(4\sum_{e\in E} w_e + 2\right)|V|\\ &\quad + \left(\frac{3}{2}\sum_{e\in E} w_e + \frac{1}{2}\sum_{(u,v)\in E, u\in S, v\in\overline{S}} w_{u,v}\right) - \left(4\sum_{e\in E} w_e + 2\right)|V|\\ &= \frac{3}{2}\sum_{e\in E} w_e + \frac{1}{2}\sum_{(u,v)\in E, u\in S, v\in\overline{S}} w_{u,v},\end{aligned}$$

where the positive contribution of $|V|(4\sum_{e\in E} w_e + 2)$ is from the $|V|$ vertex contests, the contribution of $3w_{u,v}/2$ or $2w_{u,v}$ is from the two contests corresponding to the edge (u,v), and a negative contribution of $|V|(4\sum_{e\in E} w_e + 2)$ is due to the cost of the players. The rest of the argument follows as before.

Mixed-strategy Nash equilibrium. In the MAX CUT problem, a randomized solution where each vertex is on either side of the cut with a probability of $1/2$ cannot be improved by moving just one vertex to a different side (or to a different distribution over the two sides). So, it is locally stable. We apply the same idea to find an MNE in the single-activity contest game just constructed.

Consider the following mixed strategy profile: for every $v \in V$, there is a unique player i_v who produces an output of 1 along one of the two activities α_v or β_v with a probability of $1/2$ each. As argued before, player i_v would not want to produce output along any other α_u or β_u for $u \neq v$ because of the vertex contests. Also, given that every other player is following this strategy, player i_v has the same utility for selecting either α_v or β_v, so $(1/2, 1/2)$ is as good a mixed-strategy as any other. So, we have an equilibrium. $\square$

7.6.3.1 Fixed Parameter Tractability for Single-Activity Models

Theorem 7.17. *In the single-activity models, both cost and budget, we can compute a pure-strategy Nash equilibrium in polynomial time if the number of players is a constant.*

Proof of Theorem 7.17. For the budget model, a player essentially chooses an activity among the k activities; once she selects an activity, she can w.l.o.g. produce the maximum output in her budget for that activity. So, she essentially has k actions.

Similarly, for the cost model, she essentially has at most $m \cdot k$ actions. First, the player chooses an activity among the k activities, then she chooses the level of output for that activity. For a given activity, the criterion of a contest corresponds to a minimum level of output. When choosing the level of output for an activity, a player can w.l.o.g. choose among these minimum levels of output corresponding to the contests. As she increases her output for an activity from one level to another level with a higher output, she wins a superset of contests as she was winning previously.

As there are n players and each player (effectively) has at most $m \cdot k$ actions ($k \leq m \cdot k$ for budget model), therefore there are at most $(m \cdot k)^n$ action profiles. We can enumerate over these profiles and find the PNE is polynomial time if n is a constant. $\square$

Theorem 7.18. *In the single-activity budget model, we can compute a pure-strategy Nash equilibrium in polynomial time if the number of contests is a constant.*

Proof of Theorem 7.18. If we carefully observe the potential function of the budget model, $\sum_{j \in M} \sum_{\ell \in [n_j(\boldsymbol{b})]} v_j(\ell)$, given in (7.1), we can see that for every contest $j \in M$, $\sum_{\ell \in [n_j(\boldsymbol{b})]} v_j(\ell)$ can take at most $n+1$ different values depending upon the value of $n_j(\boldsymbol{b})$. So, the potential function can take at most $(n+1)^m$ different values.

If we are not at a PNE, there is a best-response move that strictly increases the potential function, and we can find such a best-response in polynomial time (even for the multi-activity model, see Theorem 7.11). So, we can find a PNE in polynomial time by repeatedly finding a best response move and strictly increasing the potential function, and this process will terminate in at most $(n+1)^m$ many steps. $\square$

Note that idea used in the proof of Theorem 7.18 does not directly carry over to the cost model because of the additional $\left(-\sum_{i \in N} c_i(\boldsymbol{b}_i)\right)$ term in the potential function for the cost model.

