Five 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
Philipp Strack
aus Bonn

Bonn 2013

Dekan:	Prof. Dr. Klaus Sandmann
Erstreferent:	Prof. Dr. Paul Heidhues
Zweitreferent:	Prof. Dr. Benny Moldovanu

Tag der mündlichen Prüfung: 29.05.2013

Acknowledgements

In preparing this thesis, I received support from many people to whom I am grateful. First of all, I wish to express my sincere gratitude to my supervisor Paul Heidhues for his enduring support. 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, for his extraordinary support, guidance and many valuable comments. For comments and many helpful discussions I owe many thanks to Stefan Ankirchner, Dirk Bergemann, Sven Rady and Tymon Tatur.

I am particularly indebted to my friend and coauthor Christian Seel. Our time in the office was very fruitful, inspiring and let to the first three chapters of this dissertation. I am very grateful for the chance to collaborate and learn from the co-authors, Sebastian Ebert and Matthias Lang.

For material support, I want to thank the Bonn Graduate School of Economics, the Foundation of the German Economy, the DFG (SFB/TR15) and the Haussdorff-Center.

I want to thank Urs Schweizer and Silke Kinzig for their endless efforts in organizing and running the Graduate School.

Finally, I am greatly indebted to my girlfriend Linda, my family and my friends.

Contents

Introduction

This thesis consists of five chapters. The first three chapters form an entity and deal with stochastic contest models. The fourth chapter analyzes a problem of strategic experimentation with private payoffs. In the final chapter the behavior of agents with prospect theory preferences in optimal stopping problems is analyzed.

Contests and tournaments appear in many real world situations, like sports, politics, patent races, relative reward schemes in firms, or (public) procurement. In a contest multiple agents compete for a fixed prize. The first two chapters of this dissertation analyze contest models in continuous time. Both are based upon joint work with Christian Seel.

In the first two chapters we introduce a new type of contest model. In our model each agent decides when to stop a privately observed Brownian motion with drift. The processes of the different players are uncorrelated. The player who stops with the highest value wins a fixed prize. Each agent observes only his own progress, but not the progress or the stopping decision of the other players. As no deviations are observable we use Nash equilibrium as the solution concept. The model is similar to an all-pay auction with complete information as the payoffs of an agent depend only on the distribution of the stopped values of the Brownian motion of the other players. Hence the question which distributions can be implemented by stopping a Brownian motion naturally arises. This question is known in the probability theory literature as the Skorokhod embedding problem. Using insights from this literature we are able to construct an equilibrium and show uniqueness of the equilibrium outcome. Using a bound on stopping times derived in Ankirchner and Strack (2011) we are able to show in the second chapter that the equilibrium we construct is also the unique equilibrium if the contest ends at a fixed point in time – given this deadline is not to short. This result is important as many real world contests have a fixed deadline. From a technical point of view we provide a method to show that equilibria in an infinite horizon game are also equilibria in the finite horizon game given the time horizon is long enough. The next paragraph describes the applications of the model.

The model of the first chapter can be used to analyze the competition between fund managers. Each fund manager decides when to sell a risky asset. The performance of the asset follows a Brownian motion with drift. Assets are uncorrelated among fund managers and might go bankrupt. The fund manager who makes the highest profit by selling his asset wins a fixed prize. We show that the model has unique Nash equilibrium outcome. In this equilibrium the fund managers hold the risky asset even if it yields losses in expectation, i.e. the drift of the Brownian motion is negative. We show that the expected losses incurred in equilibrium by the fund managers are non-monotone in the expected return of the risky asset. Losses are highest if the underlying assets yields only moderate losses. In the last step we analyze the asymmetric situation where two fund managers hold assets with different expected return (drift of the Brownian motion) and riskiness (variance of the Brownian motion). We prove that in the unique Nash equilibrium outcome the weaker fund manager makes up for his disadvantage by using more risky strategies.

The second chapter focusses on Research and Development Tournaments. Each agent decides for how long to do research and pay the associated cost. The outcome of research is uncertain and described by a Brownian motion with drift. The research progress of different agents is uncorrelated. When all agents stop, the one who made the most progress, wins a prize. All agent have to pay the cost of their research. We assume that there exists a finite deadline at which all agents are forced to stop doing research.

We prove that the game has a unique Nash equilibrium outcome. If costs per time are constant, and if the riskiness of the research process (variance of the Brownian motion) converges to zero the equilibrium outcome converges to the symmetric equilibrium outcome of the symmetric all-pay auction. As the symmetric all-pay auction has multiple equilibria, this result provides an equilibrium selection criterion in favor of the symmetric equilibrium.

If the variance of the Brownian motion is strictly greater zero, i.e. the outcome of research is uncertain the predictions of our model are different from those of the all-pay auction. In equilibrium each agent makes positive profits, while in the all-pay auction the agents make zero profits. Furthermore we show that the profits of the agents are increasing in their costs and variance and decreasing in the drift. Hence the agents prefer to be in a situation where research is very costly, inefficient, and risky. This has possibly implications for many real world situations as the agents have incentives to make the designer choose inefficient research targets.

