Four Essays in Economic Theory

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Christian Seel
aus Troisdorf

Bonn 2011

Dekan:	Prof. Dr. Klaus Sandmann
Erstreferent:	Prof. Dr. Tymon Tatur
Zweitreferent:	Prof. Dr. Benny Moldovanu

Tag der mündlichen Prüfung: 04.10.2011

Acknowledgements

In preparing this thesis, I benefited a lot from the support of many people to whom I am grateful.

First of all, I want to thank my supervisor Tymon Tatur. He granted me complete freedom in choosing topics for each chapter of this dissertation and provided a myriad of constructive comments and suggestions. Moreover, I learned a lot from countless discussions with him about economic phenomena and possible ways to model them. It was a pleasure to have him as a supervisor.

Secondly, I want to thank Benny Moldovanu, who acts as a referee in the thesis committee. He provided several valuable comments about the first two chapters of this dissertation. Moreover, he gave me guidance on many general issues such as which journals to target or career planning in academics, and enabled me to meet many interesting seminar speakers.

I am particularly grateful to Philipp Strack. He is a good friend and coauthor of the first three chapters of this thesis. I always enjoyed to work with him, but also to spent many entertaining hours of leisure time together. I have learned a lot from the way he tends to approach economic problems, which was mostly different to mine.

I am grateful to my other two coauthors, Matthias Lang and Philipp Wichardt. The discussions with Matthias while writing Chapter 3 of this thesis also helped me improve the first two chapters. To Philipp, I am especially grateful for forcing me to work on my writing.

I owe many thanks to Paul Heidhues for his critique and comments about previous versions of the first three chapters and many helpful discussions. In the same way, I owe many thanks to Anja Schöttner for her topics course, which inspired the third chapter of this dissertation.

Moreover, this work profited from discussions with and comments from Rafael Aigner, Stefan Ankirchner, Dirk Bergemann, Jörg Bewersdorff, Norbert Christopeit, Martin Cripps, Eduardo Faingold, Fabrizio Germano, Julio González-Díaz, Faruk Gul, Martin Hellwig, Sergei Izmalkov, Philippe Jehiel, Navin Kartik, Christian Kellner, Georg Kirchsteiger, Eugen Kovac, Daniel Krähmer, Matthias Kräkel, Sebastian Kranz, Stephan Lauermann, Gábor Lugosi, Moritz Meyer-ter-Vehn, Konrad Mierendorff, Frank Riedel, Urs Schweizer, Avner Shaked, Xianwen Shi, Ron Siegel, Deszö Szalay, Nora Szech, Christoph Wagner, an anonymous associate editor, three anonymous referees, and many other people. I wish to thank Andrea, Angela, Guy, John, Linda, and Matthias for pointing out typos and syntax errors in earlier versions of the chapters in this dissertation.

The Bonn Graduate School of Economics provided financial support, for which I am grateful. Moreover, both my classmates and the administration of the BGSE, namely Urs Schweizer, Silke Kinzig, Corinna Lehmann, and Pamela Mertens created an excellent environment for economic research.

Finally, I am greatly indebted to my family, my girlfriend Eva, and my friends. Throughout the last years, they offered me invaluable support, even when I was moody. Especially Eva and my parents helped me a lot to stay focused on the dissertation.

Contents

List of Figures

Introduction

This thesis consists of four chapters. The first three chapters form an entity and the core of this dissertation. The last chapter is a note, which is not related to the others.

More specifically, chapters one to three analyze models of contests. Contests and tournaments are widespread mechanisms in many different areas of the real world. For example, they occur in sports, politics, patent races, relative reward schemes in firms, or (public) procurement. In contests, participants are incentivized by the possibility to win prizes. The winning probability of a contestant depends on her performance relative to other contestants. This structure enables the principal to commit on paying out rewards based on (relative) performance at the end of the competition, even if performances are not verifiable in court.

In this thesis, we analyze contests in which n players compete for one prize. The success of a player—in absolute terms—depends on the realization of a stochastic process in continuous time. Each player has the same stochastic process (but different realizations with probability one) and there is no correlation between the processes. A strategy of a player specifies a stopping time for his process. Hence, it measures how long he is active in the contest, e.g., exerts effort. The player who stops his process at the highest value wins the prize.

The key difference of the models in this thesis compared to most of the previous literature is a different observability assumption. Most of the literature analyzes one of two polar cases. Either players can perfectly observe of each others contest success throughout the competition (see, e.g., Harris and Vickers, 1987, Moscarini and Smith, 2007) or they do not learn anything about the contest success over time (see, e.g., Lazear and Rosen, 1981, Park and Smith, 2008, Siegel, 2009). In contrast to this, we assume that each player observes his own stochastic research success over time and can adjust his strategy accordingly, but does not receive any information about the progress of his competitors. Typical examples illustrating this setting are R&D contests for an innovation or procurement contests.

The difference in modeling compared to the full information case has implications on the choice of equilibrium concept. In particular, we do not need to consider any refinement of Nash equilibrium, since no new information about the success of the rivals arrives over time. However, the existence of Nash equilibria is not obvious in our framework, since the games have infinite strategy spaces and discontinuous payoffs. In fact, showing existence and uniqueness for these games is one of the main technical contributions of this thesis.

In the first chapter, which is based on joint work with Philipp Strack, we study risk-taking behavior in contests. To focus on the risk-taking behavior, we deviate from previous literature by abstracting from effort cost. More precisely, we analyze a contest model in which each player can decide when to stop a Brownian motion with (usually negative) drift. A player has to stop in case of bankruptcy, i.e., at the first time his process hits zero.

The equilibrium construction first characterizes the unique candidate for an equilibrium distribution over final values of the stopped process. To verify equilibrium existence, we then apply a result from probability theory by Pedersen and Peskir (2001) to show that there exist a stopping time which induces the equilibrium distribution. Moreover, we explicitly derive a corresponding strategy for the two-player case.

In equilibrium, agents do not stop immediately even if the drift is negative, since they maximize their expected winning probability rather than their expected value. As it turns out, the expected value of the equilibrium distribution of an agent is non-monotone in drift and variance. Hence, the principal incurs highest expected losses in the natural case in which the drift is only slightly negative. Potential applications of the model include competition between managers of private equity funds, competition in declining industries, and optimal strategies for roulette tournaments.

The second chapter, which is based on joint work with Philipp Strack, introduces flow costs of continuation to the setting of the first chapter. Moreover, differing from the first chapter, we now assume that the drift is positive and abstract from the bankruptcy constraint. Hence, the analysis in this chapter is a lot more in the spirit of the contest literature; natural applications of the model include R&D and procurement contests.

Imposing mild assumptions on the cost function, we prove existence and uniqueness of the Nash equilibrium outcome with similar techniques as in the first chapter. In addition, we apply a recent mathematical result from Ankirchner and Strack (2011) to construct a bounded time stopping strategy—a strategy which stops almost surely before a fixed time T—which induces the equilibrium distribution. From a technical point of view, this introduces a method to construct equilibria in continuous time games that are independent of the time horizon given the horizon is long enough. This result also reinforces the economic relevance of the model, since most real-world contests end within bounded time.

We then discuss the relation of our model to static all-pay contests. As the variance converges to zero, the equilibrium distribution of the game converges to that of the symmetric equilibrium of an all-pay auction. The implication of this result is twofold. First, it provides an equilibrium selection argument in favor of the symmetric equilibrium in the symmetric all-pay auction discussed in Baye et al. (1996). Second, the result supports the validity of all-pay models to analyze contests in which the variance is negligible.

For positive variance, each participant makes positive profits. Intuitively, he generates rents through the private information about his research progress. In the case of two players and constant costs, the profits of each player increase as each player's costs increase, variance increases, or productivity decreases. Thus, according to the model, participants prefer to have mutually worse technologies or to take part in contests whose outcome is more random. This finding, which cannot be obtained in an all-pay contest, goes along with the common intuition that competitors prefer competition to be less fierce.

In the third chapter, which is based on joint work with Matthias Lang and Philipp Strack, we scrutinize a contest model with a simpler, weakly increasing technology. More precisely, as long as a player exerts costly effort, i.e., does not stop the process, he receives draws according to a Poisson process. At a fixed deadline T, the player who has accumulated most draws wins the prize.

In this setting, we find that if the deadline is sufficiently high, the set of equilibrium distributions of the contest is equivalent to that of a single-prize common-value all-pay auction with discrete bids. We derive an explicit lower bound on the required contest length until the long-run equilibrium distribution is feasible. Hence, the all-pay auction is also a suitable model to analyze stochastic contests in a setting with a purely additive research structure and a sufficiently long time horizon. In the end, we briefly discuss how the equilibrium set of the game changes if the duration of the contest is low.

The last chapter of this dissertation, which is based on joint work with Philipp Wichardt, refines the valuation equilibrium concept originally introduced by Jehiel and Samet (2007). In this concept, (boundedly rational) players group different moves in a game tree into similarity classes. They attach a valuation, i.e., a real number, to each similarity class. Whenever a player has to make decision, she chooses only moves from similarity classes with the (locally) highest valuation. In equilibrium, valuations are confirmed by equilibrium outcomes.

Instead of keeping the grouping of moves into similarity classes entirely exogenous as Jehiel and Samet (2007) do, we start with an intuitive basic grouping, which players can refine at a lexicographic cost. The modified equilibrium concept takes the trade-off between a more costly grouping and its potential benefits into account. Hence, it makes the valuation concept applicable to games at hand without having to specify the grouping ad hoc.

We apply the modified concept to a burning money game. It predicts that adding a possibility of burning money, i.e., publicly harm oneself, has no influence on the equilibrium set of a coordination game. This result, which differs from standard solution concepts like subgame perfect equilibrium or forward induction, is roughly in line with empirical evidence for the game gathered by Huck and Müller (2005). Moreover, it highlights the high degree of rationality entailed in standard solution concepts for this game.

Chapter 1

Gambling in Contests

This chapter presents a strategic model of risk-taking behavior in contests. Formally, we analyze an n-player winner-take-all contest in which each player decides when to stop a privately observed Brownian motion with drift. A player whose process reaches zero has to stop. The player with the highest stopping point wins.

We derive a closed-form solution of the unique Nash equilibrium outcome of the game. Contrary to the explicit cost for a higher stopping time in a war of attrition, here, higher stopping times are riskier, because players can go bankrupt. In equilibrium, the trade-off between risk and reward causes a non-monotonicity: highest expected losses occur if the process decreases only slightly in expectation.

1.1 Introduction

To provide more excitement for the players, the (online) gambling industry introduced casino tournaments. The rules are simple: all participants pay a fixed amount of money prior to the tournament—the "buy-in"—that enters into the prize pool. In return, they receive chips, which they can invest in the casino gamble throughout the tournament. At the end of the tournament, the player who has most chips wins a prize, which is the sum of the buy-ins minus some fee charged by the organizers. Benefits are two-sided: players restrict their maximal loss to the buy-in and enjoy a new, strategic component of the game; the casino makes a sure profit through the fee it charges.

The observability of each other's chip stacks throughout the tournament depends on the provider. The no-observability case provides a good illustration of our model—in equilibrium, players use the gamble even though it has a negative expected value.[1]

In the model, each player decides when to stop a privately observed Brownian motion (X_t) with (usually negative) constant drift coefficient μ, constant diffusion coefficient σ, and initial endowment x_0. If a player becomes bankrupt, i.e., $X_t = 0$, she has to stop. The player who stops at the highest value wins a prize.

Instead of an explicit cost for a higher contest success (e.g., Lazear and Rosen, 1981, Hillman and Samet, 1987), here, higher prizes are riskier. In equilibrium, players maximize their winning probability rather than the expected value of the process. Hence, they do not stop immediately even if the underlying process is decreasing in expectation. Intuitively, if all other players stop immediately, it is better for the remaining player to play until she wins a small amount or goes bankrupt, since she can ensure to win an arbitrarily small amount with a probability arbitrarily close to one.

In the unique equilibrium outcome, expected losses are non-monotonic in the expected value of the gamble—a more favorable gamble can lead to higher expected losses. Intuitively, this results from the trade-off between risk and reward: if the gamble has only a slightly negative expected value, the relatively high probability of winning makes people stop later, which increases expected losses. If the principal—who might have imperfect information about the drift—obtains wins or losses of the players, contests are not a reliable compensation scheme, because even with a slightly negative drift, the principal incurs a large loss.

[1] Several online casinos use a leaderboard for the chip stacks. In most cases, however, it updates with a delay to create more tension. In this variant, players should only play close to the end of the contest to veil their realizations.

The formal analysis proceeds as follows. Proposition 1 derives a necessary formula for an implied stopping chance $F(x)$ in the symmetric equilibrium of an n-player game that pinpoints the unique candidate equilibrium distribution. To do so, we exploit that each player has to be indifferent whether to stop or to continue at any point of her support at any point in time.

For the two-player case, Proposition 2 derives the equilibrium stopping time that induces $F(x)$ explicitly. It involves mixing whether to stop with a chance that depends on the current state X_t. Proposition 4 extends Proposition 1 and 2 to a two-player game with asymmetric starting values.

For more than two players, Proposition 3 ensures the existence of a stopping time that induces $F(x)$. Its proof relies on a result in probability theory on the Skorokhod embedding problem. This literature—initiated by Skorokhod (1961, 1965)—analyzes under which conditions a stopping time of a stochastic process exists that embeds, i.e., induces, a given probability distribution; for an excellent survey article, see Obłój (2004). In the proof of Proposition 3, we verify a sufficient condition from Pedersen and Peskir (2001). This whole approach is new to game theory and the main technical contribution of this chapter.

Proposition 5 provides the main characterization result: the general shape of the expected value of the stopped processes is quasi-convex, falling, then rising in drift μ and variance σ. In particular, highest expected losses occur if the process decreases only slightly in expectation.

Apart from casino tournaments, this chapter provides a stylized model for the following applications. First, consider a private equity fund that invests in start-up companies. The value of the fund is mostly private information until maturity, because start-ups do not trade on the stock market and the composition of the fund is often unknown. The model analyzes a competition between fund managers in which, at maturity, the best performing manager gets a prize—a bonus or a job promotion.

In this application, there are several possible reasons for a downward drift. For instance, there may be no good investment opportunities in the market. Moreover, the downward drift may capture the cost of paying an expert to search for possible investments. The model predicts that return on investment is very sensitive to the profitability of investment opportunities. In particular, a slightly negative drift is most harmful for the investors. In this case, contestants behave as if they were risk-loving, which a payment based on absolute performance could avoid.

As a second example, consider a competition in a declining industry. In a duopoly, for instance, firms compete to survive and get the monopoly profit. Fudenberg and Tirole (1986) model the situation as a war of attrition—only the firm who stays alone in the market wins a prize, but both incur costs until one firm drops out.

In an interpretation of our model, managers of both firms decide if they want to make risky investments, e.g., in R&D or stocks of other firms. Investments are costly, but could improve the firm's value. When the duopoly becomes unprofitable, the firm with the higher value wins—either by a take-over battle or because the other firm cannot compete in a prize war—and its manager keeps his job.

Our model predicts that managers choose very risky strategies. In particular, investors lose most money in expectation if investment opportunities have a slightly negative expected value, which is consistent with being in a declining industry. This effect increases in the asymmetry of the firms' values. Intuitively, to satisfy the indifference condition for the stronger firm, the weaker firm has to make up for its initial disadvantage by taking higher risks.

1.1.1 Related Literature

Hvide (2002) investigates whether tournaments lead to excessive risk-taking behavior. He modifies Lazear and Rosen (1981) by assuming that players bear costs to raise their expected value, but can raise their variance without costs. In equilibrium, they choose maximum variance and low effort. Similarly, Anderson and Cabral (2007) scrutinize an infinite competition in which two players, who observe each other, can update their binary choice of variance continuously. In their model, flow payoffs depend on the difference in contest success. In equilibrium, both players choose the risky strategy until the lead of one player is above a threshold; in this case, the leader switches to the save option.

In the literature on races, players balance a higher effort cost against a higher winning probability. Moscarini and Smith (2007)—building on a discrete time model of Harris and Vickers (1987)—analyze a two-person continuous-time race with costly effort choice. In equilibrium, effort is increasing in the lead of a player up to some threshold above which the laggard resigns; for an application to political economy, see also Gul and Pesendorfer (2011). These papers assume full observability of each other's contest success over time. In our model, however, stopping decisions and realizations of the rivals are unobservable.

Regarding the assumptions on information and payoffs, the model most resembles a silent timing game—as first explored in Karlin (1953), and most recently, in Park and Smith (2008). The latter paper also generalizes the all-pay war of attrition, and so assumes that later stopping times cost linearly more. Contrary to a silent timing game, in this chapter, players do not only possess private information about their stopping decision, but also about the realization of their stochastic process.

Finally, this chapter relates to the finance literature on gambling for resurrection, e.g., Downs and Rocke (1994). In this literature, managers take unfavorable gambles for a chance to save their firms from bankruptcy. Here, however, players take high risks to veil their contest outcomes.

We proceed as follows. Section 1.2 introduces the model. Section 1.3 derives the unique equilibrium distribution. In Section 1.4, we state the main characterization result, Proposition 5, and discuss its implications. Section 1.5 concludes.

1.2 The Model

There are n agents $i \in \{1, 2, \ldots, n\} = N$ who face a stopping problem in continuous time. At each point in time $t \in \mathbb{R}_+$, agent i privately observes the realization of a stochastic process $X^i = (X^i_t)_{t \in \mathbb{R}_+}$ with

$$X^i_t = x_0 + \mu t + \sigma B^i_t \,.$$

The constant $x_0 > 0$ denotes the starting value of all processes; see Section 1.3.3 for heterogeneous starting values. The drift $\mu \in \mathbb{R}$ is the common expected change of each process X^i_t per time, i.e., $\mathbb{E}(X^i_{t+\Delta} - X^i_t) = \mu\Delta$. The noise term is an n-dimensional Brownian motion (B_t) scaled by $\sigma \in \mathbb{R}_+$.

1.2.1 Strategies

A strategy of player i is a stopping time τ^i. This stopping time depends only on the realization of his process X^i_t, as the player only observes his own process.[2] Mathematically, the agents' stopping decision until time t has to be $\mathcal{F}^i_t$-measurable, where $\mathcal{F}^i_t = \sigma(\{X^i_s : s < t\})$ is the sigma algebra induced by the possible observations of the process X^i_s before time t. We restrict agents' strategy spaces in two ways. First, we require finite expected stopping times, i.e., $\mathbb{E}(\tau^i) < \infty$. Second, a player has to stop in case of bankruptcy. More formally, we require $\tau^i \leq \inf\{t \in \mathbb{R}_+ : X^i_t = 0\}$ a.s.. To incorporate mixed strategies, we allow for randomized stopping times—progressively measurable functions $\tau^i(\cdot)$ such that, for every $r^i \in [0, 1]$, $\tau^i(r^i)$ is a stopping time. Intuitively, agents can draw a random number r^i from the uniform distribution on $[0, 1]$ before the game and play a stopping strategy $\tau^i(r^i)$.

[2]The equilibrium of the model would be the same if the stopping decision was reversible and stopped processes were constant.

1.2.2 Payoffs

The player who stops his process at the highest value wins a prize, which we normalize to one without loss of generality. Ties are broken randomly. Formally,

$$\pi^i = \frac{1}{k}\,\mathbf{1}_{\{X^i_{\tau^i} = \max_{j\in N} X^j_{\tau^j}\}}\,,$$

where $k = |\{i \in N : X^i_{\tau^i} = \max_{j\in N} X^j_{\tau^j}\}|$. As payoffs add up to one, the game is a constant sum game. All agents maximize their expected payoff, i.e., the probability to win the contest. This optimization is independent of their attitude towards risk.

1.2.3 Condition on the Parameters

To ensure equilibrium existence in finite time stopping strategies, we henceforth impose a technical condition that places a positive upper bound on μ; for a discussion, see Section 1.3.2.

Assumption 1. $\mu < \log(1 + \frac{1}{n-1})\frac{\sigma^2}{2x_0}$.

1.3 Equilibrium Analysis

In this section, we first derive the unique candidate for an equilibrium distribution. Second, we prove equilibrium existence—this is not trivial as the game has discontinuous payoffs and infinite strategy spaces. Our proof verifies the existence of a stopping time, which induces the candidate for an equilibrium distribution. We close the section with an extension to asymmetric starting values.

1.3.1 The Equilibrium Distribution

Every strategy of agent i induces a (potentially non-smooth) cumulative distribution function (cdf) $F^i : \mathbb{R}_+ \to [0,1]$ of his stopped process, where $F^i(x) = \mathbb{P}(X^i_{\tau^i} \leq x)$.

The probability of a tie is non-zero only if the distributions of at least two agents have a mass point above zero or the distributions of all agents have a mass point at zero or both.[3] The next lemma proves otherwise.

Lemma 1. *In equilibrium, for every $x > 0$, no agent $i \in N$ has a mass point at x, i.e., $\mathbb{P}(X^i_{\tau^i} = x) = 0$. At least one agent has no mass point at zero.*

We omit the proof and present a verbal argument instead, because the proof is simply a specialization of the now standard logic in static game theory with a continuous state space; e.g., Burdett and Judd (1983). As usual, mixed strategies in a competitive game can have no interior mass point at the same point in the state space (here, the same x), since this would create a profitable deviation in one direction: With a slightly higher x, one raises one's win chance a boundedly positive probability with an arbitrarily small loss, since one beats everyone with lower x and the one player with mass at x; however, an agent can have a mass point at zero, since any other player who can pass him would have already been bankrupt.