Chapter 8

Conclusion and Future Work

Contest theory provides useful models for studying competitive environments in economics and computer science. Our work contributed to specific directions focusing on learning in Tullock contests, equilibrium in simultaneous contests, and contest design.

Our work on best-response and related dynamics in Tullock contests assumes that the agents observe and react to the actions of other agents. This assumption is realistic in specific applications, e.g., firms competing for market share through advertising, but may not hold for applications like Bitcoin mining. In the latter case, the contest allocates the prize probabilistically, and an agent gets the prize with a probability proportional to their action. In each round of the game, an agent only observes whether they win the round or not, but not the actions of the other agents. An important research direction is to study learning in these games with only probabilistic feedback.

Our work on learning in Tullock contests focuses only on convex cost functions. However, some applications may have economies of scale, and the marginal costs may decrease, including computer science applications like blockchain protocols [8]. There has been little attention to learning in Tullock contests with non-convex cost functions, likely because Tullock contests with non-convex cost functions (non-concave utility functions) may not have a pure-strategy Nash equilibrium and the learning dynamics are less well-behaved. For example, best-response dynamics is not a suitable learning model in these settings and does not converge (including the empirical distribution of the actions generated by the dynamics). Dynamics that can converge empirically to mixed-strategy Nash equilibria, e.g., fictitious play, are more suitable for these models. Although these research directions are more challenging, pursuing them is crucial for a better understanding of dynamics and equilibrium in Tullock contests. Similarly, learning in all-pay auctions, and more generally, in rank-order allocation contests, has received limited attention due to similar challenges but is also an important direction of future research.

Another promising direction is to extend the techniques from Chapters 2 and 3 to other classes of games. In particular, use techniques from randomized algorithms and convex optimization to derive (rate of) convergence results for discrete learning dynamics in specific classes of games, especially games that have a unique pure-strategy Nash equilibrium, e.g., diagonally strictly concave games [93].

Our work in Chapter 4 showed that continuous-time BR dynamics converges to an ϵ-approximate equilibrium in $\Theta(\log(1/\epsilon))$ time for Tullock contests with arbitrary, possibly non-homogeneous, convex costs. However, with discretization, we proved only an $O(1/\epsilon)$ convergence rate. Can we discretize the continuous dynamics differently to get an algorithm that computes the equilibrium in $poly(\log(1/\epsilon))$ steps? More broadly, we want to understand which properties of a dynamical system allow us to find computationally efficient discretized algorithms and when it is not possible (computationally hard), with a particular focus on learning in games, e.g., extending convergence results of continuous gradient-based dynamics of [93] to discretized versions. Such results are not immediate, as indicated by the difference between discrete and continuous replicator dynamics [52]. This research direction is also related to the active research area in optimization that addresses the continuous-to-discrete challenge, e.g., of choosing the step size to have fast convergence of gradient descent [3, 4].

Contests are receiving increased attention in computer science because of their use in modeling crowdsourcing and blockchain applications. Blockchains naturally lead to contest-like games (including the Bitcoin mining application discussed previously in this thesis). For example, blockchains need to implement some form of decentralized governance. Generally, in proof-of-stake blockchains, an agent gets voting rights in proportion to the amount of their stake that gets blocked, which leads to opportunity costs. Such voting games are similar to *group* contests [97]. As another example, consider the blockchain platform Cardano that uses crowdsourcing to get quality reviews for proposals to make changes to the platform [18]. Cardano incentivizes reviewers by providing rewards for quality reviews, which leads to a game of simultaneous contests. Current literature, including our work in Chapter 7, provides a limited understanding of group and simultaneous contests, primarily due to the difficulty in analyzing these contests. We may be able to make progress in these directions by formulating them as large contests [87], where each agent has only a minimal influence on the outcome.