Lemma 1 renders the tie-breaking rule obsolete, because it implies that the probability of a tie is zero. Denote the winning probability of player i if he stops at $X^i_{\tau^i} = x$ by $u^i(x)$, where $u^i(x) : \mathbb{R}_+ \to [0,1]$. As there are no mass points away from zero, we can express $u^i(x)$ in terms of the other agents' cdf's.

$$\begin{aligned} u^i(x) &= \mathbb{P}(x > \max_{j \neq i} X^j_{\tau^j}) + \frac{1}{k} \underbrace{\mathbb{P}(x = \max_{j \neq i} X^j_{\tau^j}\})}_{=0} \\ &= \prod_{j \neq i} \mathbb{P}(X^j_{\tau^j} \leq x) = \prod_{j \neq i} F^j(x) \end{aligned} \tag{1.3.1}$$

[3] As common in economic literature, we do not consider the mathematical problem of an accumulation of mass points (Cantor Construction); we thus assume that either there is only a finite number of mass points or they have no accumulation point.

We call $u^i(\cdot)$ the utility function of agent i given the distributions of the other agents. These utility functions are helpful to derive the equilibrium—a point where each player maximizes $\mathbb{E}(u^i(X^i_{\tau^i}))$.

Denote the right endpoint of the support of the distribution of player i by $\overline{x}^i = \sup\{x : F^i(x) < 1\}$ and the left endpoint by $\underline{x}^i = \inf\{x : F^i(x) > 0\}$. The right endpoint has to be finite, because agents can only use strategies that stop almost surely in finite time. The following results establish necessary conditions on u^i and the distribution functions in equilibrium; the proofs are in the appendix.

Lemma 2. *The utility u^i of every agent $i \in N$ is strictly increasing on the interval $[\underline{x}^i, \overline{x}^i]$.*

Lemma 3. *For each player i, the utility $u^i(X^i_t)$ is a local martingale on the interior of the support of his distribution, i.e., $X^i_t \in (\underline{x}^i, \overline{x}^i) \Rightarrow \mathbb{E}(\mathrm{d}u^i(X^i_t)|\mathcal{F}^i_t) = 0$.*

Lemma 4. *The support of the cdf of each player is identical and starts at zero.*

All players share the same utility function. Hence, Lemma 3 and 4 directly imply the following corollary:

Corollary 1. *The unique equilibrium distributions are atomless and symmetric.*

As the utility u^i does not depend on time ($\frac{\partial u^i}{\partial t} = 0$), by Itô's lemma (Revuz and Yor, 2005, p.147) the expected change in utility per marginal unit of time is

$$\begin{aligned}\mathbb{E}(\mathrm{d}u^i(X^i_t)|\mathcal{F}^i_t) &= \mathbb{E}\left((\mu u^{i\prime}(X^i_t) + \frac{\sigma^2}{2}u^{i\prime\prime}(X^i_t))\mathrm{d}t + u^{i\prime}(X^i_t)\sigma \mathrm{d}B_t|\mathcal{F}^i_t\right) \\ &= \mu u^{i\prime}(X^i_t) + \frac{\sigma^2}{2}u^{i\prime\prime}(X^i_t)\mathrm{d}t\,.\end{aligned}$$

By Lemma 3, this equation is equal to zero for all x on the support of F^i, which yields the following ordinary differential equation:

$$0 = \mu u^{i\prime}(x) + \frac{\sigma^2}{2}u^{i\prime\prime}(x)\,.$$

For $\mu \neq 0$, all solutions to this equation are of the form $u^i(x) = \alpha + \beta\exp(\frac{-2\mu x}{\sigma^2})$ for all constants $\alpha, \beta \in \mathbb{R}$. To fix α and β, we use two constraints on u^i. First, all players win with probability $\frac{1}{n}$ in equilibrium (Corollary 1). In particular, they do so when they stop immediately (Lemma 3). Second, the value of the cdf at zero is zero, because the support is atomless (Corollary 1). Thus, we get:

$$\begin{aligned} \frac{1}{n} = u^i(x_0) &= \alpha + \beta \ \exp(\frac{-2\mu x_0}{\sigma^2}) \\ 0 = u^i(0) &= \alpha + \beta\,. \end{aligned}$$

This system of equations uniquely determines α and β, and thereby also u^i as

$$u^i(x) = \min\left\{1, \frac{1}{n}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}\right\}.$$

It remains to construct the corresponding equilibrium distributions. For this purpose, we insert the symmetry property of the equilibrium (Corollary 1) into equation (1.3.1) to get

$$u^i(x) = \prod_{j\neq i}^{n} F^j(x) = F(x)^{n-1} \Rightarrow F(x) = \sqrt[n-1]{u^i(x)}\,.$$

Hence, we characterize the unique candidate for an equilibrium distribution as follows (for an illustration, see Figure 1):

Proposition 1. *Assume $\mu \neq 0$. A strategy profile is a Nash equilibrium, if and only if each player's strategy induces the cumulative distribution function*

$$F(x) = \min\left\{1, \sqrt[n-1]{\frac{1}{n}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}}\right\}.$$

Proof. We have already proven that any equilibrium strategy is symmetric and induces the distribution F. Assumption 1 ensures that there exists a finite x such that $F(x) = 1$; see Section 4.2 for details. To complete the proof, we need to show that no deviation gives a player a winning probability greater than $\frac{1}{n}$. Recall that, by construction of the function $F(\cdot)$, the process $(u^i(X_t^i))_{t\in\mathbb{R}_+}$ is a supermartingale. For every stopping time $\tau < \infty$, consider the sequence of bounded stopping times $\min\{\tau, n\}$ for $n \in \mathbb{N}$. By Doob's optional stopping theorem (Revuz and Yor, 2005, p.70), $\mathbb{E}(u^i(X^i_{\min\{\tau,n\}})) \leq u^i(X_0^i)$. As

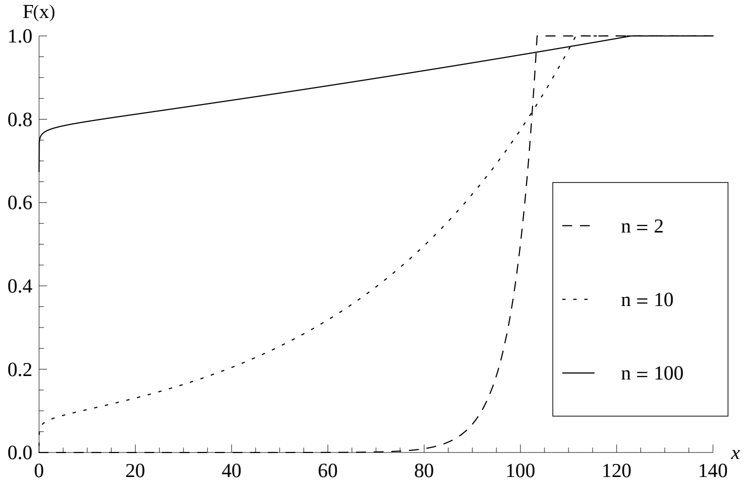

Figure 1.1: An example ($\mu = -0.1$, $x_0 = 100$, $\sigma = 1$) of the equilibrium cdf's for different sizes of players n.

$u(X_t^i) \in [0, 1]$ is bounded, we can apply the dominated convergence theorem to get

$$\mathbb{E}(u^i(X_\tau^i)) = \mathbb{E}(\lim_{n\to\infty} u^i(X_{\min\{\tau,n\}}^i)) = \lim_{n\to\infty} \mathbb{E}(u^i(X_{\min\{\tau,n\}}^i)) \leq u^i(X_0^i) = \frac{1}{n}\,.$$

□

To complete the analysis, we scrutinize the special case in which X_t^i is a martingale, i.e., $\mu = 0$. In this case, the first term in the differential equation vanishes. The same calculation as in the case $\mu \neq 0$ yields the unique equilibrium distribution, where

$$F(x) = \min\left\{1,\ \sqrt[n-1]{\frac{x}{nx_0}}\right\}\,.$$

$F(x)$ is continuous in μ at $\mu = 0$, because the same formula follows by taking limits in Proposition 1, using the approximation $e^A = 1 + A + O(A^2)$ for small A.

1.3.2 Equilibrium Strategies

So far, we have been silent about the existence of a finite time stopping strategy τ inducing the equilibrium distribution F. For a given distribution to be implementable in finite time stopping strategies, its right endpoint has to be finite. Recall that

$$1 = F(\overline{x}) = \sqrt[n-1]{\frac{1}{n}\frac{\exp(\frac{-2\mu\overline{x}}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}}\,.$$

Hence, the right endpoint $\overline{x}$ satisfies

$$\overline{x} = -\frac{\sigma^2}{2\mu}\log(n(\exp(\frac{-2\mu x_0}{\sigma^2}) - 1) + 1)\,.$$

Consequently, the right endpoint is finite if and only if $\mu < -\log(1 - \frac{1}{n})\frac{\sigma^2}{2x_0}$, i.e., Assumption 1 holds; otherwise, no equilibrium in finite time stopping strategies exists. Intuitively, if the drift becomes too large, for every point x, the strategy, which stops only at 0 and x, reaches x with a probability higher than $\frac{1}{n}$; this strategy with $x \geq \overline{x}^i$ would thus be a profitable deviation for any equilibrium candidate with a finite right endpoint.

In the next step, we derive strategies inducing the distribution F in the two-player case to convey the main intuition. The construction uses a mixture of deterministic threshold strategies. To formalize this intuition, we introduce the martingale transformation
$\phi : \mathbb{R}_+ \to \mathbb{R}_+$, where

$$\phi(x) = \frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}\,.$$

By Itô's lemma (since $\phi''/\phi' = -2\mu/\sigma^2$), the process $(\phi(X_t^i))_{t\in\mathbb{R}_+}$ is a martingale. In this case, $F(x) = \phi(x)/2$.

Proposition 2. *If agent i randomly selects a number $\alpha \in (0, 1]$ from a uniform distribution and stops if*

$$\tau^i = \inf\{t : |\phi(X_t^i) - 1| \geq \alpha\}\,,$$

then the cumulative distribution function induced by this strategy equals F, i.e., $\mathbb{P}(X_{\tau^i}^i \leq x) = F(x)$.

Proof. By the martingale property of $(\phi(X_t^i))_{t\in\mathbb{R}_+}$, we get

$$\mathbb{P}(\phi(X_{\tau^i}^i) = 1-\alpha) = \mathbb{P}(\phi(X_{\tau^i}^i) = 1+\alpha) = \frac{1}{2}\,.$$

As α is uniformly distributed on $(0,1]$ and agent i stops iff $\phi(X_t^i) = 1 \pm \alpha$, the random variable $\phi(X_{\tau^i}^i)$ is uniformly distributed on $[0,2]$. It follows that

$$\mathbb{P}(X_{\tau^i}^i \leq x) = \mathbb{P}(\phi(X_{\tau^i}^i) \leq \phi(x)) = \frac{\phi(x)}{2} = F(x)\,.$$

□

For more than two players, the feasibility proof requires an auxiliary result from probability theory on the Skorokhod embedding problem. This literature studies whether a distribution is feasible by stopping a stochastic process; in their terminology, there exists an embedding of a probability distribution in the process. Skorokhod (1961, 1965) analyzes the problem of embedding in Brownian motion without drift. In a recent contribution, Pedersen and Peskir (2001) derive a necessary and sufficient condition for general non-singular diffusions. They define the scale function $S(\cdot)$ by

$$S(x) = \int_0^x \exp(-2\int_0^u \frac{\mu(r)}{\sigma(r)}dr)du = -\frac{\sigma^2}{2\mu}(\exp(\frac{-2\mu x}{\sigma^2}) - 1)\,.$$

Lemma 5 (Pedersen and Peskir, 2001, Theorem 2.1.)**.** *Let (X_t) be a non-singular diffusion on $\mathbb{R}$ starting at zero, let $S(\cdot)$ denote its scale function satisfying $S(0) = 0$, and let ν be a probability measure on $\mathbb{R}$ satisfying $|S(x)|\nu(\mathrm{d}x) < \infty$. Set $m = \int_{\mathbb{R}} S(x)\nu(\mathrm{d}x)$. Then there exists a stopping time τ_* for (X_t) such that $X_{\tau_*} \sim \nu$ if and only if one of the following four cases holds:*
(i) $S(-\infty) = -\infty$ and $S(\infty) = \infty$;
(ii) $S(-\infty) = -\infty, S(\infty) < \infty$ and $m \geq 0$;
(iii) $S(-\infty) > -\infty, S(\infty) = \infty$ and $m \leq 0$;
(iv) $S(-\infty) > -\infty, S(\infty) < \infty$ and $m = 0$.

Hence, to prove feasibility for our distribution F, it suffices to show $m = 0$.

Proposition 3. *There exists a stopping strategy inducing the distribution $F(\cdot)$ from Proposition 1.*

Proof. To verify the condition in Pedersen and Peskir (2001), we need a process which starts in zero. Thus, we consider the process $\tilde{X}_t = X_t - X_0$. After some transformations, we get $S(x - x_0) = -\frac{\sigma^2}{2\mu}(1 - \exp(\frac{2\mu x_0}{\sigma^2}))(\phi(x) - 1)$. This gives us

$$\begin{aligned} m &= \int_{\mathbb{R}} S(x - x_0) f(x) \mathrm{d}x \\ &= -\frac{\sigma^2}{2\mu}\left(1 - \exp(\frac{2\mu x_0}{\sigma^2})\right)\left(\int_{\mathbb{R}} f(x)\phi(x)\mathrm{d}x - 1\right). \end{aligned}$$

Consequently, it remains to show $\int_{\mathbb{R}} f(x)\phi(x)\mathrm{d}x = 1$.

$$\begin{aligned} \int_{\mathbb{R}} f(x)\phi(x)\mathrm{d}x &= \int_0^{\overline{x}} \underbrace{\frac{(n^{-\frac{1}{n-1}})}{n-1}\phi(x)^{-\frac{n-2}{n-1}}\phi'(x)}_{f(x)}\phi(x)\mathrm{d}x \\ &= \int_{\phi(0)}^{\phi(\overline{x})} \frac{(n^{-\frac{1}{n-1}})}{n-1} y^{\frac{1}{n-1}}\mathrm{d}y \\ &= \left[\frac{(n^{-\frac{1}{n-1}})}{n} y^{\frac{n}{n-1}}\right]_{y=\phi(0)=0}^{y=\phi(\overline{x})=n} \\ &= 1. \end{aligned}$$

As $m = 0$, there exists an embedding for the distribution F by Theorem 2.1. in Pedersen and Peskir (2001).[4] □

Proposition 1 and 3 combined yield F as the unique equilibrium distribution of the game.

[4] An alternative proof of Proposition 3 would verify a result on embedding in Brownian with drift from Grandits and Falkner (2000) for the process $\tilde{X}_t = \frac{X_t - X_0}{\sigma}$.

1.3.3 An Extension: Asymmetric Starting Values

In this extension, we allow for heterogeneous starting values. To get an analytical solution, we restrict attention to the two-player case—without loss of generality $x_0^1 > x_0^2$. The proof of the following proposition is similar to the proof of Proposition 1.

Proposition 4. *In equilibrium, the cdf of the first player is*

$$F^1(x) = \min\left\{1, \frac{1}{2}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0^1}{\sigma^2}) - 1}\right\}.$$

The cdf of the second player is

$$F^2(x) = \min\left\{1, \rho + (1-\rho)\frac{1}{2}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0^1}{\sigma^2}) - 1}\right\}.$$

Proof. The cdf of player 1 is the same as in the symmetric case. Thus, it is feasible by Proposition 2. For player 2, consider the following strategy: First, play until $X_t^2 \in \{0, x_0^1\}$; if x_0^1 is reached, use the same stopping strategy as player 1. This induces the above cdf, where the constant ρ—probability of absorption in 0—fulfills

$$\rho = \frac{\exp(\frac{-2\mu(x_0^1 - x_0^2)}{\sigma^2}) - 1}{\exp(\frac{-2\mu(x_0^1 - x_0^2)}{\sigma^2}) - \exp(\frac{2\mu x_0^2}{\sigma^2})}.$$

As in the proof of Proposition 1, the expected winning probability for each player in the above equilibrium candidate is the same as if he stops immediately. Furthermore, as $u^i(X_t^i)$ is a supermartingale and a local martingale on the support by construction, the same reasoning as in the proof of Proposition 1 applies. Hence, no player can do better than to stop immediately, which yields the equilibrium payoff. We show uniqueness of the equilibrium in the appendix. □

Compared to the symmetric case, the player with the lower starting value takes more risks here. In particular, he loses everything with probability ρ and takes the same gamble as player 1 with probability $1-\rho$. Asymmetry in the contest leads to higher percentage losses for a negative drift, because the handicapped player takes higher risks to offset his initial disadvantage.

1.4 Comparative Statics

This section analyzes how changes in the parameters affect the expected value of the stopped processes. To determine the expected value, we first calculate the density from the cdf in Proposition 1:

$$f(x) = \frac{2\mu}{n(n-1)\sigma^2} \sqrt[\frac{2-n}{n-1}]{\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{n(\exp(\frac{-2\mu x_0}{\sigma^2}) - 1)}} \frac{\exp(\frac{-2\mu x}{\sigma^2})}{1 - \exp(\frac{-2\mu x_0}{\sigma^2})} .$$

In what follows, we restrict attention to the two-player case for tractability; in the appendix, we state the formula for the expected value for n players. We use the density f to derive the expected value of the stopped processes for two players:

$$\begin{aligned}
\mathbb{E}(X_\tau) &= \mathbb{E}_f(x) = \int_0^{\overline{x}} x f(x) \mathrm{d}x \\
&= \frac{\sigma^2}{2\mu} + (1 + \frac{1}{2(\exp(-\frac{2\mu x_0}{\sigma^2}) - 1)})(x_0 - \frac{\sigma^2 \log(2 - \exp(\frac{2\mu x_0}{\sigma^2}))}{2\mu}) .
\end{aligned}$$

The explicit formula of the expected value allows us to characterize its shape in the following proposition—the proof is in the appendix.

Proposition 5. $\mathbb{E}(X_\tau)$ *is quasi-convex, falling, then rising in* μ*. If* $\mu < 0$*,* $\mathbb{E}(X_\tau)$ *is quasi-convex, falling, then rising in* σ*.*

Hence, an increase in the drift does not imply an increase in the expected value of the stopped processes. Intuitively, for $\mu < 0$, there are two opposing effects: an increase in the drift lowers the expected losses per time but increases the expected stopping time. Similarly, as the variance increases, the gamble gets more attractive, but it also takes less time to implement the equilibrium distribution.

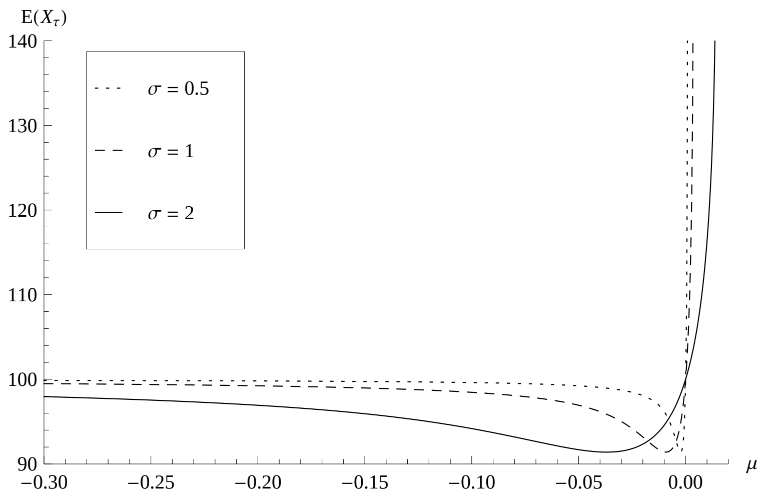

Figure 1.2: An example ($n = 2$, $x_0 = 100$) of the expected value of the stopped processes $\mathbb{E}(X_\tau)$ depending on the drift μ for different values of variance σ.

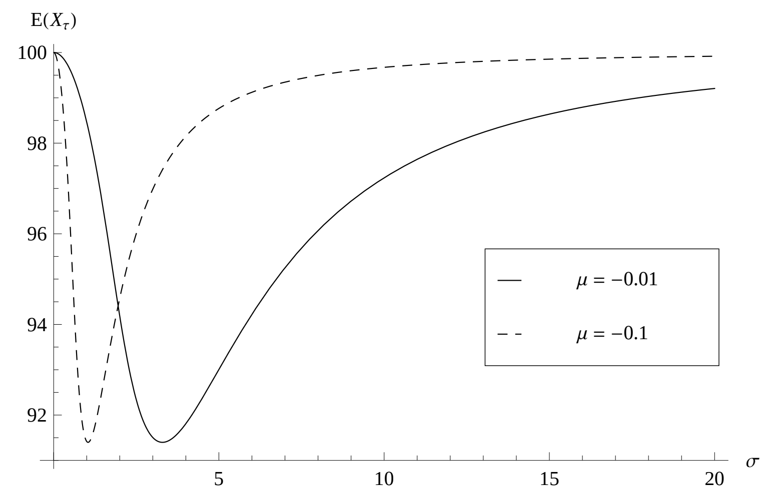

Figure 1.3: An example ($n = 2$, $x_0 = 100$) of the expected value of the stopped processes $\mathbb{E}(X_\tau)$ depending on the variance σ for different values of drift μ.

From an economic point of view, Proposition 5 illustrates a drawback of relative performance payments in risky environments: even if risky investment opportunities have only a slightly negative expected value, the principal loses a lot in expectation. Intuitively, contestants only care about outperforming each other and thus behave as if they were risk-loving. A simple linear compensation scheme based on absolute performance would not suffer from this drawback.

1.5 Conclusion

We have studied a new continuous' time model of contests. Contrary to the previous literature, players face a trade-off between a higher winning probability and a higher risk. If there are no good investment opportunities available, e.g., in a declining industry, contestants behave as if they were risk-loving—they invest in projects with negative expected returns. According to our main characterization result, Proposition 5, this problem is most severe for the natural case in which the drift is close to zero.

From a technical point of view, this chapter has developed a new method to verify equilibrium existence. The approach via Skorokhod embeddings seems promising to analyze other models without observability, because there are many sufficient conditions available in the probability theory literature.

1.6 Appendix

Proof of Lemma 2: Assume, by contradiction, there exists an interval $I = (a, b) \subset [\underline{x}^i, \overline{x}^i]$ such that $u^i(x) = \prod_{j \neq i} F^j(x)$ is constant for all $x \in I$. We distinguish three cases:

(i) For all player's $j \neq i$, $F^j(a) = 1$. Hence, by optimality, player i stops with probability 1 whenever at $\max_{j \neq i} \overline{x}^j$. This implies $\max_{j \neq i} \overline{x}^j \leq a \leq \underline{x}^i$ and player i wins for sure. Player j can deviate profitably and stop only if she hits 0 or $\overline{x}^i$, which contradicts the equilibrium assumption.

(ii) There exists a player $j \neq i$ with $F^j(b) = 0$. Hence, in equilibrium, no player ever stops in the interval $(0, \underline{x}^j)$, but at least two players stop with positive probability in every ϵ-ball around $\underline{x}^j$. To stop at $\underline{x}^j$ (with $u^i(\underline{x}^j) = 0$ by Lemma 1) is strictly worse than to continue at $\underline{x}^j$ until $X_t^i \in \{0, \max_j \overline{x}^j\}$. By continuity (Lemma 1), the argument extends to an ϵ-neighborhood of $\overline{x}^j$. This contradicts the equilibrium assumption of weak optimality of stopping in $(\underline{x}^j, \underline{x}^j + \epsilon)$.

(iii) No player $j \neq i$ stops in I, but (i) and (ii) do not hold. Hence, player i does not stop in I. Denote by $\tilde{x}$ the infimum of points above b at which a player stops. At $\tilde{x}$ (and, by continuity at an ϵ-neighborhood of $\tilde{x}$), it is strictly better to continue until $X_t^j \in \{\frac{b+a}{2}, \max_j \overline{x}^j\}$ than to stop, which contradicts the equilibrium assumption. $\square$

Proof of Lemma 3: We define $\Phi(x) = \prod_{i=1}^n F^i(x) = F^i(x)u^i(x)$. Denote the set of players who stop at x by $M(x) \subseteq N$, i.e.,

$$M(x) = \{i \in \{1, \ldots, n\} : (F^i)'(x) \neq 0\}.$$

By Lemma 2, $|M(x)| \geq 2$ for all $\min_{i \in N} \underline{x}^i < x < \max_{i \in N} \overline{x}^i$. For notational convenience, we omit the point x, at which all functions are evaluated, i.e., we write F^i and M instead of $F^i(x)$ and $M(x)$. Furthermore, we write $\mathbb{E}(\mathrm{d}u^i(x))$ shorthand for $\mathbb{E}(\mathrm{d}u^i(X_s^i)|\mathcal{F}_s^i)$ given $X_s^i = x$. For every agent $k \notin M$, we have:

$$|M|\Phi' = \sum_{i \in M} (F^i u^i)' = \sum_{i \in M} \left(F^i u^{i\prime} + F^{i\prime} u^i\right) = \sum_{i \in M} F^i u^{i\prime} + \underbrace{\sum_{i \in N} F^{i\prime} u^i}_{\Phi'}$$

$$\begin{aligned}
\Leftrightarrow \quad & (|M| - 1)\Phi' = \sum_{i \in M} F^i u^{i\prime} \\
\Rightarrow \quad & (|M| - 1)F^k u^{k\prime} = \sum_{i \in M} F^i u^{i\prime} \\
\Rightarrow \quad & (|M| - 1)F^k u^{k\prime\prime} = \sum_{i \in M} \left(F^i u^{i\prime\prime} + F^{i\prime} u^{i\prime}\right).
\end{aligned}$$

We calculate the expected change in winning probability of player k if he continues to play for an infinitesimally small amount of time $\mathbb{E}(\mathrm{d}u^k)$:

$$\begin{aligned}(|M|-1)F^k\mathbb{E}(\mathrm{d}u^k) &= (|M|-1)F^k(\mu u^{k\prime}+\frac{\sigma^2}{2}u^{k\prime\prime})\\ &= \mu(|M|-1)F^k u^{k\prime}+\frac{\sigma^2}{2}(|M|-1)F^k u^{k\prime\prime}\\ &= \mu\sum_{i\in M}F^i u^{i\prime}+\frac{\sigma^2}{2}\sum_{i\in M}\left(F^i u^{i\prime\prime}+F^{i\prime}u^{i\prime}\right)\\ &= \sum_{i\in M}\underbrace{(\mu u^{i\prime}+\frac{\sigma^2}{2}u^{i\prime\prime})}_{=\mathbb{E}(\mathrm{d}u^i)=0}F^i+\sum_{i\in M}\underbrace{F^{i\prime}u^{i\prime}}_{>0}\\ &> 0\,.\end{aligned}$$

As agent $i \in M$ stops with strictly positive probability in any neighborhood of x, he is indifferent between the strategy that stops at x and any other strategy that stops in a small neighborhood of x. Thus, $\mathbb{E}(\mathrm{d}u^i(x)) = 0$.
So far, we have shown that $\mathbb{E}(\mathrm{d}u^i(x)) = 0$ if $i \in M(x)$ and $\mathbb{E}(\mathrm{d}u^i(x)) > 0$ if $i \notin M(x)$. For every agent i, there exists an interval $I \subset [\underline{x}^i, \overline{x}^i]$ such that $i \in M(x)$ for every $x \in I$. Whenever $X_t^i = x \in I$, agent i is indifferent between the strategy that stops immediately and the strategy $\tau = \inf\{t : t \in \{\underline{x}^i, \overline{x}^i\}\}$. Formally,

$$\begin{aligned}0 &= u^i(x)-\mathbb{E}(u^i(X_\tau))\\ &= u^i(x)-\mathbb{E}(u^i(x)+\int_t^\tau \mu u^{i\prime}(X_s^i)+\frac{\sigma^2}{2}u^{i\prime\prime}(X_s^i)\mathrm{d}s+\int_t^\tau u^{i\prime}(X_s^i)\sigma\mathrm{d}B_s)\\ &= \mathbb{E}(\int_t^\tau \mu u^{i\prime}(X_s^i)+\frac{\sigma^2}{2}u^{i\prime\prime}(X_s^i)\mathrm{d}s)=\mathbb{E}(\int_t^\tau \mathbb{E}(\mathrm{d}u^i(X_s^i)))\,.\end{aligned}$$

The process enters every interval and $\mathbb{E}(\mathrm{d}u^i(x))$ is non-negative for all $x \in [\underline{x}^i, \overline{x}^i]$. Hence, the expectation $\mathbb{E}(\int_t^\tau \mathbb{E}(\mathrm{d}u^i(X_s^i)))$ can only be zero if $\mathbb{E}(\mathrm{d}u^i(x)) = 0$ almost surely. □

Proof of Lemma 4: By contradiction, assume $\max_i \underline{x}^i \neq 0$. Thus, to stop at $X_t^j = \underline{x}^i$ (and, by continuity in a neighborhood of this point) is strictly worse than to continue until $X_t^j \in \{0, \max_i \overline{x}^i\}$; this contradicts optimality.

Assume there exist players i and j such that $\overline{x}^i > \overline{x}^j$. Assume player j reaches his right endpoint at time t, $X_t^j = \overline{x}^j$. By the same argument as in Lemma 3, the continuation strategy $\tau = \inf\{s \geq t : X_s^j \in \{\overline{x}^j - \epsilon, \overline{x}^i\}\}$ is strictly better than to stop at $\overline{x}^j$, which contradicts optimality. $\square$

Proof of Proposition 4: To prove uniqueness, note that Lemma 1-4 do not rely on any symmetry arguments and are also valid for asymmetric starting values. Hence, the equation $u^i(x) = F^j(x)$ fixes the above construction uniquely given the right endpoint. The minmax property (constant sum game) implies that each player must receive the same payoff in any equilibrium. Thus, the local martingale condition uniquely determines $\overline{x}$. By Lemma 1, only one agent might set a mass point at 0. Feasibility implies that the agent with the lower starting value sets the mass point at zero and uniquely determines the size of the mass point. $\square$

Formula for the Expected Value in the n-Player Case:

Let *Hyp* denote the Gauss hypergeometric function.

$$\begin{aligned}
\mathbb{E}(x) &= \int_0^{\overline{x}} x f(x)\mathrm{d}x = (\overline{x}F(\overline{x}) - 0F(0)) - \int_0^{\overline{x}} F(x)\mathrm{d}x \\
&= \overline{x} - \int_0^{\overline{x}} \sqrt[n-1]{\frac{1}{n}\frac{\exp(-2\mu x) - 1}{\exp(-2\mu x) - 1}}\mathrm{d}x \\
&= \overline{x} + \frac{\sqrt[n-1]{1 - \exp(-2\mu\overline{x})}}{2\mu}(n-1)Hyp(\frac{1}{n-1}, \frac{1}{n-1}, \frac{n-2}{n-1}, \exp(2\mu\overline{x}))\,.
\end{aligned}$$

Proof of Proposition 5. We apply the monotone transformation $y = \exp(\frac{2\mu x_0}{\sigma^2})$ to $\mathbb{E}(X_\tau)$ to get

$$\begin{aligned}
\mathbb{E}(X_\tau) &= \frac{x_0}{\log(y)} + (1 + \frac{y}{2(1-y)})(x_0 - \frac{x_0 \log(2-y)}{\log(y)})\,, \\
&= x_0\left(\frac{1}{\log(y)} + (1 + \frac{y}{2(1-y)})(1 - \frac{\log(2-y)}{\log(y)})\right)\,.
\end{aligned}$$

for $y \neq 1$. This expression is convex if and only if it is convex for $x_0 = 1$. Assumption 1 implies $y \in (0,2)$.

$$
\begin{aligned}
\frac{\partial^2 \mathbb{E}(X_\tau)/x_0}{\partial y^2} &= \frac{4(-2+y)(-1+y)^3 + 2(-1+y)^2\left(2-5y+2y^2\right)\log(y)}{2(-2+y)(-1+y)^3 y^2 \log(y)^3} \\
&+ \frac{y^2\left(3-4y+y^2\right)\log(y)^2 - 2(-2+y)y^2\log(y)^3}{2(-2+y)(-1+y)^3 y^2 \log(y)^3} \\
&- \frac{(-2+y)\log(2-y)(2(-2+y)(-1+y)^2 - 2y^2\log(y)^2)}{2(-2+y)(-1+y)^3 y^2 \log(y)^3} \\
&- \frac{(-2+y)\log(2-y)\log(y)\left(-2+7y-6y^2+y^3\right)}{2(-2+y)(-1+y)^3 y^2 \log(y)^3}
\end{aligned}
$$

with the continuous extension $\frac{\partial^2 \mathbb{E}(X_\tau)/x_0}{\partial y^2} = \frac{1}{6}$ at $y = 1$. Simple algebra shows that nominator and denominator are negative on $y \in (0,2), y \neq 1$. Hence, the function is convex on $(0,2)$. As y is monotone increasing in μ, $\mathbb{E}(X^i_{\tau^i})$ is quasi-convex in μ. As y is also monotone increasing (decreasing) in σ for $\mu < 0$ ($\mu > 0$), $\mathbb{E}(X^i_{\tau^i})$ is quasi-convex (quasi-concave) in σ if $\mu < 0$ ($\mu > 0$).

It remains to show that $\mathbb{E}(X_\tau)$ is first decreasing, then increasing. For $\mu \to -\infty$ and $\mu \to 0$, $\mathbb{E}(X_\tau) \to x_0$. For any negative value of μ, the expected value of the stopped processes is smaller than x_0, because the process is a supermartingale. Hence, by quasi-convexity, $\mathbb{E}(X_\tau)$ has to be first decreasing, then increasing.

$\square$

Chapter 2

Continuous Time Contests

This chapter introduces a contest model in continuous time in which each player decides when to stop a privately observed Brownian motion with drift and incurs costs depending on his stopping time. The player who stops his process at the highest value wins a prize. Under mild assumptions on the cost function, we prove existence and uniqueness of the Nash equilibrium outcome, even if players have to choose bounded stopping times. We derive a closed form of the equilibrium distribution. If the noise vanishes, the equilibrium outcome converges to—and thus selects—the symmetric equilibrium outcome of an all-pay auction. For positive noise levels, results differ from those of all-pay contests with complete information—for instance, participants make positive profits. We show that for two players and constant costs, the profits of each participant increase for higher costs of research, higher volatility, or lower productivity of each player. Hence, participants prefer a contest design which impedes research progress.

2.1 Introduction

Two types of models are predominant in the literature on contests, races, and tournaments. In one of these, there is no learning about the performance measure or standings throughout the competition at all, while the other one considers full feedback about the performance of each player at all points in time. The former category includes all-pay contests with complete information (Hillman and Samet, 1987, Siegel, 2009, 2010), Tullock contests (Tullock, 1980), silent timing games (Karlin, 1953, Park and Smith, 2008), and models with additive noise in the spirit of Lazear and Rosen (1981). The latter category contains wars of attrition (Maynard Smith, 1974, Bulow and Klemperer, 1999), races (Aoki, 1991, Hörner, 2004, Anderson and Cabral, 2007), and contest models with full observability such as Harris and Vickers (1987) and Moscarini and Smith (2007).

In this chapter, we want to analyze an intermediate case in which there is partial feedback about the performance measure. More precisely, a player observes his own stochastic research progress over time, but he does not observe the progress of the other players or their effort decisions. A good example for this setting is an R&D contest. Each participant is well-informed about his own progress, but often uninformed about the progress of his competitors. For concrete examples of such competitions, see, e.g., Taylor (1995).

Formally, our model is an n-player contest in which each player decides when to stop a privately observed Brownian motion (X_t) with drift μ and volatility σ. As long as a player exerts effort, i.e., does not stop the process, he incurs flow costs $c(X_t)$. The player who stops his process at the highest value wins a prize.

Under mild assumptions on the cost function—it has to be continuous and bounded away from zero—the game has a unique Nash equilibrium outcome. This outcome is implementable in stopping strategies which stop almost surely before a fixed time $T < \infty$. Hence, provided the contest length is above a threshold, the equilibrium outcome is independent of the contest length. In equilibrium, each player makes positive expected profits. For two players and constant costs, these profits increase if the productivity (drift) of both players decreases, the volatility increases, or the costs increase. Hence, participants prefer a contest design which impedes research progress.

The formal analysis proceeds as follows. Proposition 6 and Theorem 1 establish existence and uniqueness of the equilibrium distribution. The existence proof first characterizes the equilibrium distribution $F(x)$ of values at the stopping time $X_\tau = x$ uniquely up to its endpoints. We then use a Skorokhod embedding approach to show that there exists a stopping strategy, which induces this distribution. This technique from probability theory (e.g., Skorokhod, 1961, 1965; for a survey, see Oblój, 2004) has already been introduced in the first chapter.

Moreover, we verify a condition from a recent paper in mathematics (Ankirchner and Strack, 2011) to show that there exists a bounded time stopping strategy—a strategy that stops almost surely before a fixed time $T < \infty$—which induces the equilibrium distribution. As most real-world contests have a fixed deadline, this result fortifies the predictions of the model. It is also one of main technical contributions of this chapter, since the technique is also applicable to other models without observability. However, for tug-of-war models with full observability (Harris and Vickers, 1987, Moscarini and Smith, 2007, Gul and Pesendorfer, 2011), one cannot construct bounded time equilibria in a similar way, because, for any fixed deadline, there is a positive probability that no player has a sufficiently lead until the deadline.

We then analyze the shape of the equilibrium distribution. As uncertainty vanishes, the distribution converges to the symmetric equilibrium distribution of an all-pay auction by Theorem 2. On the one hand, the model offers a microfoundation for the use of all-pay auctions to scrutinize environments in which uncertainty is not a crucial ingredient; on the other hand, it gives an equilibrium selection result between the equilibria of the symmetric all-pay auction analyzed in Baye et al. (1996). Moreover, this result serves as a benchmark to discuss how our predictions differ from all-pay models if volatility is strictly positive.

For any $\sigma > 0$, Proposition 7 states that all players make positive expected profits in equilibrium. Intuitively, agents use the private information about their progress to generate rents. The intuition is similar to an all-pay contest, in which players have incomplete information about the valuation or effort cost of their rivals, see, e.g., Hillman and Riley (1989), Amann and Leininger (1996), Krishna and Morgan (1997), or Moldovanu and Sela (2001).

Finally, we analyze the special case of two players and constant costs. We derive a closed-form solution for the profits of each player, which depends only on the ratio $\frac{2\mu^2}{c\sigma^2}$ (Proposition 10). In particular, profits increase as costs c increase, volatility σ^2 increases, or productivity μ decreases (Theorem 3). Hence, contestants prefer to have mutually worse technologies. This result, which does not hold in a static all-pay contest, goes along with the common intuition that players prefer competition to be less fierce.

2.1.1 Related Literature

In the first chapter, we have analyzed a model in which players do not have any costs of research, but have a (usually negative) drift and face a bankruptcy constraint. The driving forces of both models differ substantially. In particular, in this chapter, contestants trade off higher costs against a higher winning probability, whereas in the first, the trade-off is between winning probability and risk. Also, the applications of the first chapter are related to finance and managerial compensation, while this chapter is in spirit of the contest literature.

This chapter entails a direct extension of the literature on *silent timing games*—see, e.g., Karlin (1953). Among others, the literature on silent timing games scrutinizes our setting for the case without uncertainty. Intuitively, adding uncertainty allows us to have a model with partial learning throughout the contest.

With a similar motivation, Taylor (1995) also analyzes a model in which players only learn about their own stochastic research success. In his T-period model, however, the highest draw in a single period determines this success. The resulting equilibrium stopping rule is a threshold strategy, which stops whenever a player has a draw above a deterministic, time-independent value.

We proceed as follows. Section 2.2 sets up the model. In Section 2.3, we prove that an equilibrium exists and is unique. Section 2.4 discusses the relation to all-pay contests and derives the main comparative statics results. Section 2.5 concludes. Most proofs are relegated to the appendix.

2.2 The Model

There are $n < \infty$ agents indexed by $i \in \{1, 2, \ldots, n\} = N$ who face a stopping problem in continuous time. At each point in time $t \in \mathbb{R}_+$, agent i privately observes the realization of a stochastic process $(X_t^i)_{t \in \mathbb{R}_+}$ with

$$X_t^i = x_0 + \mu t + \sigma B_t^i .$$

The constant x_0 denotes the starting value of all processes; without loss of generality, we assume $x_0 = 0$. The drift $\mu \in \mathbb{R}_+$ is the common expected change of each process X_t^i per time, i.e., $\mathbb{E}(X_{t+\Delta}^i - X_t^i) = \mu\Delta$. The noise term is an n-dimensional Brownian motion (B_t) scaled by $\sigma \in \mathbb{R}_+$.

2.2.1 Strategies

A pure strategy of player i is a stopping time τ^i. This stopping time depends only on the realization of his process X_t^i, as the player only observes his own process.Mathematically, the agents' stopping decision until time t has to be $\mathcal{F}_t^i$-measurable, where $\mathcal{F}_t^i = \sigma(\{X_s^i : s < t\})$ is the sigma algebra induced by the possible observations of the process X_s^i before time t. In contrast to the first chapter, we now require stopping times to be bounded by a real number $T < \infty$ such that $\tau^i < T$ almost surely.

To incorporate mixed strategies, we allow for randomized stopping times—progressively $\mathcal{F}_t^i$-measurable functions $\tau^i(\cdot)$ such that, for every $r^i \in [0, 1]$, the value $\tau^i(r^i)$ is a stopping time. Intuitively, agents draw a random number r^i from the uniform distribution on $[0, 1]$ before the game and play a stopping strategy $\tau^i(r^i)$.[1]

[1] Although the unique equilibrium outcome of this chapter can be obtained in pure strategies, we allow for mixing to obtain the results in a more general framework.