Recent papers have used contests to study competition for college seats [21, 68, 27]. The data analysis in these papers indicates that agents may not follow the traditional contest models. Proposed intuitive explanations include participants not exactly knowing their own valuations or not fully following rational behavior (within the scope of the modeling parameters). The pressing research problem here is to develop suitable contest models guided by real-life data and use these models to recommend policy interventions that help mitigate the practical issues (e.g., with the JEE college admission exam in India [27]). Tools from behavioral economics may play an important role. A related line of work is to study fairness and diversity objectives, particularly when applied to contests induced by scholarships, funding, and other education-related competitions like hackathons; our work in Chapters 5 and 6 makes progress along these directions, but a lot remains to explore.

Contests are practical models to study both classical economic environments and modern strategic interactions enabled by the internet. Our work contributes to a better understanding of these models and highlights important directions of future research, but a lot remains to explore.

Bibliography

[1] Tim Adamo and Alexander Matros. A Blotto game with incomplete information. *Economics Letters*, 105(1):100–102, 2009.

[2] Laurence Ales, Soo-Haeng Cho, and Ersin Körpeoğlu. Optimal award scheme in innovation tournaments. *Operations Research*, 65(3):693–702, 2017.

[3] Jason M. Altschuler and Pablo A. Parrilo. Acceleration by stepsize hedging i: Multi-step descent and the silver stepsize schedule, 2023.

[4] Jason M. Altschuler and Pablo A. Parrilo. Acceleration by stepsize hedging ii: Silver stepsize schedule for smooth convex optimization, 2023.

[5] Erwin Amann and Wolfgang Leininger. Asymmetric all-pay auctions with incomplete information: the two-player case. *Games and Economic Behavior*, 14 (1):1–18, 1996.

[6] Erwin Amann and Heng Qiao. Parallel contests. Technical report, Thurgauer Wirtschaftsinstitut, University of Konstanz, 2008.

[7] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the 13th International Conference on Information Systems*, page 200, 01 2009.

[8] Nick Arnosti and S Matthew Weinberg. Bitcoin: A natural oligopoly. *Management Science*, 68(7):4755–4771, 2022.

[9] Yakov Babichenko and Aviad Rubinstein. Settling the complexity of nash equilibrium in congestion games. In *Proceedings of the 53rd ACM Symposium on Theory of Computing*, pages 1426–1437, 2021.

[10] Ravi Bapna, Chrysanthos Dellarocas, and Sarah Rice. Vertically differentiated simultaneous Vickrey auctions: Theory and experimental evidence. *Management Science*, 56(7):1074–1092, 2010.

[11] Yasar Barut and Dan Kovenock. The symmetric multiple prize all-pay auction with complete information. *European Journal of Political Economy*, 14(4):627–644, 1998.

[12] Michael R Baye, Dan Kovenock, and Casper G De Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, 1996.

[13] Richard Bellman. On "colonel Blotto" and analogous games. *SIAM Review*, 11 (1):66–68, 1969.

[14] Piotr Berman and Marek Karpinski. On some tighter inapproximability results. In *Proceedings of the 26th International Colloquium on Automata, Languages, and Programming*, pages 200–209. Springer, 1999.

[15] Dimitri Bertsekas. *Convex Optimization Algorithms*. Athena Scientific, Nashua, NH, USA, 2015.

[16] Dimitri P Bertsekas and John N Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.

[17] Vittorio Bilò, Laurent Gourvès, and Jérôme Monnot. Project games. In *Proceedings of the 11th International Conference on Algorithms and Complexity*, pages 75–86. Springer, 2019.

[18] Georgios Birmpas, Lyudmila Kovalchuk, Philip Lazos, and Roman Oliynykov. Parallel contests for crowdsourcing reviews: Existence and quality of equilibria. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, pages 268—-280, 2023.

[19] Donald W Blackett. Some Blotto games. *Naval Research Logistics Quarterly*, 1(1):55–60, 1954.

[20] Donald W Blackett. Pure strategy solutions of Blotto games. *Naval Research Logistics Quarterly*, 5(2):107–109, 1958.

[21] Aaron L Bodoh-Creed and Brent R Hickman. College assignment as a large contest. *Journal of Economic Theory*, 175:88–126, 2018.

[22] E Borel and J Ville. Application de la théorie des probabilitiés aux jeux de hasard, Gauthier-Villars; reprinted in Borel, E., & Chéron, A. (1991). *Théorie mathematique du bridge à la portée de tous, Editions Jacques Gabay*, 1938.

[23] Emile Borel. La théorie du jeu et les équations intégrales à noyau symétrique gauche. *Comptes rendus de l'Académie des Sciences*, 173:1304–1308, 1921.

[24] Renaud Bourlès, Yann Bramoullé, and Eduardo Perez-Richet. Altruism in networks. *Econometrica*, 85(2):675–689, 2017.

[25] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.

[26] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180, 1996.

[27] L Elisa Celis, Amit Kumar, Anay Mehrotra, and Nisheeth K Vishnoi. Bias in evaluation processes: An optimization-based model. In *Proceedings of the 37th Conference on Neural Information Processing Systems*, volume 36, 2023.

[28] Shuchi Chawla and Jason D Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 181–196, New York, USA, 2013. ACM.

[29] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. *Games and Economic Behavior*, 113:80–96, 2019.

[30] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):1–57, 2009.

[31] Xi Chen, Christos Papadimitriou, and Tim Roughgarden. An axiomatic approach to block rewards. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 124–131, New York, NY, USA, 2019. ACM.

[32] Yun Kuen Cheung, Stefanos Leonardos, and Georgios Piliouras. Learning in markets: Greed leads to chaos but following the price is right. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pages 111–117, 2021.

[33] Subhasish M Chowdhury, Patricia Esteve-González, and Anwesha Mukherjee. Heterogeneity, leveling the playing field, and affirmative action in contests. *Southern Economic Journal*, 89(3):924–974, 2023.

[34] Ani Dasgupta and Kofi O Nti. Designing an optimal contest. *European Journal of Political Economy*, 14(4):587–603, 1998.

[35] Partha Dasgupta. The theory of technological competition. In *New Developments in the Analysis of Market Structure: Proceedings of a Conference Held by the International Economic Association in Ottawa, Canada*, pages 519–549. Springer, 1986.

[36] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39 (1):195–259, 2009.

[37] Shanglyu Deng, Qiang Fu, and Zenan Wu. Optimally biased Tullock contests. *Journal of Mathematical Economics*, 92:10–21, 2021.

[38] Robert Deschamps. An algorithm of game theory applied to the duopoly problem. *European Economic Review*, 6(2):187–194, 1975.

[39] Dominic DiPalantino and Milan Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 119–128, 2009.

[40] Avinash K Dixit et al. Strategic behavior in contests. *American Economic Review*, 77(5):891–898, 1987.

[41] Davide Dragone, Luca Lambertini, and Arsen Palestini. Static and dynamic best-response potential functions for the non-linear cournot game. *Optimization*, 61(11):1283–1293, 2012.

[42] Pradeep Dubey, Ori Haimanko, and Andriy Zapechelnyuk. Strategic complements and substitutes, and potential games. *Games and Economic Behavior*, 54(1):77–94, 2006.

[43] Edith Elkind, Abheek Ghosh, and Paul Goldberg. Contest design with threshold objectives. In *Proceedings of the 17th Conference on Web and Internet Economics*, page 554, 2021. arXiv:2109.03179.

[44] Christian Ewerhart. The lottery contest is a best-response potential game. *Economics Letters*, 155:168–171, 2017.

[45] Christian Ewerhart. Ordinal potentials in smooth games. *Economic Theory*, 70 (4):1069–1100, 2020.

[46] Christian Ewerhart and Federico Quartieri. Unique equilibrium in contests with incomplete information. *Economic Theory*, 70(1):243–271, 2020.

[47] Christian Ewerhart and Kremena Valkanova. Fictitious play in networks. *Games and Economic Behavior*, 123:182–206, 2020.

[48] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th ACM Symposium on Theory of Computing*, pages 604–612, 2004.

[49] John Fearnley, Paul Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: Cls= ppad∩pls. *Journal of the ACM*, 70(1): 1–74, 2022.