2.2.2 Payoffs

The player who stops his process at the highest value wins a prize $p > 0$. Ties are broken randomly. Until he stops, each player incurs flow costs $c : \mathbb{R} \to \mathbb{R}_{++}$, which depend on the current value of the process X_t, but not on the time t. The payoff π^i is thus

$$\pi^i = \frac{p}{k}\,\mathbf{1}_{\{X^i_{\tau^i} = \max_{j \in N} X^j_{\tau^j}\}} - \int_0^{\tau^i} c^i(X^i_t)dt \ ,$$

where $k = |\{i \in N : X^i_{\tau^i} = \max_{j \in N} X^j_{\tau^j}\}|$ is the number of agents who stop at the highest value. All agents maximize their expected profit $\mathbb{E}(\pi^i)$. We henceforth normalize p to 1, since agents only care about the trade-off between winning probability and cost-prize ratio. The cost function satisfies the following mild assumption:

Assumption 2. For every $x \in \mathbb{R}$, the cost function $c : \mathbb{R} \to \mathbb{R}_{++}$ is continuous and bounded away from zero on $[x, \infty)$.

There are several possible interpretations for the production technology. For instance, one can interpret the drift as progress in research and the martingale part as learning. Alternatively, the process could measure the progress in the production of a prototype. In this interpretation, the variance might be due to different market prizes of each component, which influence the value of the prototype. Apart from that, the prototype might turn out to require more or less components compared to the construction plan.

2.3 Equilibrium Construction

In this section, we first establish some necessary conditions on the distribution functions in equilibrium. In a second step, we prove existence and uniqueness of the Nash equilibrium outcome and determine the equilibrium distributions depending on the cost function.

Every strategy of agent i induces a (potentially non-smooth) cumulative distribution function (cdf) $F^i : \mathbb{R} \to [0, 1]$ of his stopped process $F^i(x) = \mathbb{P}(X^i_{\tau^i} \leq x)$. Denote the endpoints of the support of the equilibrium distribution of player i by

$$\begin{aligned} \underline{x}^i &= \inf\{x : F^i(x) > 0\} \\ \overline{x}^i &= \sup\{x : F^i(x) < 1\}. \end{aligned}$$

Let $\underline{x} = \max_{i\in N} \underline{x}^i$ and $\overline{x} = \max_{i\in N} \overline{x}^i$. In the next step, we establish a series of auxiliary results that are crucial to prove uniqueness of the equilibrium distribution.

Lemma 6. *At least two players stop with positive probability on every interval $I = (a, b) \subset [\underline{x}, \overline{x}]$.*

Lemma 7. *No player places a mass point in the interior of the state space, i.e., for all i, for all $x > \underline{x}$: $\mathbb{P}(X^i_{\tau^i} = x) = 0$. At least one player has no mass at the left endpoint, i.e., $F^i(\underline{x}) = 0$, for at least player i.*

We omit the proof of Lemma 7, since it is just a specialization of the standard logic in static game theory with a continuous state space; see, e.g., Burdett and Judd (1983). Intuitively, in equilibrium, no player can place a mass point in the interior of the state space, since no other player would then stop slightly below the mass point. This contradicts Lemma 6.

Lemma 7 implies that the probability of a tie is zero. Thus, we can express the winning probability of player i if he stops at $X^i_{\tau^i} = x$, given the distributions of the other players, as

$$u^i(x) = \mathbb{P}(\max_{j\neq i} X^j_{\tau^j} \le x) = \prod_{j\neq i} F^j(x).$$

Lemma 8. *All players have the same right endpoint, $\overline{x}^i = \overline{x}$, for all i.*

Lemma 9. *All players have the same expected profit in equilibrium. Moreover, each player loses for sure at $\underline{x}$, i.e., $u^i(\underline{x}) = 0$, for all i.*

Lemma 10. *All players have the same equilibrium distribution function, $F^i = F$, for all i.*

As players have symmetric distributions, we henceforth drop the superscript i. The previous lemmata imply that each player is indifferent between any stopping strategy which remains on his support. By Itô's lemma, it follows from the indifference inside the support that, for every point $x \in (\underline{x}, \overline{x})$, the function $u(\cdot)$ must satisfy the second order ordinary differential equation (ODE)

$$c(x) = \mu u'(x) + \frac{\sigma^2}{2} u''(x) \,. \tag{2.3.1}$$

As (2.3.1) is a second order ODE, we need two boundary conditions to determine $u(\cdot)$ uniquely. One boundary condition is $u(\underline{x}) = 0$ from Lemma 9. We determine the other one in the following lemma:

Lemma 11. *In equilibrium, $u'(\underline{x}) = 0$.*

The idea of the proof in the appendix is simple. If the derivative was negative, $u'(\underline{x}) < 0$, there would a profitable deviation at $\underline{x}$, which stops in the neighborhood of $\underline{x}$ rather than at the point itself.

Imposing the two boundary conditions, the solution to equation (2.3.1) is unique. To calculate it, we define $\phi(x) = \exp(\frac{-2\mu x}{\sigma^2})$ as a solution of the homogeneous equation $0 = \mu u'(x) + \frac{\sigma^2}{2} u''(x)$. To solve the inhomogeneous equation (2.3.1), we apply the variation of the constants formula. We then use the two boundary conditions to calculate the unique solution candidate. Finally, we rearrange with Fubini's Theorem to get

$$u(x) \quad = \quad \begin{cases} 0 & \text{for } x < \underline{x} \\ \frac{1}{\mu} \int_{\underline{x}}^{x} c(z)(1 - \phi(x - z)) \mathrm{d}z & \text{for } x \in [\underline{x}, \overline{x}] \\ 1 & \text{for } x > \overline{x} \,. \end{cases}$$

By symmetry of the equilibrium strategy, the function $F : \mathbb{R} \to [0, 1]$ satisfies $F(x) = \sqrt[n-1]{u(x)}$. Consequently, the unique candidate for an equilibrium distribution is

$$F(x) = \begin{cases} 0 & \text{for all } x < \underline{x} \\ \sqrt[n-1]{\frac{1}{\mu} \int_{\underline{x}}^{x} c(z)(1 - \phi(x - z)) \mathrm{d}z} & \text{for all } x \in [\underline{x}, \overline{x}] \\ 1 & \text{for all } x > \overline{x} \,. \end{cases}$$

In the next step, we verify that F is a cumulative distribution function, i.e., that F is nondecreasing and that $\lim_{x\to\infty} F(x) = 1$.

Lemma 12. *F is a cumulative distribution function.*

Proof. By construction of F, $F(\underline{x}) = 0$. Clearly, F is increasing on $(\underline{x}, \overline{x})$, as the derivative with respect to x,

$$F'(x) = \frac{F(x)^{2-n}}{(n-1)}\left(\frac{2}{\sigma^2}\int_{\underline{x}}^{x} c(z)\phi(x-z)\mathrm{d}z\right),$$

is greater than zero for all $x > \underline{x}$. It remains to show that there exists an $x > \underline{x}$ such that $F(x) = 1$.

$$\begin{aligned}
F(x)^{n-1} &= \frac{1}{\mu}\int_{\underline{x}}^{x} c(z)(1-\phi(x-z))\mathrm{d}z \\
&\geq \frac{1}{\mu}\inf_{y\in[\underline{x},\infty)} c(y)\left(x-\underline{x}-\frac{\sigma^2}{2\mu}(1-\phi(x-\underline{x}))\right) \\
&\geq \frac{1}{\mu}\inf_{y\in[\underline{x},\infty)} c(y)(x-\underline{x}-\frac{\sigma^2}{2\mu})
\end{aligned}$$

Assumption 1 implies that the cost function $c(\cdot)$ is bounded away from zero. Consequently, $\inf_{y\in[\underline{x},\infty)} c(y)$ is strictly greater than zero. Continuity of F implies that there exists a point $\overline{x} > \underline{x}$ such that $F(\overline{x}) = 1$. $\square$

The next lemma derives a necessary condition for a distribution F to be the outcome of a strategy τ.

Lemma 13. *If $\tau \leq T < \infty$ is a bounded stopping time that induces the continuous distribution $F(\cdot)$, i.e., $F(z) = \mathbb{P}(X_\tau \leq z)$, then $1 = \int_{\underline{x}}^{\overline{x}} \phi(x)F'(x)\mathrm{d}x$.*

Proof. Observe that $(\phi(X_t))_{t\in\mathbb{R}_+}$ is a martingale. Hence, by Doob's optional stopping theorem, for any bounded stopping time τ,

$$1 = \phi(X_0) = \mathbb{E}[\phi(X_\tau)] = \int_{\underline{x}}^{\overline{x}} \phi(x)F'(x)\mathrm{d}x\,.$$

$\square$

We now use the necessary condition from Lemma 13 to prove that the equilibrium distribution is unique.

Proposition 6. *There exists a unique pair* $(\underline{x}, \overline{x}) \in \mathbb{R}^2$ *such that the distribution*

$$F(x) = \begin{cases} 0 & \text{for all } x \leq \underline{x} \\ \sqrt[n-1]{\frac{1}{\mu} \int_{\underline{x}}^{x} c(z)(1-\phi(x-z))dz} & \text{for all } x \in (\underline{x}, \overline{x}) \\ 1 & \text{for all } x \geq \overline{x} \end{cases}$$

is the unique candidate for an equilibrium distribution.

Proof. As F is continuous, the right endpoint $\overline{x}$ satisfies $1 = \int_{\underline{x}}^{\overline{x}} F'(x; \underline{x}, \overline{x}) \mathrm{d}x$. Since $F'(x; \underline{x}, \overline{x})$ is independent of $\overline{x}$, we henceforth drop the dependency in our notation. By the implicit function theorem,

$$\frac{\partial \overline{x}}{\partial \underline{x}} = -\frac{-\overbrace{F'(\underline{x}; \underline{x})}^{=0} + \int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \mathrm{d}x}{F'(\overline{x}; \underline{x})} = -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \mathrm{d}x}{F'(\overline{x}; \underline{x})}. \tag{2.3.2}$$

Lemma 13 states that any feasible distribution satisfies $1 = \int_{\underline{x}}^{\overline{x}} F'(x; \underline{x}) \phi(x) \mathrm{d}x$. Applying the implicit function theorem to this equation gives us

$$\begin{aligned} \frac{\partial \overline{x}}{\partial \underline{x}} &= -\frac{-\overbrace{F'(\underline{x}; \underline{x})}^{=0} + \int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \phi(x) \mathrm{d}x}{F'(\overline{x}; \underline{x}) \phi(\overline{x})} \\ &= -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \overbrace{\phi(x - \overline{x})}^{<1} \mathrm{d}x}{F'(\overline{x}; \underline{x})} \\ &< -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \mathrm{d}x}{F'(\overline{x}; \underline{x})}. \end{aligned} \tag{2.3.3}$$

The last inequality follows from $\frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \geq 0$. Hence, conditions (2.3.2) and (2.3.3) intersect exactly once. Thus, in equilibrium, the left and right endpoint are unique. $\square$

Hence, each equilibrium strategy induces the distribution F. The next lemma shows that this condition is also sufficient.

Lemma 14. *Every strategy that induces the unique distribution F from Proposition 6 is an equilibrium strategy.*

Proof. Define $\Psi(\cdot)$ as the unique solution to (2.3.1) with the boundary conditions $\Psi(\underline{x}) = 0$ and $\Psi'(\underline{x}) = 0$. By construction, the process $\Psi(X^i_t) - \int_0^t c(X^i_s)\mathrm{d}s$ is a martingale and $\Psi(x) = u(x)$ for all $x \in [\underline{x}, \overline{x}]$. As $\Psi'(x) < 0$ for $x < \underline{x}$ and $\Psi'(x) > 0$ for $x > \overline{x}$, $\Psi(x) > u(x)$ for all $x \notin [\underline{x}, \overline{x}]$. For every stopping time S, we use Itô's Lemma to calculate the expected value

$$\begin{aligned}\mathbb{E}[u(X_S) - \int_0^S c(X_t)\mathrm{d}t] &\leq \mathbb{E}[\Psi(X_S) - \int_0^S c(X_t)\mathrm{d}t] \\ &= \Psi(X_0) = u(X_0) = \mathbb{E}(u(X_\tau))\,.\end{aligned}$$

The last equality results from the indifference of every agent to stop immediately with the expected payoff $u(X_0)$ or to play the equilibrium strategy with the expected payoff $\mathbb{E}(u(X_\tau))$. □

The intuition is simple. By construction of F, all agents are indifferent between all stopping strategies which stop inside the support $[\underline{x}, \overline{x}]$. As every agent wins with probability one at the right endpoint, it is strictly optimal to stop there. The condition $F'(\underline{x}) = 0$ ensures that it is also optimal to stop at the left endpoint.

So far, we have verified that a bounded stopping time $\tau \leq T < \infty$ is an equilibrium strategy if and only if it induces the distribution $F(\cdot)$, i.e., $F(z) = \mathbb{P}(X_\tau \leq z)$. To show that the game has a Nash equilibrium, it remains to establish the existence a bounded stopping time inducing $F(\cdot)$. The problem of finding a stopping time τ such that a Brownian motion stopped at τ has a given centered probability distribution F, i.e., $F \sim B_\tau$, is known in the probability literature as the Skorokhod embedding problem (SEP). Since its initial formulation in Skorokhod (1961, 1965), many solutions have been derived; for a survey article, see Obłój (2004). In a recent mathematical paper, Ankirchner and Strack (2011) find conditions guaranteeing the existence of stopping times τ that are *bounded* by some real number $T < \infty$, and embed a given distribution in Brownian motion, possibly with drift.[2] In addition to the assumption stated in the next lemma, Ankirchner and Strack (2011) assume that the condition in Lemma 13 holds, which we have already imposed. They define $g(x) = F^{-1}(\Phi(x))$, where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} \exp(\frac{z^2}{2})\mathrm{d}z$ is the density function of the normal distribution.

Lemma 15 (Ankirchner and Strack, 2011, Theorem 2)**.** *Suppose that $g(\cdot)$ is Lipschitz-continuous with Lipschitz constant $\sqrt{T}$. Then the distribution F can be embedded in $X_t = \mu t + B_t$, with a stopping time that stops almost surely before T.*

The main conceptual innovation is the bounded time requirement $\tau < T$. This is not trivial, as for any fixed time horizon T, there exists a positive probability that X_t does not leave any interval $[a, b]$ with $a < X_0 < b$. Hence, a similar construction as in Proposition 2 in the first chapter, which uses a mixture over cutoff strategies of the form

$$\tau_{a,b} = \inf\{t : \mathbb{R}_+ : X_t \notin [a, b]\} \,,$$

cannot be used to implement F. The proof in Ankirchner and Strack (2011) constructs a pure strategy for any distribution, which meets the Lipschitz condition.

This type of equilibrium, which is independent of the deadline provided it is sufficiently high, cannot be obtained in tug-of-war models with full observability (Harris and Vickers, 1987, Moscarini and Smith, 2007, Gul and Pesendorfer, 2011). Intuitively, for any fixed deadline, in these models there is a positive probability that no player has a sufficient lead until the deadline, which detains a similar result.

The previous lemma enables us to prove the main result of this section:

Theorem 1. *The game has a Nash equilibrium.*

[2]Ankirchner and Strack (2011) use a construction of the stopping time introduced for Brownian motion without drift in Bass (1983) and for the case with drift in Ankirchner et al. (2008).

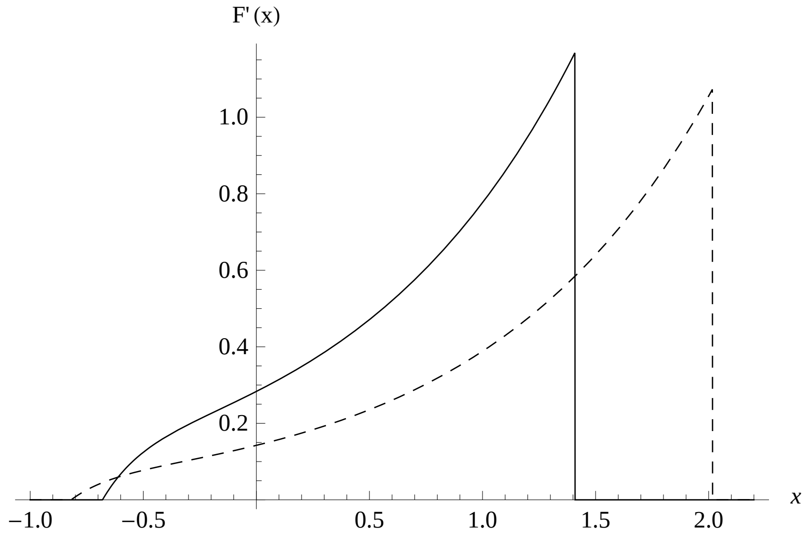

Figure 2.1: The density function $F'(\cdot)$ for the parameters $n = 2$, $\mu = 3$, $\sigma = 1$ and the cost-functions $c(x) = \exp(x)$ solid line and $c(x) = \frac{1}{2}\exp(x)$ dashed line.

The proof in the appendix verifies Lipschitz continuity of the function g, which makes Lemma 15 applicable. Thus, a Nash equilibrium in bounded time stopping strategies exists, and, by Proposition 6, the equilibrium distribution F is unique.[3]

2.4 Equilibrium Analysis

2.4.1 Convergence to the All-Pay Auction

This subsection considers the relationship between the literature on all-pay contests and our model for vanishing noise. We first establish an auxiliary result about the left endpoint:

Lemma 16. *If the noise vanishes, the left endpoint of the equilibrium distribution converges to zero, i.e.,* $\lim\limits_{\sigma \to 0} \underline{x} = 0$.

Proof. For any bounded stopping time, for any $\sigma > 0$, feasibility implies that $\underline{x} \leq 0$. By contradiction, assume there exists a constant ϵ such that $\underline{x} \leq \epsilon < 0$ for some sequence

[3]It is straightforward to show that the distribution F is also the unique equilibrium distribution in the space of finite time stopping strategies.

$(\sigma_k)_{k\in\mathbb{N}}$ with $\lim_{k\to\infty}\sigma_k = 0$. Then F' is bounded away from zero by

$$\begin{aligned} F'(x) &= \frac{F(x)^{2-n}}{n-1}\frac{2}{\sigma^2}\int_{\underline{x}}^{x} c(z)\phi(x-z)\mathrm{d}z \\ &\geq \frac{1}{n-1}\frac{2}{\sigma^2}\int_{\epsilon}^{x} c(z)\phi(x-z)\mathrm{d}z \\ &= \frac{1}{\mu(n-1)}\left(\inf_{y\in[\epsilon,\infty)} c(y)\right)(1-\phi(x-\epsilon)). \end{aligned}$$

For every point $x < 0$, $\lim_{\sigma_k\to 0}\phi(x) = \infty$. Thus, $\lim_{\sigma_k\to 0}\int_{\underline{x}}^{0} F'(x)\phi(x)\mathrm{d}x > 1$, which contradicts feasibility, because $\int_{\underline{x}}^{0} F'(x)\phi(x)\mathrm{d}x \leq \int_{\underline{x}}^{\overline{x}} F'(x)\phi(x)\mathrm{d}x = 1$. $\square$

Taking the limit $\sigma \to 0$, the symmetric equilibrium distribution converges to

$$\lim_{\sigma\to 0} F(x) = \sqrt[n-1]{\frac{1}{\mu}\int_0^x c(z)\mathrm{d}z}.$$

In a static n-player all-pay auction, the equilibrium distribution is

$$F(x) = \sqrt[n-1]{\frac{x}{v}},$$

where x is the total outlay of a participant and v is her valuation; see, e.g., Hillman and Samet (1987). In our case, the total outlay depends on the flow costs at each point, the speed of research μ, and the stopping time τ. More precisely, it is $\int_0^x \frac{c(z)}{\mu}\mathrm{d}z$. The valuation v in the analysis of Hillman and Samet (1987) coincides with the prize p—which we have normalized to one—in our contest. This yields us the following proposition:

Theorem 2. *For vanishing noise, the equilibrium distribution converges to the symmetric equilibrium distribution of an all-pay auction.*

Thus, our model supports the use of all-pay auctions to analyze contests in which the variance is negligible. Figure 2.2 illustrates the similarity to the all-pay auction equilibrium if variance σ and costs $c(\cdot)$ are small in comparison to the drift μ.

Moreover, the symmetric all-pay auction has multiple equilibria—for a full characterization see Baye et al. (1996). This chapter offers a selection criterion in favor of the symmetric

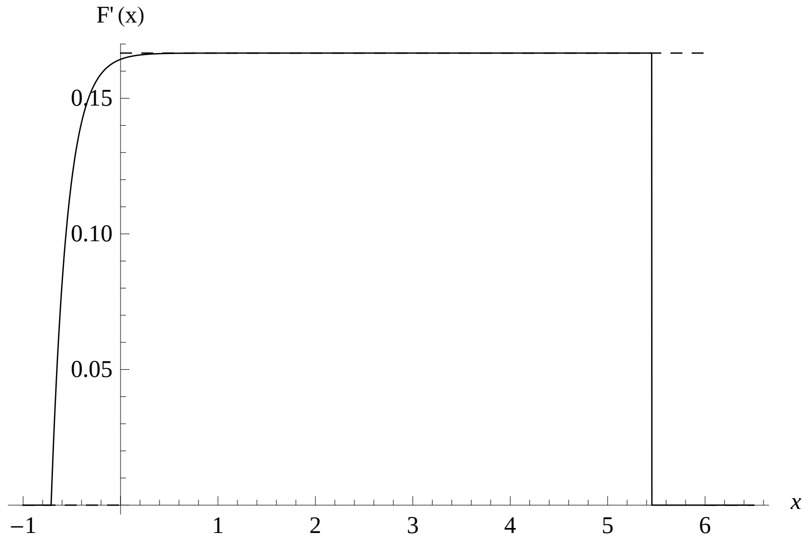

Figure 2.2: This picture shows the density function $F'(\cdot)$ with support $[-0.71, 5.45]$ for the parameters $n = 2$, $\mu = 3$, $\sigma = 1$ and the cost-functions $c(x) = \frac{1}{2}$ (solid line) and for the same parameters the equilibrium density of the all-pay auction with support $[0, 6]$ (dashed line).

equilibrium, in which no participant places a mass point at zero. Intuitively, all other equilibria of the symmetric all-pay auction include mass points at zero for some players, which is not possible in our model for any positive σ by Lemma 7.

2.4.2 Comparative Statics and Rent Dispersion

Proposition 2 has linked all-pay contests with complete information to our model for the case of vanishing noise. In the following, we scrutinize how the predictions differ for positive noise. In a symmetric all-pay contests with complete information, agents make zero profits in equilibrium. This does not hold true in our model for any positive level of variance σ:

Proposition 7. *In equilibrium, all agents make strictly positive expected profits.*

Proof. In equilibrium, agents are indifferent between stopping immediately and the equilibrium strategy. Their expected profit is thus $u(0)$, which is strictly positive as $\underline{x} < 0$. $\square$

Intuitively, agents generate informational rents through their private information about the research progress. A similar result is known in the literature on all-pay contests with incomplete information, see, e.g., Hillman and Riley (1989), Amann and Leininger (1996), Krishna and Morgan (1997), and Moldovanu and Sela (2001). In these models, participants take a draw from a distribution prior to the contest, which determines their effort cost or valuation. The outcome of the draw is private information. In contrast to this, private information about one's progress arrives continuously over time in our model.

In the following, we derive comparative statics in the number of players for constant costs. We define the support length as $\Delta = \overline{x} - \underline{x}$.

Lemma 17. *If the number of players n increases and $c(x) = c$, the support length Δ remains constant and both endpoints increase.*

Proof. If $c(x) = c$, $F(\overline{x}) - F(\underline{x})$ clearly depends only on Δ. Hence, for $F(\overline{x}) - F(\underline{x}) = 1$, Δ has to be constant. As F gets more concave if n increases, by feasibility, $\underline{x} \nearrow$ and $\overline{x} \nearrow$. □

Proposition 8. *If the number of agents n increases and $c(x) = c$, the expected profit of each agent decreases.*

Proof. The function $u(x)$ depends only on $x - \underline{x}$. As n increases, $\underline{x}$ increases by Lemma 17. Thus, the expected value of stopping immediately, $u(0)$, which is an optimal strategy in both cases, decreases as n increases. □

Hence, in accordance with most other models, each player is worse off if the number of contestants increases.

2.4.3 The Special Case of Two Players and Constant Costs

We now restrict attention to the case $n = 2$ and $c(x) = c$ to get more explicit results. For this purpose, we require additional notation. In particular, we denote by $W_0 : [-\frac{1}{e}, \infty) \rightarrow$

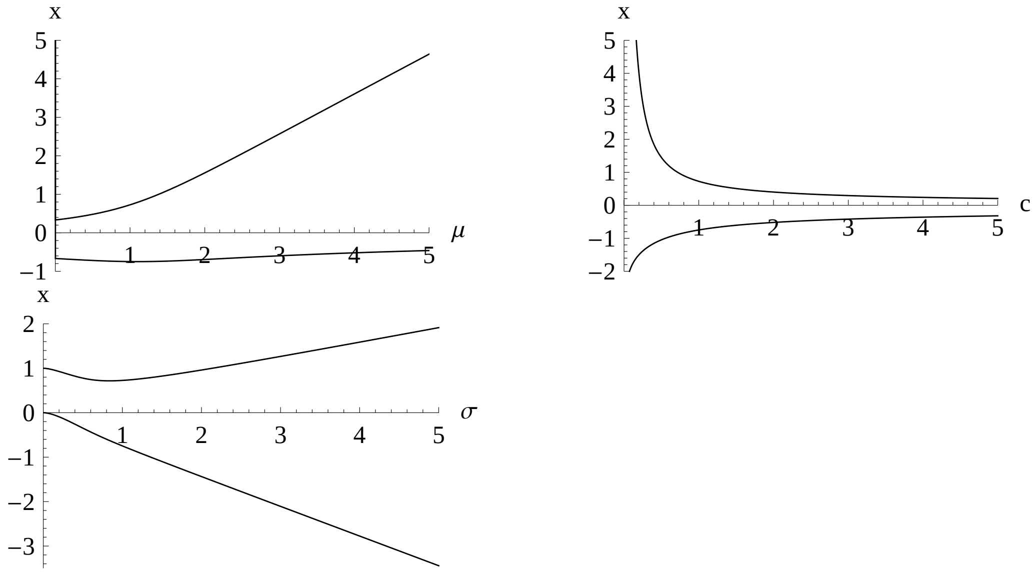

Figure 2.3: This picture shows the left endpoint $\underline{x}$ and the right endpoint $\overline{x}$ for $n = 2, \sigma = 1, \mu = 1$ and constant cost $c = 1$ varying the productivity μ in the first picture, the costs c in the second picture and the variance σ in the third.

$\mathbb{R}_+$ the principal branch of the Lambert W-function. This branch is implicitly defined on $[-\frac{1}{e}, \infty)$ as the unique solution of $x = W(x)\exp(W(x))$, $W \geq -1$. Define $h : \mathbb{R}_+ \to [0,1]$ by

$$h(y) = \exp(-y - 1 - W_0(-\exp(-1-y)))\,.$$

The next proposition pins down the left and right endpoint of the support of the players.

Proposition 9. *The left and right endpoint are*

$$\begin{aligned}
\underline{x} &= \frac{\sigma^2}{2\mu}\left(2\log(1 - h(\frac{2\mu^2}{c\sigma^2})) - \log(\frac{4\mu^2}{c\sigma^2})\right) \\
\overline{x} &= \frac{\sigma^2}{2\mu}\left(2\log(1 - h(\frac{2\mu^2}{c\sigma^2})) - \log(\frac{4\mu^2}{c\sigma^2}) - \log(h(\frac{2\mu^2}{c\sigma^2}))\right)\,.
\end{aligned}$$

For an illustration how the endpoints change depending on the parameters, see Figure 2.3. The next proposition derives a closed-form solution of the profits π of each player.

Proposition 10. *The equilibrium profit of each player depends only on the ratio* $y = \frac{2\mu^2}{c\sigma^2}$. *It is given by*

$$\pi = \frac{(1-h(y))^2}{2y^2} - \frac{2\log(1-h(y)) - \log(2y) - 1}{y}\,.$$

Given the previous proposition, it is simple to establish the main comparative statics result of this chapter.

Theorem 3. *The equilibrium profit of each player increases if costs increase, variance increases, or drift decreases.*

To get an intuition for the result, we decompose the term $\frac{2\mu^2}{c\sigma^2}$, which determines the equilibrium profit of the players, into two parts:

$$\frac{2\mu^2}{c\sigma^2} = \underbrace{\frac{\mu}{c}}_{\text{Productivity}} \times \underbrace{\frac{2\mu}{\sigma^2}}_{\text{Signal to noise ratio}}\,.$$

The first term is the productivity $\frac{\mu}{c}$ of the agents. As firms get more productive, competition gets more fierce and each firm makes less profits. The second term $\frac{2\mu}{\sigma^2}$ is the signal to noise ratio, which measures the informativeness of $X^i_{\tau^i}$. Intuitively, if the signal to noise ratio decreases, the outcome $X^i_{\tau^i}$ becomes less correlated with agent i's effort choice τ^i. In turn, this reduces his incentives to exert effort and thereby the cost of his expected stopping time. As his winning probability in equilibrium remains constant, his profits are decreasing in the signal to noise ratio. Summarizing, participants prefer to have mutually worse—more costly, more random, or less productive—technologies.

Even for a perfectly uninformative signal, however, agents cannot extract the full surplus:

Proposition 11. *The equilibrium profit of each agent is bounded from above by* $4/9$.

Proof. The agents profit is decreasing in $y = \frac{2\mu^2}{c\sigma^2} \geq 0$ by Theorem 3. Hence, profits are bounded from above by $\lim_{y\to 0} u(0)$. By l'Hôpital's rule,

$$\lim_{y\to 0} \frac{(1-h(y))^2}{2y^2} - \frac{2\log(1-h(y)) - \log(2y) - 1}{y} = \frac{4}{9}\,.$$

$\square$

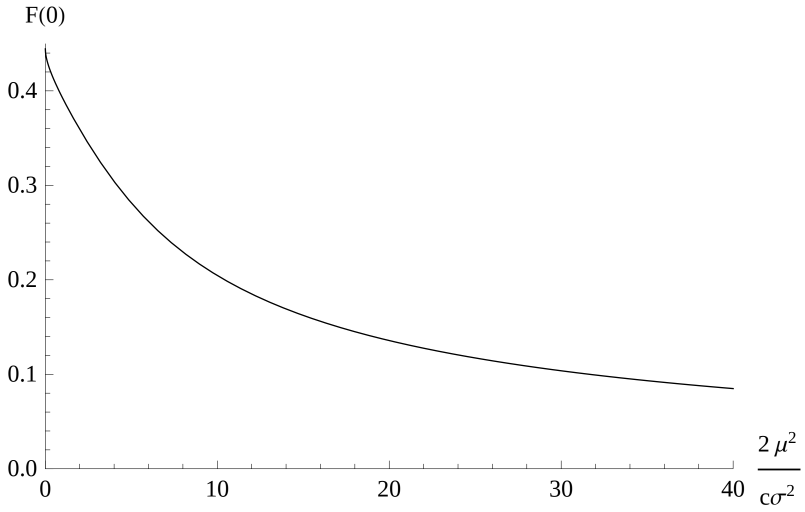

Figure 2.4: This picture shows the equilibrium profit $F(0)$ of the agents on the y-Axes for $n = 2$ constant cost-functions $c(x) = c \in \mathbb{R}_+$ with $y = \frac{2\mu^2}{c\sigma^2}$ on the x-Axes.

The expected equilibrium effort $\mathbb{E}(\tau^i)$ is bounded from below by

$$\begin{aligned} & \frac{4}{9} \geq \mathbb{E}(F(X^i_{\tau^i}) - c\tau^i) = \frac{1}{2} - c\mathbb{E}(\tau^i) \\ \Leftrightarrow \quad & \mathbb{E}(\tau^i) \geq \frac{1}{18c}\,. \end{aligned}$$

2.5 Conclusion and Discussion

In this chapter, we have introduced a model of contests in continuous time in which each player learns only about his own research progress. Under mild assumptions on the cost function, a Nash equilibrium outcome exists and is unique. If the research progress contains little uncertainty, the equilibrium is close to the symmetric equilibrium of a static all-pay auction. If the research outcome is uncertain, players prefer mutually higher costs of research, worse technologies, and higher uncertainty. These comparative statics, which go along with the intuition that players prefer competition to be less fierce, cannot be obtained in a static all-pay contest.

From a technical perspective, we have introduced a method to construct equilibria in continuous time games that are independent of the time horizon. Furthermore, we have introduced a constructive method to calculate a minimal time horizon that ensures the existence of such equilibria. These methodological contributions may prove fruitful in future research, and add to the general understanding of continuous time models.

2.6 Appendix

Proof of Lemma 6. As players have to use bounded time stopping strategies, each player i stops with positive probability on every subinterval of $[\underline{x}^i, \overline{x}^i]$. Hence, it suffices to show that at least two players have $\overline{x}$ as their right endpoint. Assume, by contradiction, only player i has $\overline{x}$ as his right endpoint. Denote $\overline{x}^{-i} = \max_{j \neq i} \overline{x}^j$. Then, for any $\epsilon > 0$, at $\overline{x}^{-i} + \epsilon$, player i strictly prefers to stop, which yields him the maximal possible winning probability of 1 without any additional costs. This contradicts the optimality of a strategy, which stops at $\overline{x}^i > \overline{x}^{-i} + \epsilon$. $\square$

Remark 1. *We write* $\tau^i_{(a,b)}(x)$ *shorthand for* $\inf\{t : X^i_t \notin (a,b) | X^i_s = x\}$ *in the next three proofs. Clearly,* $\tau^i_{(a,b)}(x)$ *is not a bounded time strategy, but we use it to bound the payoffs. Note that, for sufficiently large time horizon* T*, the payoff from stopping at* $\min\{\tau^i_{(a,b)}(x), T\}$ *is arbitrarily close to that of* $\tau^i_{(a,b)}(x)$*.*

Proof of Lemma 8. Assume $\overline{x}^j > \overline{x}^i$. For at least two players j, j', the payoff from $\tau^j_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)$ is weakly higher than from stopping at $X^j_t = \overline{x}^i$ by Lemma 6. By Lemma 7, at least one of these players—denote it j—wins with probability zero at $\underline{x}^j$. Note that $u^i(\overline{x}^i) = \prod_{h \neq i} F^h(\overline{x}^i) < \prod_{h \neq j} F^h(\overline{x}^i) = u^j(\overline{x}^i)$, because $F^i(\overline{x}^i) = 1 > F^j(\overline{x}^i)$.

Optimality of $\tau^j_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)$ implies

$$u^j(\overline{x}^i) \leq \mathbb{P}(X^j_{\tau^j} = \overline{x}^j | \tau^j_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)) u^j(\overline{x}^j) - \mathbb{E}(c(\tau^j_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)))\,.$$

On the other hand,

$$u^i(\overline{x}^i) < u^j(\overline{x}^i) \leq \mathbb{P}(X^i_{\tau^i} = \overline{x}^j | \tau^i_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)) u^i(\overline{x}^j) - \mathbb{E}(c(\tau^i_{(\underline{x}^j, \overline{x}^j)}(\overline{x}^i)))\,.$$

Hence, at $X^i_t = \overline{x}^i$, for a sufficiently long time horizon T, player i can profitably deviate by stopping at $\min\{\tau^i_{(a,b)}(x), T\}$. This contradicts the equilibrium assumption. □

Proof of Lemma 9. To prove the first statement, we distinguish two cases.
(i) If at least two players have $F^i(\underline{x}) = 0$, then $u^i(\underline{x}) = 0\ \forall i$. Assume there exists a player j who makes less profit than a player i, where $\pi^i \leq \mathbb{P}(X^i_{\tau^i} = \overline{x} | \tau^i_{(\underline{x}^i, \overline{x})}(0)) - \mathbb{E}(c(\tau^i_{(\underline{x}^i, \overline{x})}(0)))$. If player j deviates to the strategy $\min\{\tau^j_{(\underline{x}^j, \overline{x})}(0), T\}$, player j gets a profit arbitrarily close to π^i; this contradicts optimality of player j's strategy.
(ii) If only one player has $F^i(\underline{x}) = 0$, then $u^i(\underline{x}) > 0$. We now consider the case in which this player i makes a weakly higher payoff than the remaining players, who make the same payoff each—otherwise the argument in the first part of the proof leads to a contradiction. For any interval $I \in [\underline{x}, \overline{x}]$ in which player i stops with positive probability, by Lemma 6, there exists another player j who also stops in the interval. In particular, for $x \in I$, for any $\epsilon > 0$, we get

$$\mathbb{P}(X^i_{\tau^i} = \overline{x} | \tau^i_{(\underline{x}, \overline{x})}(x)) + \mathbb{P}(X^i_{\tau^i} = \underline{x} | \tau^i_{(\underline{x}, \overline{x})}(x)) u^i(\underline{x}) - \mathbb{E}(c(\tau^i_{(\underline{x}, \overline{x})}(x))) < u^i(x) + \epsilon$$

and

$$\mathbb{P}(X^j_{\tau^j} = \overline{x} | \tau^j_{(\underline{x}, \overline{x})}(x)) - \mathbb{E}(c(\tau^j_{(\underline{x}, \overline{x})}(x))) \geq u^j(x)\ \forall j \neq i.$$

For $\epsilon \to 0$, the two equations imply that $u^i(x) > u^j(x)$, for all $j \neq i$, for all x in the support of player i. Hence, $F^i(x) \leq F^j(x)\ \forall j$, for all x on the support of player i, and, by monotonicity of F^j, on $[\underline{x}, \overline{x}]$. Thus, the distribution of player i stochastically dominates that of all other players. This contradicts feasibility, since all players start at the same value and stopping times have to be bounded.
The second statement of the lemma follows immediately from the proof of (ii). □

Proof of Lemma 10. Recall that all players have the same profit and $u^i(\underline{x}) = 0\ \forall i$. Each player stops on any interval $I \subset [\underline{x}, \overline{x}]$ with positive probability, since stopping times are bounded. By contradiction, assume $u^i(x) > u^j(x)$ for some players i, j and some value x. As it is weakly optimal for player i to continue at x with $\tau^i_{(\underline{x},\overline{x})}(x)$, this strategy is strictly optimal for player j. At x, player j thus has a bounded time stopping strategy whose expected payoff is arbitrarily close to $u^i(x)$, which contradicts $u^i(x) > u^j(x)$. Hence, $u^i(x) = u^j(x)$ holds globally, which implies $F^i(x) = F$, for all i. $\square$

Proof of Lemma 11. By definition, $u(x) = 0$ for all $x \leq \underline{x}$. Hence, the left derivative $\partial_- u(\underline{x})$ is zero. It remains to prove that the right derivative $\partial_+ u(\underline{x})$ is also zero. For a given $u : \mathbb{R} \to \mathbb{R}_+$, let $\Psi : \mathbb{R} \to \mathbb{R}$ be the unique function that satisfies the second order ordinary differential equation $c(x) = \mu\Psi'(x) + \frac{\sigma^2}{2}\Psi''(x)$ with the boundary conditions $\Psi(\underline{x}) = u(\underline{x})$ and $\Psi'(\underline{x}) = \partial_+ u(\underline{x})$. Clearly, $\Psi'(\underline{x}) \geq 0$ because u is weakly increasing. The remaining proof goes by contradiction. Assume $\Psi'(\underline{x}) > 0$. Then there exists a point $\hat{x} < \underline{x}$ such that $\Psi(x) < 0 = u(x)$ for all $x \in (\hat{x}, \underline{x})$. Consider the strategy S that stops when either the point $\hat{x}$ or $\overline{x}$ is reached or at 1,

$$S = \min\{1, \inf\{t \in \mathbb{R}_+ : X^i_t \notin [\hat{x}, \overline{x}]\}\}.$$

As $u(x) \geq \Psi(x)$ with strictly inequality for $x \in (\hat{x}, \underline{x})$, it follows that $\mathbb{E}(u(X_S)) > \mathbb{E}(\Psi(X_S))$. Thus,

$$\mathbb{E}(u(X_S) - \int_0^S c(X^i_t)\mathrm{d}t) > \mathbb{E}(\Psi(X_S) - \int_0^S c(X^i_t)\mathrm{d}t)\,.$$

Note that, by Itô's lemma, the process $\Psi(X^i_t) - \int_0^t c(X^i_s)\mathrm{d}s$ is a martingale. By Doob's optional sampling theorem, agent i is indifferent between the equilibrium strategy τ and the bounded time strategy S, i.e.,

$$\begin{aligned}\mathbb{E}(\Psi(X_S) - \int_0^S c(X^i_t)\mathrm{d}t) &= \mathbb{E}(\Psi(X_\tau) - \int_0^\tau c(X^i_t)\mathrm{d}t) \\ &= \mathbb{E}(u(X_\tau) - \int_0^\tau c(X^i_t)\mathrm{d}t)\,.\end{aligned}$$

The last step follows because $u(x)$ and $\Psi(x)$ coincide for all $x \in (\underline{x}, \overline{x})$. Consequently, the strategy S is a profitable deviation, which contradicts the equilibrium assumption. $\square$

Proof of Theorem 1. The function Φ is Lipschitz continuous with constant $\frac{1}{\sqrt{2\pi}}$. Consequently, it suffices to prove Lipschitz continuity of F^{-1} to get the Lipschitz continuity of $F^{-1} \circ \Phi$. The density $f(\cdot)$ is

$$\begin{aligned} f(x) &= \frac{F(x)^{-n+2}}{n-1} \frac{2}{\sigma^2} \int_{\underline{x}}^{x} c(z)\phi(x-z)\mathrm{d}z \\ &= \frac{F(x)^{-n+2}}{n-1} \frac{2}{\sigma^2} \left(\int_{\underline{x}}^{x} c(z)\mathrm{d}z - \mu F(x)^{n-1} \right) \end{aligned}$$

As $f(x) > 0$ for all $x > \underline{x}$, it suffices to show Lipschitz continuity of F^{-1} at 0. We substitute $x = F^{-1}(y)$ to get

$$(f \circ F^{-1})(y) \geq \frac{1}{n-1} \frac{2}{\sigma^2} \left(y^{2-n} \underbrace{(\min_{z \in [\underline{x}, \overline{x}]} c(z))}_{=\underline{c}} (F^{-1}(y) - F^{-1}(0)) - \mu y \right).$$

Rearranging with respect to $F^{-1}(y) - F^{-1}(0)$ gives

$$\begin{aligned} F^{-1}(y) - F^{-1}(0) &\leq \left(\frac{(n-1)\sigma^2}{2} (f \circ F^{-1})(y) + \mu y \right) \frac{y^{n-2}}{\underline{c}} \\ &\leq \left(\frac{(n-1)\sigma^2}{2} f(\overline{x}) + \mu \right) \frac{y^{n-2}}{\underline{c}}. \end{aligned}$$