[50] Michal Feldman, Kevin Lai, and Li Zhang. A price-anticipating resource allocation mechanism for distributed shared clusters. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 127–136, 2005.

[51] Qiang Fu and Zenan Wu. On the optimal design of biased contests. *Theoretical Economics*, 15(4):1435–1470, 2020.

[52] Drew Fudenberg and David K Levine. *The Theory of Learning in Games*, volume 2. MIT Press, Cambridge, MA, USA, 1998.

[53] Arieh Gavious and Yizhaq Minchuk. Revenue in contests with many participants. *Operations Research Letters*, 42(2):119–122, 2014.

[54] Amihai Glazer and Refael Hassin. Optimal contests. *Economic Inquiry*, 26(1): 133–143, 1988.

[55] Oliver Gross and Robert Wagner. A continuous colonel Blotto game. Technical report, RAND PROJECT AIR FORCE SANTA MONICA CA, 1950.

[56] Marina Halac, Navin Kartik, and Qingmin Liu. Contests for experimentation. *Journal of Political Economy*, 125(5):1523–1569, 2017.

[57] Sergiu Hart. Discrete colonel Blotto and general lotto games. *International Journal of Game Theory*, 36(3):441–460, 2008.

[58] Josef Hofbauer and William H Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002.

[59] Jonas Holmerin. Vertex cover on 4-regular hyper-graphs is hard to approximate within 2-ε. In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pages 544–552, 2002.

[60] Rafael Hortala-Vallve and Aniol Llorente-Saguer. Pure strategy Nash equilibria in non-zero sum colonel Blotto games. *International Journal of Game Theory*, 41(2):331–343, 2012.

[61] Zhengao Huang. Fictitious play for games with a continuum of strategies. Technical report, State University of New York at Stony Brook, Stony Brook, NY, USA, 2002.

[62] Nicole Immorlica, Greg Stoddard, and Vasilis Syrgkanis. Social status and badge design. In *Proceedings of the 24th International Conference on World Wide Web*, pages 473–483, 2015.

[63] Martin Kaae Jensen. Aggregative games and best-reply potentials. *Economic Theory*, 43(1):45–66, 2010.

[64] David S Johnson, Christos H Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.

[65] Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory.* MIT Press, Cambridge, MA, USA, 1994.

[66] René Kirkegaard. Favoritism in asymmetric contests: Head starts and handicaps. *Games and Economic Behavior*, 76(1):226–248, 2012.

[67] Kai A Konrad. *Strategy and Dynamics in Contests.* Oxford University Press, Oxford, UK, 2009.

[68] Kala Krishna, Sergey Lychagin, Wojciech Olszewski, Ron Siegel, and Chloe Tergiman. Pareto improvements in the contest for college admissions. Technical report, National Bureau of Economic Research, 2022.

[69] Nikolai S Kukushkin. Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48(1):94–110, 2004.

[70] Karim R Lakhani, David A Garvin, and Eric Lonstein. Topcoder (a): Developing software through crowdsourcing. Technical Report 610-032, Harvard Business School General Management Unit Case, 2010.

[71] Alice LaPlante. Microsoft News. Bridging the STEM divide: AI hackathon helps young women excel in computer science. `https://bit.ly/microsoft-hackathon`, May 2021. Accessed: 2022-09-19.

[72] Jacob D Leshno and Philipp Strack. Bitcoin: An axiomatic approach and an impossibility theorem. *American Economic Review: Insights*, 2(3):269–86, 2020.

[73] Qiang Liu. Stein variational gradient descent as gradient flow. In *Proceedings of the 27th Conference on Neural Information Processing Systems*, volume 30, 2017.

[74] Glenn C Loury. Market structure and innovation. *The Quarterly Journal of Economics*, 93(3):395–410, 1979.

[75] Steven Matthews. On the implementability of reduced form auctions. *Econometrica*, 52(6):1519–22, 1984.

[76] Avner May, Augustin Chaintreau, Nitish Korula, and Silvio Lattanzi. Filter & follow: How social media foster content curation. In *Proceedings of the 2014 ACM Conference on Measurement and Modeling of Computer Systems*, pages 43–55, 2014.