This proves the Lipschitz continuity of $F^{-1}(\cdot)$ for $n > 2$. Note that, for two agents, the function $F^{-1}(\cdot)$ is not Lipschitz continuous as $f(\underline{x}) = 0$. However, we show in the following paragraph that $F^{-1} \circ \Phi$ is Lipschitz continuous for $n = 2$.

$$\begin{aligned} F(x) &= \int_{\underline{x}}^{x} \frac{c(z)}{\mu} (1 - \phi(x-z)\mathrm{d}z \\ &\leq \underbrace{\left(\sup_{z \in [\underline{x}, \overline{x}]} \frac{c(z)}{\mu} \right)}_{=\overline{c}} \left(x - \underline{x} - \frac{\sigma^2}{2\mu} (1 - \phi(x - \underline{x})) \right) \end{aligned}$$

A second order Taylor expansion around $\underline{x}$ yields that, for an open ball around $\underline{x}$ and $\underline{x} < x$, we have the following upper bound

$$x - \underline{x} - \frac{\sigma^2}{2\mu}(1 - \phi(x - \underline{x})) \leq \frac{2\mu}{\sigma^2}(1 - \phi(x - \underline{x}))^2 \,.$$

For an open ball around $\underline{x}$, we get an upper bound on $F(x) \leq \frac{2\overline{c}}{\sigma^2}(1-\phi(x-\underline{x}))^2$ and hence the following estimate

$$1 - \phi(x - \underline{x}) \geq \sqrt{\frac{\sigma^2}{2\overline{c}}F(x)} \,.$$

We use this estimate to obtain a lower bound on $f(\cdot)$ depending only on $F(\cdot)$

$$\begin{aligned} f(x) &= \frac{2}{\sigma^2}\left(\int_{\underline{x}}^{x} c(z)\phi(x-z)\mathrm{d}z\right) \geq \frac{2\underline{c}}{\sigma^2}\left(\frac{\sigma^2}{2\mu}(1-\phi(x-\underline{x}))\right) \\ &\geq \frac{\underline{c}}{\mu}\sqrt{\frac{\sigma^2}{2\overline{c}}F(x)} \,. \end{aligned}$$

Consequently, there exists an $\epsilon > \underline{x}$ such that, for all $x \in [\underline{x}, \epsilon)$, we have an upper bound on $\frac{(\phi\circ\Phi^{-1}\circ F)(x)}{f(x)}$. Taking the limit $x \to \underline{x}$ yields

$$\lim_{x\to\underline{x}} \frac{(\phi\circ\Phi^{-1}\circ F)(x)}{f(x)} \quad\leq\quad \lim_{x\to\underline{x}} \frac{(\phi\circ\Phi^{-1}\circ F)(x)}{\frac{\underline{c}}{\mu}\sqrt{\frac{\sigma^2}{2\overline{c}}F(x)}} \leq \sqrt{\frac{2\overline{c}\mu^2}{\underline{c}^2\sigma^2}} \lim_{y\to 0}\frac{(\phi\circ\Phi^{-1})(y)}{\sqrt{y}} = 0 \,.$$

□

Proof of Proposition 9 *and* 10. Rearranging the density condition $1 = F(\overline{x}) = \frac{c}{\mu}[\Delta - \frac{\sigma^2}{2\mu}(1-\phi(\Delta))]$ yields

$$\exp(\frac{-2\mu\Delta}{\sigma^2}) = -\frac{2\mu}{\sigma^2}[\Delta - (\frac{\mu}{c} + \frac{\sigma^2}{2\mu})].$$

The solution to the transcendental algebraic equation $e^{-a\Delta} = b(\Delta - d)$ is $\Delta = d + \frac{1}{a}W_0(\frac{ae^{-ad}}{b})$, where $W_0 : [-\frac{1}{e}, \infty) \to \mathbb{R}_+$ is the principal branch of the Lambert W-function. This branch is implicitly defined on $[-\frac{1}{e}, \infty)$ as the unique solution of

$x = W(x)\exp(W(x))$, $W \geq -1$. Hence,

$$\Delta = \frac{\mu}{c} + \frac{\sigma^2}{2\mu}[1 + W_0(-\exp(-1 - \frac{2\mu^2}{c\sigma^2}))]$$

and

$$\begin{aligned}\phi(\Delta) &= \exp(\frac{-2\mu^2}{c\sigma^2} - 1 - W_0(-\exp(-1 - \frac{2\mu^2}{c\sigma^2}))) \\ &= \exp(-1 - y - W_0(-\exp(-1 - y))) \\ &= h(y)\,.\end{aligned}$$

Note that $\phi(\Delta)$ only depends on $y = \frac{2\mu^2}{c\sigma^2}$. Moreover, $h(y)$ is strictly decreasing in y, as $W_0(\cdot)$ and $\exp(\cdot)$ are strictly increasing functions. For constant costs, the feasibility condition from Lemma 13 reduces to

$$\begin{aligned}1 &= \int_{\underline{x}}^{\overline{x}} F'(x)\phi(x)\mathrm{d}x \\ &= \frac{c\sigma^2}{2\mu^2}[\frac{1}{2}\phi(\underline{x}) + \frac{1}{2}\phi(2\overline{x} - \underline{x}) - \phi(\overline{x})]\,.\end{aligned}$$

Dividing by $\phi(\underline{x})$ gives

$$\begin{aligned}\phi(-\underline{x}) &= \frac{c\sigma^2}{2\mu^2}[\frac{1}{2} + \frac{1}{2}\phi(\Delta)^2 - \phi(\Delta)] \\ &= \frac{1}{y}[\frac{1}{2} + \frac{1}{2}h(y)^2 - h(y)] \\ &= \frac{1}{2y}(1 - h(y))^2 \\ &= g(y)\end{aligned}$$

Note that $g : \mathbb{R}_+ \to [0, 1]$ is strictly decreasing in y. We calculate $\underline{x}$ as

$$\underline{x} = -\phi^{-1}(\phi(-\underline{x})) = -\frac{\sigma^2}{2\mu}\log(\frac{2\mu^2}{c\sigma^2[\frac{1}{2} + \frac{1}{2}\phi(\Delta)^2 - \phi(\Delta)]})\,.$$

Simple algebraic transformations yield the expression for $\underline{x}$ and $\overline{x}$ (inserting Δ) in Proposition 9.

We plug in $\underline{x}$ to get:

$$\begin{aligned}
F(0) &= \frac{c}{\mu}[-\underline{x} - \frac{\sigma^2}{2\mu}(1 - \phi(-\underline{x}))] \\
&= \frac{c\sigma^2}{2\mu^2}[\log(\frac{\frac{2\mu^2}{c\sigma^2}}{\frac{1}{2} + \frac{1}{2}\phi(\Delta)^2 - \phi(\Delta)}) + \frac{\frac{1}{2} + \frac{1}{2}\phi(\Delta)^2 - \phi(\Delta)}{\frac{2\mu^2}{c\sigma^2}} - 1] \\
&= \frac{1}{y}[\log(\frac{y}{\frac{1}{2} + \frac{1}{2}h(y)^2 - h(y)}) + \frac{\frac{1}{2} + \frac{1}{2}h(y)^2 - h(y)}{y} - 1] \\
&= \frac{1}{y}[g(y) - \log(g(y)) - 1]
\end{aligned}$$

Hence, the value of $F(0)$ depends on the value of the fraction $y = \frac{2\mu^2}{c\sigma^2}$ in the above way, which completes the proof of Proposition 10. □

Proof of Theorem 3. By Proposition 10, it suffices to show that the profit $F(0)$ is increasing in y. Consider the expression from the previous proof:

$$F(0) = \frac{1}{y}[g(y) - \log(g(y)) - 1]$$

The function $x - \log(x)$ is increasing in x. Hence, $g(y) - \log(g(y)) - 1$ is decreasing in y, because $g(y)$ is decreasing in y. As $\frac{1}{y}$ is also decreasing in y, the product $\frac{1}{y}[g(y) - \log(g(y)) - 1]$ is decreasing in y. □

Chapter 3

Equilibrium Equivalence of Stochastic Contests with Poisson Arrivals and All-Pay Auctions

We compare the n-player single-prize common-value all-pay auction with discrete bids to a stochastic contest model in continuous time. In the continuous time model, ideas arrive according to a Poisson process as long as a player exerts effort. At the termination date T, the player who has most arrivals wins a prize.

If T is above a threshold value, the symmetric equilibrium distributions of both games coincide. We derive an upper bound for the threshold value. Finally, we briefly discuss the equilibrium structure if the termination date falls short of the threshold.

3.1 Introduction

There is a large amount of literature on all-pay auctions that are often motivated as reduced-form models of a contest.[1] The bid in the auction serves as a proxy for the effort

[1] For example, Hillman and Samet (1987), Hillman and Riley (1989), Baye et al. (1993, 1996), and Che and Gale (1998) study all-pay auctions with a continuous bid space, while Dechenaux et al. (2003, 2006), and Cohen and Sela (2007) scrutinize all-pay auctions in which the set of bids is countable.

or production cost each participant incurs in the contest; think, for example, of an R&D, job promotion, or lobbying contest. This chapter contributes to the question if the all-pay auction is a valid model to analyze stochastic contests without observability.

To do so, we contrast two models. One is the single-prize common-value all-pay auction with discrete bids. In the other model, players can wait to build up a "stack", whose size determines their contest success. The current size of the stack and the stopping decision are private information. Contrary to a silent timing game, however, the arrival of successes, which add up to the stack, is not deterministic. Instead, as long as a player exerts costly effort, i.e., does not stop his process, successes arrive according to a Poisson process. At time T, the player who accumulated most successes wins a prize. Ties are broken randomly.

A stopping strategy in the contest induces a probability distribution over the number of successes at the stopping time. In an infinitely long contest, any such distribution can be induced by a stopping time. Hence, the choice of a stopping time is equivalent to the choice of a distribution in this case. If the designer adjusts the cost-prize ratio to the Poisson arrival rate, both games can be represented by the same normal form. Thus, they share the same set of Nash equilibrium distributions.[2]

For a finite contest, however, the normal form differs from that of the all-pay auction, since some distributions cannot be induced by a stopping strategy. For instance, a distribution that puts all probability mass on one success cannot be replicated in bounded time, since, for any finite value T, there is a positive probability that no success occurs until time T.

Nevertheless, under a genericity assumption on the parameters, Proposition 13 shows that if T is large enough, any symmetric Nash equilibrium distribution of the discrete all-pay auction can be implemented in the contest. Moreover, by Proposition 14, for sufficiently large T, the game has no other symmetric equilibria. Hence, the set of symmetric equilibrium distributions coincides for both games. In Proposition 15, we provide a time bound

[2] Thompson (1952) describes the general principles of coalescing of moves; for a discussion, see also Kohlberg and Mertens (1986).

suggesting that the critical level T such that the equilibrium distribution can be induced is moderate. Hence, an all-pay auction is a valid model to analyze stochastic contests, if (i) the contest lasts sufficiently long and (ii) it has a cumulative structure, i.e., ideas that arrive add up to the final success.

For small values of T, the long-run equilibrium distribution cannot be induced. This constraint is similar to the introduction of a binding bidding cap in an all-pay auction in that it also restricts the set of feasible distributions. However, the effect of both restrictions is diametrically opposed. We use an example to illustrate that a restriction on the time T leads to more probability mass at lower points, since players do not have enough time to transfer their mass. With a binding bidding cap, however, there is more mass at the highest point, because players can choose any distribution on the given interval, but compensate for the restriction on the maximal bid by placing more mass on the highest feasible bid.

3.1.1 Related Literature

Compared to the analysis in Chapter 2, this chapter considers a simpler, weakly increasing research process. This structure allows us to get a direct equivalence to an all-pay auction and an explicit time bound such that it holds. Moreover, we provide some intuition how equilibria look like in short contests, which remains an open question for the setting in Chapter 2.

The main difference to the seminal paper by Taylor (1995) is that in his model the best draw in a single period determines the contest success. The equilibrium distributions of his game differ substantially from those of an all-pay auction.

Clearly, the question which model is more appropriate for an application is context-specific. For example, consider a university department that wants to hire a new researcher. If the department hires the applicant who has more (major) publications, a cumulative structure is appropriate. If it hires the applicant with the single best publication, Taylor's model is a good description. Similarly, the measure of success in a job promotion contest for managers may be either the number of successful projects or the best project. Another example fitting our criterion are typical R&D competitions, e.g., for a fighter jet or tank in which the government usually awards the prize depending on a number of attributes.

We proceed as follows. Section 3.2 discusses the discrete all-pay auction. In Section 3.3, we formally introduce the stochastic contest and derive the main equivalence results. Section 3.4 concludes. We relegate most proofs to the appendix.

3.2 The All-Pay Auction

Consider a model with n risk-neutral players indexed by $i \in \{1, \ldots, n\} = N$. A pure strategy of each player i is a bid $x^i \in \mathbb{N}_0$ entailing costs cx^i. We henceforth normalize c to 1. A mixed strategy of player i is a probability measure $f^i \colon \mathbb{N}_0 \mapsto [0,1]$. Denote the associated cumulative distribution function by $F^i(z) = \mathbb{P}(x^i \leq z) = \sum_{y=0}^{z} f^i(y)$. The agent with the highest bid wins a prize p. Ties are broken randomly. Hence, the utility of player i is

$$u^i(x^1, \ldots, x^n) = \begin{cases} p - x^i & \text{if } x^i > x^j \quad \forall j \neq i \\ \frac{p}{m} - x^i & \text{if } i \text{ ties for the highest bid with } m-1 \text{ others} \\ -x^i & \text{otherwise.} \end{cases}$$

We henceforth impose the following assumption:

Assumption 3. If $n = 2$, then $\frac{p}{2} \notin \mathbb{N}$.

This assumption rules out non-generic parameter settings, in which the prize is an even integer. In those cases, infinitely many equilibria exist; see Baye et al. (1994) and Cohen and Sela (2007). In the next lemma, we derive a global indifference property for the set of equilibrium strategies.

Lemma 18. *In every symmetric Nash equilibrium, all players are indifferent between the pure strategies* $0, 1, \ldots, \overline{x}$ *with* $\overline{x} = \min\{x \in \mathbb{N}_0 : F(x) = 1\}$*. Moreover,* $f(x) > 0$ *if and only if* $x \leq \overline{x}$*.*

The next proposition characterizes the equilibrium for the two-player case.[3]

Proposition 12. *Assume $n = 2$. The unique Nash equilibrium of the all-pay auction is*

$$f^i(x) = \begin{cases} 1 - \frac{\overline{x}}{p} & \text{if } x \leq \overline{x} \text{ and } x \text{ is even} \\ \frac{1}{p}(2 + \overline{x}) - 1 & \text{if } x < \overline{x} \text{ and } x \text{ is odd} \\ 0 & \text{otherwise,} \end{cases}$$

for both players, where $\overline{x} = \max\{x \in \mathbb{N}_0 \colon x \text{ even and } \frac{x}{p} < 1\}$.

In contrast to a previous analysis of the two-player case in Baye et al. (1994), we get a unique equilibrium. The different result stems from Assumption 3, which rules out the parameter setting of Baye et al. (1994) in which the prize has to be an even integer.

The probability measure $f(x)$ alternates due to the tie-breaking rule. For the appropriate choices of the prize—$p = y$ for any odd integer y—the above expression becomes a uniform distribution.

3.3 The Stochastic Contest

In the contest, at every point $t \leq T \in \mathbb{R}_+$, each player $i \in \{1, 2, \ldots, n\} = N$ privately observes a time-homogeneous Poisson process X_t^i with intensity $\lambda \in \mathbb{R}_+$ and jump size 1. A strategy of player i is a stopping time $\tau^i \leq T$ with respect to the natural filtration $\mathcal{F}_t^i$ generated by the process X_t^i. This stopping time induces a probability distribution over values of the process at the stopping point. We denote this distribution by $F^i \colon \mathbb{N}_0 \to [0, 1]$, where $F^i(x) := \mathbb{P}(X_{\tau^i}^i \leq x)$, and the associate probability measure by $f^i(x)$. Stopping at time t entails costs of $\tilde{c}t$. The player who has more Poisson arrivals at her stopping time τ^i wins a prize $\tilde{p}$. Thus, each player's profit is

[3] Dechenaux et al. (2003, 2006) analyze the game with bidding caps, in which the prize is foregone in case of a tie. Cohen and Sela (2007) consider the case of asymmetric valuations.

$$\pi^i = \begin{cases} \tilde{p} - \tilde{c}\tau^i & \text{if } X^i_{\tau^i} > \max\limits_{j \neq i} X^j_{\tau^j} \\ \frac{\tilde{p}}{m} - \tilde{c}\tau^i & \text{if i ties for the highest value with } m-1 \text{ others} \\ -\tilde{c}\tau^i & \text{otherwise.} \end{cases}$$

Formally, the number of bidders in a tie is $m = |\{j \in N : X^j_{\tau^j} = \max_{i \in N} X^i_{\tau^i}\}|$. Define the payoff process $(\Pi^i_t)_{t \in \mathbb{R}_+}$ of player i as his expected payoff π^i of stopping immediately

$$\Pi^i_t \quad = \quad \mathbb{E}(\pi^i | X^i_{\tau^i} = X^i_t) = \tilde{p}\ \mathbb{P}(X^i_t = \max_{j \in N} X^j_{\tau^j})\ \mathbb{E}(\frac{1}{m} | X^i_t = \max_{j \in N} X^j_{\tau^j}) - \tilde{c}t\,.$$

Hence, it is the prize times a weighted probability of winning it alone or sharing it minus the cost. In the next step, we characterize distributions F for which the process can be stopped such that the probability distribution at the stopping value equals F.

Definition 1. A distribution $F\colon \mathbb{N}_0 \to [0,1]$ is *feasible* in a contest of length $T \in \mathbb{R}_+$ if there exists a stopping time τ, with $\tau \leq T$ almost surely, that induces F.

The following lemma gives a characterization of feasible distributions; the proof in the appendix constructs a corresponding strategy.

Lemma 19. *Consider the process X^i_t as defined above. Take any distribution F with $f(x) > 0$ if and only if $x \leq \overline{x} \in \mathbb{N}_0$. There exists a time bound $T' \in \mathbb{R}_+$ such that, for all $T \geq T'$, F is feasible.*

In particular, all symmetric equilibrium distributions of the all-pay auction are feasible if the contest lasts long enough. The following proposition establishes the main result.

Proposition 13. *Let F be a symmetric equilibrium distribution of the all-pay auction, T' as defined in Lemma 19 and $\frac{\tilde{c}}{\tilde{p}} = \frac{\lambda}{p}$. Then, for all $T \geq T'$, F is also a Nash equilibrium distribution of the stochastic contest.*

Proof. By Lemma 19, there exists a strategy in the stochastic contest that induces the Nash equilibrium distribution of the all-pay auction for T sufficiently high. To show that this is indeed an equilibrium distribution for the stochastic contest, we verify that no player has a profitable deviation. Define $f^{-i}(x) = \mathbb{P}(\max_{j \neq i} X^j_{\tau^j} = x)$ and

$$m^i(x) = |\{j \in N : j \neq i, X^j_{\tau^j} = x = \max_{l \neq i} X^l_{\tau^l}\}| + 1.$$

The expectation of the infinitesimal generator of the process Π^i is

$$\begin{aligned}
&\frac{\mathbb{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t}(X_t = x) \\
&= \tilde{p}\lambda\left(f^{-i}(x)(1 - \mathbb{E}(\frac{1}{m^i(x)})) + f^{-i}(x+1)\mathbb{E}(\frac{1}{m^i(x+1)})\right) - \tilde{c} \\
&= \tilde{p}\lambda\left(f^{-i}(x)(1 - \mathbb{E}(\frac{1}{m^i(x)})) + f^{-i}(x+1)\mathbb{E}(\frac{1}{m^i(x+1)}) - \frac{1}{p}\right).
\end{aligned}$$

The formula for the infinitesimal generator omits multiple jumps, because they occur with probability zero. With a slight abuse of notation, we also denote the number of players in a tie for the all-pay auction by $m^i(x) = |\{j \in N : j \neq i, x^j = x = \max_{l \neq i} x^l\}| + 1$. The indifference property for the all-pay auction (Lemma 18) is mathematically equivalent to

$$0 = (f^{-i}(x)(1 - \mathbb{E}(\frac{1}{m^i(x)})) + f^{-i}(x+1)\mathbb{E}(\frac{1}{m^i(x+1)}))p - 1.$$

It follows that $\frac{\mathbb{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t} = 0$ for all $X_t < \overline{x}$. By optimality of a bid $\overline{x}$ compared to $\overline{x}+1$ in the all-pay auction, we have $\frac{\mathbb{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t} \leq 0$ at $\overline{x}$. For all points $x > \overline{x}$ the winning probability does not increase after a success, which implies $\frac{\mathbb{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t} = -\tilde{c} < 0$. Hence, the process Π^i_t is a supermartingale and a local martingale on the support. Denote by τ^i the stopping time inducing the equilibrium distribution F and by τ^i_0 the strategy which stops immediately. By Doob's optional stopping theorem (Rogers and Williams, 2000, p.189), for any bounded time stopping time $\hat{\tau}^i$, $\mathbb{E}(\Pi^i_{\hat{\tau}^i}) \leq \mathbb{E}(\Pi^i_{\tau^i_0}) = \mathbb{E}(\Pi^i_{\tau^i})$. Thus, there is no profitable deviation. $\square$

Hence, if the duration of the contest is long enough and the designer adjusts the prize to the Poisson arrival rate, any equilibrium distribution of the all-pay auction is also an equilibrium distribution in the stochastic contest. Intuitively, players use the stochasticity to replicate the equilibrium distribution of the all-pay auction. As a minor difference, players in the contest model face uncertainty about the realization of the process, whereas in the all-pay auction, each player can determine her bid deterministically.