[77] Jürgen Mihm and Jochen Schlapp. Sourcing innovation: On feedback in contests. *Management science*, 65(2):559–576, 2019.

[78] Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

[79] Dylan Minor. Increasing effort through rewarding the best less. Technical report, Northwestern University, 2011.

[80] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis.* Cambridge University Press, Cambridge, UK, 2017.

[81] Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001.

[82] Benny Moldovanu and Aner Sela. Contest architecture. *Journal of Economic Theory*, 126(1):70–96, 2006.

[83] Dov Monderer and Lloyd S Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[84] Hervé Moulin and J P Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7:201–221, 1978.

[85] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[86] Yurii Nesterov et al. *Lectures on Convex Optimization*, volume 137. Springer, 2018.

[87] Wojciech Olszewski and Ron Siegel. Large contests. *Econometrica*, 84(2):835–854, 2016.

[88] Wojciech Olszewski and Ron Siegel. Performance-maximizing large contests. *Theoretical Economics*, 15(1):57–88, 2020.

[89] Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48 (3):498–532, 1994.

[90] Tönu Puu. Chaos in duopoly pricing. *Chaos, Solitons & Fractals*, 1(6):573–581, 1991.

[91] Dénes Pálvölgyi, Hans Peters, and Dries Vermeulen. A strategic approach to multiple estate division problems. *Games and Economic Behavior*, 88:135–152, 2014.

[92] Brian Roberson. The colonel Blotto game. *Economic Theory*, 29(1):1–24, 2006.

[93] J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, pages 520–534, 1965.

[94] Alejandro A Schäffer. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1):56–87, 1991.

[95] Jeff S Shamma and Gürdal Arslan. Unified convergence proofs of continuous-time fictitious play. *IEEE Transactions on Automatic Control*, 49(7):1137–1141, 2004.

[96] LS Shapley. Stochastic games. In *Proceedings of the National Academy of Sciences of the United States of America*, pages 1095–1100. JSTOR, 1953.

[97] Roman M Sheremeta. Behavior in group contests: A review of experimental research. *Journal of Economic Surveys*, 32(3):683–704, 2018.

[98] Martin Shubik and Robert James Weber. Systems defense games: Colonel Blotto, command and control. *Naval Research Logistics Quarterly*, 28(2):281–287, 1981.

[99] Ron Siegel. All-pay contests. *Econometrica*, 77(1):71–92, 2009.

[100] Margaret Slade. What does an oligopoly maximize? *Journal of Industrial Economics*, 42(1):45–61, 1994.

[101] Curtis R Taylor. Digging for golden carrots: An analysis of research tournaments. *The American Economic Review*, 85(4):872–890, 1995.

[102] Lars Thorlund-Petersen. Iterative computation of cournot equilibrium. *Games and Economic Behavior*, 2(1):61–75, 1990.

[103] John W Tukey. A problem of strategy. *Econometrica*, 17(1):73, 1949.

[104] Gordon Tullock. Efficient rent-seeking. In James Buchanan, Robert Tollison, and Gordon Tullock, editors, *Toward a Theory of the Rent Seeking Society*, pages 131–146. Texas A & M University, College Station, Texas, USA, 1980.

[105] Adrian Vetta. 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*, pages 416–425, 2002.

[106] Milan Vojnović. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press, Cambridge, UK, 2016.

[107] Mark Voorneveld. Best-response potential games. *Economics letters*, 66(3): 289–295, 2000.

[108] Karl Wärneryd. Chaotic dynamics in contests. *Economic Inquiry*, 56(3):1486–1491, 2018.

[109] Alan Washburn. Blotto politics. *Operations Research*, 61(3):532–543, 2013.

[110] Stephen J Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.

[111] Jiang Yang, Lada A Adamic, and Mark S Ackerman. Crowdsourcing and knowledge sharing: strategic user behavior on taskcn. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 246–255, 2008.

[112] Li Zhang. The efficiency and fairness of a fixed budget resource allocation game. In *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming*, pages 485–496. Springer, 2005.