In the next proposition, we establish the reverse direction of the previous result. Hence, for long enough contests both games have the same equilibria.

Proposition 14. *Assume* $\frac{\tilde{c}}{\tilde{p}} = \frac{\lambda}{p}$. *There exists a time bound* $T'' \in \mathbb{R}_+$ *such that, for all* $T \geq T''$, *the set of symmetric equilibrium distributions in the contest and in the all-pay auction coincide.*

Therefore, the all-pay auction is suitable for the analysis of stochastic contests if the contest lasts long enough and the research process has a cumulative structure.

The next paragraph scrutinizes the amount of time needed for the equilibrium distributions to be feasible. More precisely, we derive an upper bound on the minimal time to reach the uniform distribution; recall that in a two-player contest this is the unique Nash equilibrium for the appropriate choice of prize.

Proposition 15. *The uniform distribution on* $\{0, 1, \ldots, \overline{x}\}$ *is feasible if the contest lasts at least* $T = \frac{\overline{x}\log(\overline{x}+1)}{\lambda}$ *periods.*

The construction of the upper bound in the proof can be applied to derive an upper bound for any other (equilibrium) distribution.

While it is beyond the scope of this chapter to give a full characterization of the equilibria for small values of T, we illustrate the main intuition with a simple example. In the example, we let $n = 2, \tilde{p} = 3, \lambda = 1, \tilde{c} = 1$. If $T < \log(3) \approx 1.0986$, it is straightforward to verify with the infinitesimal generator condition that, in the unique equilibrium, players continue until they have one success or the game ends. For larger values of T, players also continue if they have one success, since the probability mass of the other player at one or two successes is sufficiently high. For $T \geq \log(3) - \log(\frac{1}{\log(3)}) \approx 1.1927$, the equilibrium distribution of the all-pay auction with $p = 3$, which places probability one-third on zero, one, and two successes, respectively, is feasible.

Hence, in short contests, players do not have enough time to transfer as much probability mass to higher values compared to the all-pay auction equilibrium. This restriction on the strategies leads to a different result than restricting the strategy space with a binding bidding cap in the all-pay auction. More specifically, a binding bidding cap leads to a higher probability mass at the bidding cap, while a low value of contest length leads to a concentration of the probability mass at lower values.

3.4 Conclusion

In this chapter, we have compared the static all-pay auction with discrete bids to a stochastic contest model in which the research success is weakly increasing as long as a player exerts effort. If the contest lasts long enough—a moderate contest length is sufficient—and has a cumulative structure, the symmetric equilibrium distributions of the contest coincide with those of the discrete all-pay auction.

Hence, the applicability of all-pay auction models to analyze stochastic contests depends on the structure of the research process. For a cumulative process, the all-pay auction is appropriate, while it is not suited to model a contest in which only the best research outcome in a single period counts as in Taylor (1995).

3.5 Appendix

Proof of Lemma 18. By contradiction, assume $\exists x \in \mathbb{N}_0 : x < \overline{x}$ such that the pure strategy x is strictly worse than all $\tilde{x} \in supp\{f\}$. Hence, in any symmetric equilibrium $f^i(x) = 0$ for all i. Define $x' := \min\{\hat{x} > x\colon$ player i is indifferent between $\hat{x}$ and any $\tilde{x} \in supp\{f\}\}$ as the lowest optimal bid above x. Hence, $f^i(x'-1) = 0$ for all players, which implies that $\mathbb{P}(\max_{j \neq i} x^j = x'-1) = 0$. Yet, as player i strictly prefers x' to $x'-1$, we get

$$\mathbb{P}(\max_{j \neq i} x^j = x')\mathbb{E}(\frac{1}{m^i(x')}) > \frac{1}{p} \tag{3.5.1}$$

with $m^i(x') = |\{j \in N : j \neq i, x^j = x' = \max_{l \neq i} x^l\}| + 1$. On the other hand, player i prefers, at least weakly, x' to $x'+1$, i.e.,

$$\left(1 - \mathbb{E}(\frac{1}{m^i(x')})\right) \mathbb{P}(\max_{j \neq i} x^j = x') + \mathbb{P}(\max_{j \neq i} x^j = x'+1)\mathbb{E}(\frac{1}{m^i(x'+1)}) \leq \frac{1}{p}. \tag{3.5.2}$$

Due to $\mathbb{E}(\frac{1}{m^i(x')}) \leq \frac{1}{2}$, the left hand side of this equation is at least as large as $\mathbb{P}(\max_{j \neq i} x^j = x')\mathbb{E}(\frac{1}{m^i(x')})$. Thus, (3.5.1) and (3.5.2) yield a contradiction, which completes the proof of the first part of Lemma 18.

If $n > 2$, then $\mathbb{E}(\frac{1}{m^i(x')}) < \frac{1}{2}$ for all points in $\{0, 1, \ldots, \overline{x}\}$ for which $f(x) > 0$. It follows that $f(x) > 0$ if and only $x \leq \overline{x}$, as the contradiction is valid even if (3.5.1) holds with equality. For two players, we establish this result in the proof of Proposition 12. □

Proof of Proposition 12. In any equilibrium in which both players are indifferent between all strategies $\{0, 1, \ldots, \overline{x}\}$, we have

$$\begin{aligned} &\frac{f^i(x)}{2} + \frac{f^i(x+1)}{2} = \frac{1}{p} \;\; \forall x < \overline{x} - 1, \; \forall i, \quad \text{and} \\ &\sum_0^{\overline{x}} f^i(x) = 1. \end{aligned} \tag{3.5.3}$$

The unique solution to this system of equations for $\frac{p}{2} \notin \mathbb{N}$ is given in Proposition 12. Hence, by the indifference argument in Lemma 18, the equilibrium is the unique symmetric equilibrium. This also implies that $f(x) > 0$ if and only if $x \leq \overline{x}$ in the symmetric equilibrium for two players, which completes the missing step in the proof of Lemma 18.

In the following, we show that the two-player game does not have an asymmetric equilibrium. First, we prove that $f^i(x+2) > 0 \implies f^i(x+1) > 0$ or $f^i(x) > 0$, for all $x \in \mathbb{N}_0$, for all i. Assume to the contrary

$$\exists x \in \mathbb{N}, i \in N \colon f^i(x) = 0, f^i(x+1) = 0 \text{ and } f^i(x+2) > 0.$$

Optimality for player j implies

$$f^j(x) \geq 0, f^j(x+1) = 0,$$

because her winning probability does not change on $\{x, x+1\}$, but she has higher costs for any point above x. Moreover, if $f^i(x+2) > 0$ is optimal, $f^j(x+2) > 0$, because player i would prefer $x+1$ to $x+2$ otherwise. Yet, in this case, player j prefers, at least weakly, to bid $x+2$ compared to x, which means

$$\frac{2}{p} \leq \frac{f^i(x+2)}{2}. \tag{3.5.4}$$

However, comparing the profits of (the supposedly optimal) $x+2$ and $x+3$ for player j gives

$$p\,\frac{f^i(x+2) + f^i(x+3)}{2} - 1\,,$$

which is strictly positive by equation (3.5.4). Thus, player j strictly prefers to bid $x+3$ compared to $x+2$. This yields the required contradiction to optimality of bidding $x+2$. Hence, for any two subsequent points smaller than $\overline{x}$, there is at most one point without mass.

In the next step, we show that players are indifferent on $\{0, \dots, \overline{x}\}$. Assume to the contrary that player j finds it strictly worse to bid $x < \overline{x}$. Thus, by the first step, $f^j(x-1) > 0, f^j(x) = 0$ and $f^j(x+1) > 0$. Furthermore, as it is strictly worse for player j to play x compared to $x+1$ and $x-1$, we have

$$\frac{f^i(x+1)}{2} + \frac{f^i(x)}{2} > \frac{1}{p} \quad \text{ and } \quad \frac{f^i(x-1)}{2} + \frac{f^i(x)}{2} < \frac{1}{p},$$

which in turn implies $f^i(x+1) > f^i(x-1)$. Since player i (weakly) prefers $x+1$ to $x-1$, we have

$$\frac{f^j(x-1) + f^j(x+1)}{2} \geq \frac{2}{p}. \tag{3.5.5}$$

As player i also (weakly) prefers $x+1$ to x and $x+2$ respectively, it holds

$$\frac{f^j(x) + f^j(x+1)}{2} \geq \frac{1}{p} \quad \text{ and } \quad \frac{f^j(x+1) + f^j(x+2)}{2} \leq \frac{1}{p}. \tag{3.5.6}$$

This yields $f^j(x-1) = f^j(x+1) = \frac{2}{p}$ and $f^j(x+2) = 0$. If $f^i(x-1) > 0$, equation (3.5.5) holds with equality and, by optimality for player i, we also have $f^j(x-2) = 0$ (if existent). If, on the other hand, $f^i(x-1) = 0$, the same result, $f^j(x-2) = 0$, holds by the first part of the proof. This argument extends to the whole support and any resulting function f which alternates between 0 and $\frac{2}{p}$ violates the probability measure condition (3.5.3) as $\frac{p}{2} \notin \mathbb{N}$ by Assumption 3.

It remains to show indifference on 0. Assume that to bid 0 is strictly worse than to bid 1 for player i. If both players have $f^i(0) = f^j(0) = 0$, weak optimality of a bid of 1 compared to 2 implies that $f(1) \leq \frac{2}{p}$. This has to hold with equality to guarantee non-negative profits. Thus, players are indifferent between bids of 0 and 1.

Therefore, it remains to consider the case in which exactly one player, say player j, has $f^j(0) > 0$. In this case, however, her expected profit is zero at zero and it has to be zero for any strategy in the support because she is indifferent between all strategies in the support. Given $f^i(0) = 0$, this implies $f^i(1) = \frac{2}{p}$. As before, the argument extends to the whole support and any resulting function f which alternates between 0 and $\frac{2}{p}$ violates the probability measure condition (3.5.3).

Thus, in any equilibrium, players have to be indifferent on $\{0, \ldots \overline{x}\}$. In addition, the probability measure condition (3.5.3) must hold. By the first part of the proof, this determines the equilibrium distributions uniquely. $\square$

Proof of Lemma 19. Rost (1976) constructs a stopping time τ for general right continuous Markov processes $(X_t)_{t\in\mathbb{R}_+}$ that minimizes the residual expectation $\mathbb{E}(\int_{\min\{t,\tau\}}^{\tau} \mathrm{d}s) = \int_t^{\infty} \mathbb{P}(\tau > s)\mathrm{d}s$ for all $t \in \mathbb{R}_+$ and embeds F, i.e., $X_\tau \sim F$. To apply Rost, we need to verify the condition on page 198 (Rost, 1976):
For every positive function $g\colon \mathbb{N}_0 \to \mathbb{R}$,

$$\mathbb{E}(\int_0^{\infty} g(X_t)\mathrm{d}t) \geq \mathbb{E}^F(\int_0^{\infty} g(X_t)\mathrm{d}t).$$

In the following, we show that this condition is fulfilled

$$\mathbb{E}^F(\int_0^\infty g(X_t)\mathrm{d}t) = \sum_{x=0}^{\overline{x}} f(x) \sum_{i=x}^{\infty} \frac{1}{\lambda} g(i) = \sum_{x=0}^{\overline{x}} f(x) \left(\sum_{i=0}^{\infty} \frac{1}{\lambda} g(i) - \sum_{i=0}^{x-1} \frac{1}{\lambda} g(i) \right)$$
$$\leq \sum_{i=0}^{\infty} \frac{1}{\lambda} g(i) = \mathbb{E}(\int_0^\infty g(X_t) dt) \,.$$

By the Lemma on page 201 in Rost (1976), the associated stopping time is the first hitting time of a set A

$$\tau = \inf\{t \in \mathbb{R}_+ | (t, X_t) \in A\} \,. \tag{3.5.7}$$

We can equivalently formulate (3.5.7) as

$$\tau = \inf\{t \in \mathbb{R}_+ | t \geq H(X_t)\} \,,$$

with some function $H\colon \mathbb{N}_0 \to \mathbb{R}_+$. As the density f is positive for all points in the support, $H(x)$ is finite. Since the support of the distribution F is a finite number of points, the minimum of H on the support exists and is finite $H^\star = \min_{0 \leq x \leq \overline{x}} H(x)$. This stopping time embeds F in a minimal time $H^\star$ with $\tau \leq H^\star$ almost surely. □

Proof of Proposition 14. By Proposition 13, any symmetric equilibrium distribution in the all-pay auction is also an equilibrium distribution in the stochastic contest. Therefore, it remains to scrutinize whether there are additional symmetric equilibria in the stochastic contest. To be feasible, an equilibrium distribution in the stochastic contest has $f(x) > 0$ if and only if $x \leq \overline{x}$.

Assume to the contrary that there exists an equilibrium distribution in the stochastic contest that is not an equilibrium distribution in the all-pay auction. In this case, a player strictly prefers to continue until the next success for at least one point. We denote the largest of these points by $\tilde{x}$. In the following, we show that the player strictly wants to continue at $\tilde{x}$, but is at most indifferent whether to continue at $\tilde{x} + 1$, contradict each other for $n > 2$,

$$f^{-i}(\tilde{x})(1 - \mathbb{E}(\frac{1}{m^i(\tilde{x})})) + f^{-i}(\tilde{x}+1)\mathbb{E}(\frac{1}{m^i(\tilde{x}+1)}) > \frac{\tilde{c}}{\tilde{p}\lambda}$$
$$f^{-i}(\tilde{x}+1)(1 - \mathbb{E}(\frac{1}{m^i(\tilde{x}+1)})) + f^{-i}(\tilde{x}+2)\mathbb{E}(\frac{1}{m^i(\tilde{x}+2)}) \leq \frac{\tilde{c}}{\tilde{p}\lambda}.$$

As $f^i(\tilde{x})$ approaches zero as $T \to \infty$ and $\mathbb{E}(\frac{1}{m^i(x)}) < 1 - \mathbb{E}(\frac{1}{m^i(x)})$ for all $x \leq \bar{x}$ and $n > 2$, these equations contradict each other for T large enough.

We now consider the remaining case $n = 2$. As before, denote the highest point at which a player strictly prefers to continue by $\tilde{x}$ and $f^i(\tilde{x})$ by $\tilde{\epsilon}$. As a player prefers to continue at $\tilde{x}$, but weakly prefers to stop at $\tilde{x}+1$, $f(\tilde{x}+1) \in (\frac{2}{p} - \tilde{\epsilon}, \frac{2}{p}]$. Consequently, the expected gain of continuing at $\tilde{x}$ is at most $\tilde{\epsilon}p$ – the expected gain of continuation until the next success, if the player could continue to play forever. To make continuation at least weakly optimal at $\tilde{x}-1$ for some $t < T$, $f(\tilde{x}-1) \in [\frac{2}{p} - 2\tilde{\epsilon}, \frac{2}{p} - \tilde{\epsilon}]$. Hence, either $f(\tilde{x}-2) \leq 2\tilde{\epsilon}$ or the player always continues to play there. This argument extends to the whole support.

Hence, $\sum_{k=0}^{\overline{x}} f(k) = \frac{2}{p}l + \hat{\epsilon}$, with a natural number $l \in \mathbb{N}$ and $\hat{\epsilon}$ the sum over the differences from 0 or $\frac{2}{p}$ at all points of the support,

$$\hat{\epsilon} = \sum_{k=0}^{\overline{x}} (-1)^{k+\mathbf{1}_{\{p<f(0)\}}} \min\{f(k), \frac{2}{p} - f(k)\}.$$

As $T \to \infty$, $\overline{x}$ is bounded, since $\sum_{i=0}^{\overline{x}} f(k) > 1$ otherwise. Thus, $\hat{\epsilon}$ converges to zero as $T \to \infty$. By Assumption 3, $\frac{p}{2} \neq \mathbb{N}$. Consequently, there exists a time T', such that, for all $T \geq T'$, $\sum_{i=0}^{\overline{x}} f(k) \neq 1$, which contradicts the probability measure condition (3.5.3).

$\square$

Proof of Proposition 15. We construct a strategy which waits at some points in time, calculate its implementation time, and argue that it is larger than that of the optimal strategy.

The strategy proceeds as follows: continue at 0 until $\mathbb{P}(X_t = 0) = \frac{1}{\overline{x}+1}$ and stop at 1. Then continue at 1 until $\mathbb{P}(X_t = 1) = \frac{1}{\overline{x}+1}$ and stop at 2. Continue in the same way for the whole support.

To implement the first step, we get the condition $e^{-\lambda t} = \frac{1}{\overline{x}+1}$, i.e., $t = \frac{\log(\overline{x}+1)}{\lambda}$. Any later step takes less time, as less mass needs to be transferred. Hence, we can bound the required time by $T \leq \frac{\overline{x}\log(\overline{x}+1)}{\lambda}$. Clearly, this strategy is slower than the optimal strategy, which never stops in order to continue later. $\square$

Chapter 4

How Burning Money Requires a Lot of Rationality To Be Effective

In this chapter, we propose an extension of the valuation equilibrium concept (Jehiel and Samet, 2007). Contrary to the original concept, we do not treat the grouping of moves as entirely exogenous, but start with an intuitive basic grouping which players can refine at a lexicographical cost. We then adapt the valuation equilibrium concept to this setting. This approach yields predictions for extensive-form games without specifying an underlying grouping ad hoc. In an application to a burning money game, we find that adding a possibility to burn money prior to a coordination game does not affect the set of equilibrium outcomes. This prediction, which differs from that of standard solution concepts, is roughly in line with laboratory evidence gathered by Huck and Müller (2005).

4.1 Introduction

An interesting recent approach to model bounded rationality is the valuation equilibrium concept (Jehiel and Samet, 2007). The idea of the concept is that (boundedly rational) players, instead of considering all moves in a game separately, bunch "similar" moves into groups ("similarity classes") and attach a valuation to each similarity class. In equilibrium, each player chooses actions from a similarity class with locally maximal valuation, where “locally maximal” refers to the highest valuation of all similarity classes available at that point. The underlying grouping in terms of similarity, however, remains entirely exogenous.

We propose a refinement of the valuation equilibrium concept that partially endogenizes the grouping of actions. More specifically, we start with an intuitive basic grouping, which players can refine. A more complex grouping, however, enters the payoff function of a player at a lexicographical cost. The measure for complexity of a grouping is the number of similarity classes it contains. Intuitively, this criterion is similar to the complexity criterion in Rubinstein’s games played by finite automata (Rubinstein, 1986, Abreu and Rubinstein, 1988), where the number of states of an automaton enters the payoff in a lexicographical way. The refined equilibrium concept requires that no player has a benefit of using a different partition and strategy.

To illustrate the virtues of the modified definition, we compare its predictions with traditional predictions for the burning money game introduced by Ben-Porath and Dekel (1992).[1] The comparison emphasizes how standard equilibrium concepts rely on player 2 fully exploiting all available information. In particular, if differentiating moves comes at a lexicographical cost for player 2, money is never burned in equilibrium and all three equilibria of the one-shot coordination game are possible in the second stage. Hence, the modified valuation concept selects those subgame perfect equilibria in which the mere possibility to burn money does *not* affect the outcome of the coordination game. In fact, this is roughly in line with empirical evidence from a repeated version of the game gathered by Huck and Müller (2005): money is almost never burned (about 6 percent); on the second stage, all strategy profiles are observed in significant fractions, but the equilibrium strategy profile player 1 prefers occurs with highest probability; this profile is also selected by the forward induction solution in Ben-Porath and Dekel (1992).

Finally, we discuss a sequential version of matching pennies to provide more intuition on how the modified valuation concept works. Again, the concept selects a subset of subgame

[1]The burning money game combines a common coordination game with a first stage in which player 1 has the opportunity to “burn money”, i.e., choose a strategy which reduces his payoff.

perfect equilibria.[2] In contrast to the first example, there is no modified valuation equilibrium in pure strategies and the second player uses a refined grouping in any equilibrium.

4.2 The Model

The first part of this section briefly introduces the sequential valuation equilibrium concept by Jehiel and Samet (2007) with a slightly simplified notation. The second part introduces the notion of a *robust valuation equilibrium.*

4.2.1 The Basic Setup of Jehiel and Samet (2007)

Following Jehiel and Samet (2007), we consider finite extensive-form games, where $\mathcal{I}$ is the finite set of players $i = 1, \ldots, I$ and Z is the finite set of terminal histories induced by the players' actions. The payoff of each player at each terminal history is specified by a function $f_i : Z \to \mathbb{R}$. Each player i uses a behavioral strategy σ_i. The probability distribution over terminal histories induced by the player's strategies is denoted $\mathbb{P}^\sigma$. Moreover, the set of moves M_i is partitioned into similarity classes $\lambda_i \in \Lambda_i$. Each player attaches a *valuation* $\nu_i : \Lambda_i \to \mathbb{R}$ to each similarity class. Denote the probability that at least one move from similarity class λ_i is played, i.e., the probability that a terminal history containing a move in λ_i is reached, by $\mathbb{P}^\sigma(Z(\lambda_i))$.

In a sequential valuation equilibrium, strategies have to be *optimal* given valuations and valuations have to be *sequentially consistent* with the strategies. Roughly speaking, consistency requires that the valuation of a similarity class equals the expected payoff (given all players' strategies) a player obtains conditional on having used a move from the respective class.

[2]In general, the set of modified valuation equilibria is not a subset of the set of subgame perfect equilibria.

More formally, a valuation ν_i is *consistent* with a strategy profile σ if, for all $\lambda_i \in \Lambda_i$ with $\mathbb{P}^\sigma(Z(\lambda_i)) > 0$,

$$\nu_i(\lambda_i) = \mathbb{E}^\sigma(f_i|Z(\lambda_i)) = \sum_{z \in Z(\lambda_i)} \frac{\mathbb{P}^\sigma(z) f_i(z)}{\mathbb{P}^\sigma(Z(\lambda_i))}.$$

ν_i is *sequentially consistent*, if there is a fully mixed sequence $(\sigma^k)_{k=1}^\infty$ such that $\sigma^k \to \sigma$ and $\nu_i^k \to \nu_i$, where ν_i^k is the unique consistent valuation for σ^k.

Moreover, a strategy σ_i is *optimal* given ν_i, if player i only chooses actions from similarity classes with the (locally) highest valuation.

As in Wichardt (2009), we additionally restrict the set of admissible strategies to those which use the same probability distribution whenever the set of optimal actions is identical. We refer to this as *uniform tie-breaking*.

For a more extensive introduction and additional intuition about the basic setup, see Jehiel and Samet (2007) or Wichardt (2009).

4.2.2 Restricting the Partitions

To make the concept applicable to a game without having to specify a partition entirely ad hoc, we restrict the set of partitions in two ways. First, we define a primary partition of moves, which captures a priori similarity of actions. For the present purposes, we assume that a priori similarity is reflected in the labeling of actions.

Definition 2. In the *primary partition* Λ_i^0 of player i, two moves are in the same similarity class if and only if they have the same label. In any (admissible) partition, two moves are in different similarity classes if they have different labels.

Moreover, we define a measure of complexity on the set of (admissible) partitions:

Definition 3. A partition Λ_i is (weakly) finer than Λ_i' if Λ_i contains at least as many similarity classes as Λ_i'. Otherwise, Λ_i is coarser than Λ_i'.

Intuitively, a more complex partition requires closer attention to the game, which is more "costly" in terms of memory. We include the measure of complexity in the equilibrium concept:[3]

Definition 4. A profile (σ, ν, Λ) is a *robust valuation equilibria (RVE)* if, for all i,

1. ν_i is sequentially consistent with σ,
2. σ_i is optimal given ν_i,
3. for any coarser partition Λ_i' and all pairs (σ_i', ν_i') allowed by Λ_i' such that σ_i' is optimal given ν_i' and ν_i' is sequentially consistent with σ_i', we have $\mathbb{E}[f_i|(\sigma_i', \sigma_{-i})] < \mathbb{E}[f_i|\sigma]$,
4. for any (weakly) finer partition Λ_i'' and all pairs (σ_i'', ν_i'') allowed by Λ_i'' such that σ_i'' is optimal given ν_i'' and ν_i'' is sequentially consistent with σ_i'', we have $\mathbb{E}[f_i|(\sigma_i'', \sigma_{-i})] \leq \mathbb{E}[f_i|\sigma]$.

The first two conditions ensure that the profile is a sequential valuation equilibrium given the partition. Conditions 3 and 4 impose additional robustness criteria. In particular, condition 3 states that no player can obtain the same payoff with a "cheaper" partition; condition 4 states that no player can get a strictly higher payoff with a finer partition.

Intuitively, the robust valuation concept resembles Rubinstein's concept for repeated games played by finite automata (Rubinstein, 1986, Abreu and Rubinstein, 1988) in which the complexity (number of states) of an automaton enters the payoff at a lexicographical cost. Note that in Rubinstein's concept, transitions between the states of an automaton occur depending on the sequence of moves, whereas here the grouping is not necessarily tied to the preceding history.

[3]The modified concept is equivalent to introducing a lexicographic cost for choosing a finer partition.

4.3 Applications

4.3.1 The Burning Money Game

In this section, we discuss an application of the robust valuation equilibrium concept to the burning money game introduced by Ben-Porath and Dekel (1992). In this 2-stage, 2-player game, player 1 first decides whether to play b ("burn money") or nb ("not burn money"), where b reduces the payoff of player 1 in stage 2, but has no effect otherwise. The second stage is a coordination game, where player 1 can choose between U and D and player 2 can choose L or R. The payoffs are specified in Figure 4.1, where subindices are added for notational convenience.

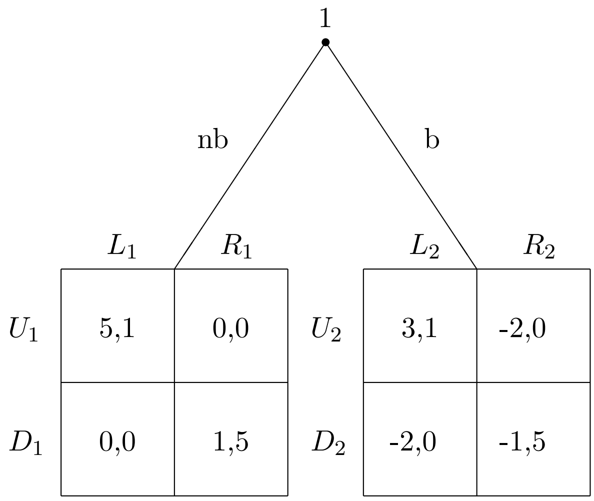

Figure 4.1: The burning money game; in stage 2 player 1 chooses rows and player 2 chooses columns.

Intuitively, it is not clear why an option to "burn money", i.e., publicly harm oneself, should have an influence on the equilibrium set—at least if players are not assumed to be fully rational. Yet, for standard equilibrium concepts such as Nash equilibrium, subgame perfect equilibrium, or forward induction, the additional option of "burning money" does matter.

The game has a myriad of subgame perfect equilibria. In particular, "burning money" can be sustained in equilibrium, e.g., in the profile $\sigma = ((b, D_1U_2), R_1L_2)$. The RVE concept, however, predicts that money is never burned and equilibrium play at the second stage is independent of the first stage:

Proposition 16. *In any RVE, each player uses his primary partition Λ_i^0, where $\Lambda_1^0 = \{\{nb\}, \{b\}, \{U_1, U_2\}, \{D_1, D_2\}\}$ and $\Lambda_2^0 = \{\{L_1, L_2\}, \{R_1, R_2\}\}$ and equilibrium strategies are given by*

(i) player 1 choosing nb and U_1U_2, player 2 choosing L_1L_2,

(ii) player 1 choosing nb and D_1D_2, player 2 choosing R_1R_2,

(iii) player 1 choosing nb and $\frac{5}{6}U_1 + \frac{1}{6}D_1$; $\frac{5}{6}U_2 + \frac{1}{6}D_2$, player 2 choosing $\frac{1}{6}L_1 + \frac{5}{6}R_1$; $\frac{1}{6}L_2 + \frac{5}{6}R_2$.

Proof. If player 2 uses the primary partition Λ_2^0, consistent valuations for every fully mixed sequence of strategies have to satisfy $\nu_1^k(\{b\}) + 2 = \nu_1^k(\{nb\})$, since player 2 uses the same probability distribution over actions after b and nb by uniform tie-breaking. Hence, nb is played in any RVE for Λ_2^0. Furthermore, by sequential consistency, mutual best responses are played on the equilibrium path at the second stage. Hence, (i)-(iii) are the only candidates for an RVE for Λ_2^0, since off-equilibrium path (after b is played), both players play the same actions as on the equilibrium path by uniform tie-breaking. As each player obtains the highest possible payoff against the strategy of the other player with his primary partition, the strategy-partition combinations in the proposition are robust to finer partitioning and robust to coarser partitioning, and, hence an RVE.

It remains to establish that the game possesses no RVE for a partition that is finer than Λ_2^0. To see this, note that unless player 1 randomizes at the first stage, player 2 has no incentive to use a refined partition: one branch of the tree is never followed and he can ensure himself whatever payoff he gets using the (cheaper) primary partition. Hence, in any RVE, for player 2 to use a partition finer than Λ_2^0, player 1 has to randomize at the first stage. This implies $\nu_1(\{b\}) = \nu_1(\{nb\})$ by optimality. Moreover, both players have to play mutual best responses after b and nb in the second stage. Taken together, this leads to a contradiction, as for any two equilibria in the one-shot game after b and nb, the payoff of player 1 is different (so that the valuation for $\{b\}$ and $\{nb\}$ have to differ). Accordingly, the above characterization captures all RVE of the game. □

Intuitively, "burning money" does not have an effect because player 2 would have to invest in a finer partitioning in order to take it into account. In that sense, the robust valuation equilibrium concept highlights the comparably high degree of rationality implicitly ascribed to player 2 in standard game-theoretic concepts.

Huck and Müller (2005) tested the predictions of the burning money game in an experiment. Their findings are roughly in line with the predictions of the robust valuation equilibrium concept: In a repeated version of the game, money is almost never burned (about 6 percent of the games) and all possible action profiles at the second stage occur in significant fraction, while the equilibrium player 1 prefers occurs with the highest probability.[4] This equilibrium is also the outcome of the forward induction solution by Ben-Porath and Dekel (1992).

4.3.2 Sequential Matching Pennies

Finally, we discuss a sequential version of matching pennies to further clarify how the RVE concept works.

In this sequential matching pennies game, player 1 can first choose whether to play heads (H) or tails (T). At stage 2, player 2 decides whether to play H or T. Final payoffs correspond to the classical matching pennies game and are given in Figure 4.2, where subindices are again added for notational convenience.

The standard equilibrium prediction for this game (subgame perfect equilibrium) is that player 2 plays his best response H_1T_2 and player 1 can use any strategy profile in equilibrium. The RVE predictions are slightly different:

[4] Exact frequencies in Huck and Müller (2005) vary depending on the exact assignment of the game.

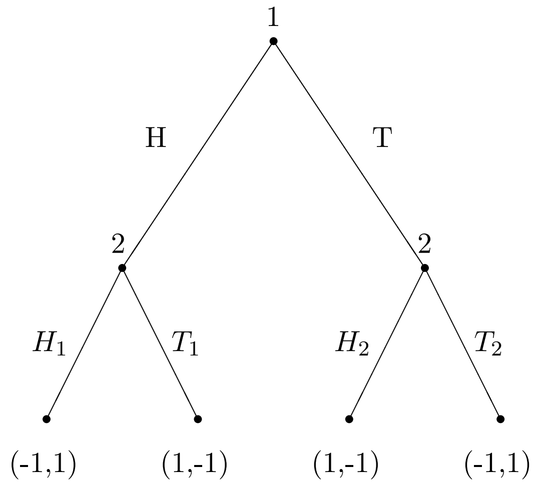

Figure 4.2: Sequential matching pennies game

Proposition 17. *In any RVE, player 2 uses a partition, which contains three similarity classes (either $\Lambda_2 = \{\{H_1\}, \{H_2\}, \{T_1, T_2\}\}$ or $\Lambda_2 = \{\{H_1, H_2\}, \{T_1\}, \{T_2\}\}$). Player 1 randomizes between H and T with probability $\mathbb{P}(H) \in (0, 1)$. Player 2 plays H_1T_2.*

Proof. If player 2 uses a partition with three similarity classes and the strategy H_1T_2 (with consistent valuations $\nu_2 = (1, -1, 1)$), she can guarantee herself the maximal payoff of one. If she does so, any strategy of player 1 leads to payoff of -1 and is thus sequentially consistent with a valuation of -1 for both similarity classes. It remains to verify whether the strategy of player 2 is robust to coarser partitioning. If player 1 does not play a pure strategy, any deviation to the coarser partition Λ_2^0 decreases the payoff of player 2, because she has to use a uniform tie-breaking over actions at both nodes. Hence, the strategy profiles in the proposition are RVE.

On the other hand, if player 1 plays a pure strategy, player 2 is better off deviating to the coarser partition with the best response to the pure strategy. Hence, there are no other RVE in which player 2 uses a partition with three similarity classes.

Clearly, no strategy profile for the finest partition $\Lambda_2 = \{\{H_1\}, \{H_2\}, \{T_1\}, \{T_2\}\}$ can be robust to coarser partitioning, since, for any strategy of player 1, there is a partition containing three similarity classes for which player 2 can obtain the maximal payoff.

In any RVE candidate for Λ_2^0, the expected payoff of player 2 has to be 1, since she would deviate to a finer partition and H_1T_2 otherwise. This implies that player 2 plays either H_1H_2 or T_1T_2. In this case, however, player 1 has a strategy, which guarantees her a payoff of 1 and player 2 a payoff of -1. Hence, there are no RVE in which player 2 uses the coarsest partition. $\square$

In this example, the RVE criterion rules out those subgame perfect equilibria in which player 1 uses a pure strategy. Intuitively, the mixed-strategy equilibria survive because player 1 forces player 2 to use his finest partition.

4.4 Discussion

This chapter has proposed a refinement of the valuation equilibrium concept, which yields predictions for extensive-form games without requiring an ad hoc specification of each player's groupings. In the main application, the predictions about the equilibrium set differ from those of standard concepts, but are intuitive and roughly in line with empirical evidence.

In future research, one might investigate the relation of RVE to other equilibrium concepts in more detail. In this respect, it would be useful to extend the definition of the concept to infinitely repeated games and compare its predictions to the predictions of subgame perfect equilibria and equilibria in games played by finite automata (Rubinstein, 1986, Abreu and Rubinstein, 1988).

Another open question is to provide general conditions for existence of an RVE. For similar reasons as in other evolutionary or bounded rationality concepts, it is possible to construct a game in which no RVE exists. Relatedly, it would be interesting to have an explicit evolutionary foundation for the RVE concept based on a model in which players experiment with and learn about strategies and partitions over time.

Bibliography

Abreu, D. and A. Rubinstein (1988): "The Structure of Nash Equilibria in Repeated Games with Finite Automata," *Econometrica*, 56, 1259–1281.

Amann, E. and W. Leininger (1996): "Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case," *Games and Economic Behavior*, 14, 1–18.

Anderson, A. and L. M. B. Cabral (2007): "Go for broke or play it safe? Dynamic competition with choice of variance," *RAND Journal of Economics*, 38, 593–609.

Ankirchner, S., G. Heyne, and P. Imkeller (2008): "A BSDE approach to the Skorokhod embedding problem for the Brownian motion with drift," *Stochastics and Dynamics*, 8, 35–46.

Ankirchner, S. and P. Strack (2011): "Skorokhod Embeddings in Bounded Time," *Stochastics and Dynamics*, forthcoming.

Aoki, R. (1991): "R&D Competition for Product Innovation: An Endless Race," *American Economic Review*, 81, 252–256.

Bass, R. F. (1983): "Skorokhod Imbedding via Stochastic Integrals," in *Seminar on probability, XVII*, Berlin: Springer, vol. 986 of *Lecture Notes in Math.*, 221–224.

Baye, M., D. Kovenock, and C. de Vries (1993): "Rigging the Lobbying Process," *American Economic Review*, 83, 289–294.

——— (1994): "The solution to the Tullock rent-seeking game when $R > 2$: Mixed-strategy equilibria and mean dissipation rates," *Public Choice*, 81, 363–380.

——— (1996): "The all-pay auction with complete information," *Economic Theory*, 8, 291–305.

Ben-Porath, E. and E. Dekel (1992): "Signaling Future Actions and the Potential for Sacrifice," *Journal of Economic Theory*, 57, 36–51.

Bulow, J. and P. Klemperer (1999): "The Generalized War of Attrition," *American Economic Review*, 89, 175–189.

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

Che, Y.-K. and I. Gale (1998): "Caps on Political Lobbying," *American Economic Review*, 88, 643–651.

Cohen, C. and A. Sela (2007): "Contests with Ties," *The B.E. Journal of Theoretical Economics*, 7, 1–16.

Dechenaux, E., D. Kovenock, and V. Lugovskyy (2003): "A Comment on 'David vs. Goliath: An Analysis of Asymmetric Mixed-Strategy Games and Experimental Evidence'," *mimeo*.

——— (2006): "Caps on bidding in all-pay auctions: Comments on the experiments of A. Rapoport and W. Amaldoss," *Journal of Economic Behavior and Organization*, 61, 276–283.

Downs, G. W. and D. M. Rocke (1994): "Conflict, Agency, and Gambling for Resurrection: The Principal-Agent Problem Goes to War," *American Journal of Political Science*, 38, 362–380.

Fudenberg, D. and J. Tirole (1986): "A Theory of Exit in Duopoly," *Econometrica*, 54, 943–960.

Grandits, P. and N. Falkner (2000): "Embedding in Brownian motion with drift and the Azéma-Yor construction," *Stochastic Processes and their Applications*, 85, 249–254.

GUL, F. AND W. PESENDORFER (2011): “The War of Information,” *Review of Economic Studies*, forthcoming.

HARRIS, C. AND J. VICKERS (1987): “Racing with Uncertainty,” *Review of Economic Studies*, 54, 1–21.

HILLMAN, A. AND J. RILEY (1989): “Politically Contestable Rents and Transfers,” *Economics and Politics*, 1, 17–39.

HILLMAN, A. AND D. SAMET (1987): “Dissipation of contestable rents by small numbers of contenders,” *Public Choice*, 54, 63–82.

HÖRNER, J. (2004): “A Perpetual Race to Stay Ahead,” *Review of Economic Studies*, 71, 1065–1088.

HUCK, S. AND W. MÜLLER (2005): “Burning money and (pseudo) first-mover advantages: an experimental study on forward induction,” *Games and Economic Behavior*, 51, 109–127.

HVIDE, H. (2002): “Tournament Rewards and Risk Taking,” *Journal of Labor Economics*, 20, 877–898.

JEHIEL, P. AND D. SAMET (2007): “Valuation Equilibrium,” *Theoretical Economics*, 2, 163–185.

KARLIN, S. (1953): “Reduction of certain classes of games to integral equations,” in: H. Kuhn & A. Tucker (Eds.), ‘Contributions to the Theory of Games, II.’, Vol.28 of *Annals of Mathematical Studies*, Princeton University Press, Princeton, pp.125-128.

KOHLBERG, E. AND J.-F. MERTENS (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54, 1003–1037.

KRISHNA, V. AND J. MORGAN (1997): “An Analysis of the War of Attrition and the All-Pay Auction,” *Journal of Economic Theory*, 72, 343–362.

LAZEAR, E. P. AND S. ROSEN (1981): “Rank-Order Tournaments as Optimum Labor Contracts,” *Journal of Political Economy*, 89, 841–864.

MAYNARD SMITH, J. (1974): "The Theory of Games and the Evolution of Animal Conflicts," *Journal of Theoretical Biology*, 47, 209–221.

MOLDOVANU, B. AND A. SELA (2001): "The Optimal Allocation of Prizes in Contests," *American Economic Review*, 91, 542–558.

MOSCARINI, G. AND L. SMITH (2007): "Optimal Dynamic Contests," *Working Paper.*

OBŁÓJ, J. (2004): "The Skorokhod embedding problem and its offspring," *Probability Surveys*, 1, 321–392.

PARK, A. AND L. SMITH (2008): "Caller Number Five and related timing games," *Theoretical Economics*, 3, 231–256.

PEDERSEN, J. L. AND G. PESKIR (2001): "The Azéma-Yor embedding in non-singular diffusions," *Stochastic Processes and their Applications*, 96, 305–312.

ROGERS, L. C. G. AND D. WILLIAMS (2000): *Diffusions, Markov Processes and Martingales: 1*, Cambridge University Press.

ROST, H. (1976): "Skorokhod stopping times of minimal variance," *Séminaire de Probabilités (Strasbourg)*, 10, 194–208.

RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83–96.

SIEGEL, R. (2009): "All-Pay Contests," *Econometrica*, 77, 71–92.

——— (2010): "Asymmetric Contests with Conditional Investments," *American Economic Review*, 100, 2230–2260.

SKOROKHOD, A. V. (1961): "Issledovaniya po teorii sluchainykh protsessov (Stokhasticheskie differentsialnye uravneniya i predelnye teoremy dlya protsessov Markova)," *Izdat Kiev University.*

——— (1965): "Studies in the theory of random processes," *Translated from the Russian by Scripta Technica, Inc. Addison-Wesley Publishing Co.*

Taylor, C. (1995): "Digging for Golden Carrots: An Analysis of Research Tournaments," *American Economic Review*, 85, 872–890.

Thompson, F. B. (1952): "Equivalence of Games in Extensive Form," RM 759, RAND Corporation.

Tullock, G. (1980): *Efficient rent seeking, in: J.Buchanan, R.Tollison and G.Tullock, (eds.), Toward a Theory of Rent-Seeking Society*, Texas A&M University Press.

Wichardt, P. (2009): "Existence of Valuation Equilibria when Strategies Cannot Differentiate Between Equal Ties," Working Paper.