The third chapter of this dissertation is joint work with Matthias Lang and Christian Seel. It provides a micro-foundation of the discrete-bid all-pay auction using a continuous time contest model. More precisely we analyze the model where each agent decides when to stop a privately observed Poisson process. The processes of the agents are uncorrelated and the stopping decision of other players is not observable. As long as an agent did not stop she pays constant flow cost. There exists a finite deadline where each agent is forced to stop. We show that if the deadline is long enough the distribution over final outcomes in any Nash equilibrium equals the distribution of bids in a discrete-bid all-pay auction.

Chapter four is joint work with Paul Heidhues and Sven Rady and deals with a problem of strategic experimentation. Precisely, an optimal stopping game with an infinite time horizon is analyzed in which multiple players face an identical two-armed bandit problem. At all points in time, agents choose between the deterministic payoff of zero and a risky payoff whose distribution depends on the state of the world. In the bad state of the world the risky payoff is always negative, representing the cost of experimentation. If the world is in the good state, with a small probability an experiment is successful and yields a high payoff. Because a success can only happen in the good state of the world, it fully reveals the state of the world.

Similar models analyzed in literature show a free-rider problem, which means that less than the socially efficient amount of information is acquired in all Markov perfect equilibria. This free-rider problem is a consequence of the ability of agents to observe each others payoffs, which results in information being a public good.

We show that when payoffs are public information even in subgame-perfect equilibria, the ensuing free-rider problem is so severe that the number of experiments is at most one plus the number of experiments that a single agent would perform.

When payoffs are private information and players can communicate via cheap talk, the social optimal symmetric experimentation profile can be supported as a perfect Bayesian equilibrium for sufficiently optimistic prior beliefs.

Chapter five deals with optimal stopping under prospect theory preferences and is based upon joint work with Sebastian Ebert. While expected utility theory is the leading normative theory of decision making under risk, cumulative prospect theory is a popular descriptive theory. Expected utility theory is well-studied in both static and dynamic settings, ranging from game theory over investment problems to institutional economics. In contrast, for cumulative prospect theory most research so far has focused on the static case.

In this chapter, we investigate cumulative prospect theory's predictions in the dynamic context and point out fundamental properties of cumulative prospect theory. We show that already a small amount of probability weighting has strong implications for the application of prospect theory in the dynamic context.

More precisely we analyze the behavior of a naive agent with cumulative prospect theory preferences in continuous time optimal stopping problems. We prove that naive agent will never stop a non-degenerate diffusion process that represents his wealth. This holds for a very large class of processes, and independently of the reference point and the curvatures of the value and weighting functions.

This dynamic result is a consequence of a static result that we call skewness preference in the small: At any wealth level there exists an arbitrarily small right-skewed gamble that a prospect theory agent wants to take. By choosing a proper stopping strategy the agent can always implement such a gamble and thus never stops. We illustrate the implications for dynamic decision problems such as irreversible investment, casino gambling, and the disposition effect.

CHAPTER 1

Gambling in Contests

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 is 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 μ, constant variance σ, and initial value 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 positive amount or goes bankrupt, since she can ensure she wins an arbitrarily small positive 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 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.

The formal analysis proceeds as follows. Proposition 3.1 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 between stopping and continuing at any point of her support at any point in time.

[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 resulting equilibrium distributions are equivalent to the no-observability case.

For the two-player case, Proposition 3.2 derives the equilibrium stopping time that induces $F(x)$ explicitly. The equilibrium strategy mixes over strategies that stop if the process leaves a fixed interval. Proposition 3.5 extends Proposition 3.1 and 3.2 to a two-player game with asymmetric starting values.

For more than two players, Proposition 3.3 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 the conditions under which 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 3.3, 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 4.1 provides the main characterization result: the general shape of the expected value of the stopped processes is quasi-convex, falling, then rising in the drift μ and in the 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 the 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 success 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—into 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.

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 (2012). 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 the present chapter, players do not only possess private information about their stopping decision, but also about the realization of their stochastic process.

Finally, the 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 2 introduces the model. Section 3 derives the unique equilibrium distribution. In Section 4, we state the main characterization result, Proposition 5, and discuss its implications. Section 5 concludes.

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 $t \in \mathrm{R}_+$, agent i privately observes the realization of a stochastic process $X^i = (X^i_t)_{t \in \mathrm{R}_+}$ with

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

The constant $x_0 > 0$ denotes the starting value of all processes. The drift $\mu \in \mathrm{R}$ is the common expected change of each process X^i_t per time, i.e., $\mathrm{E}(X^i_{t+\Delta} - X^i_t) = \mu\Delta$. The noise term is an n-dimensional Brownian motion (B_t) scaled by $\sigma \in \mathrm{R}_+$. In Section 3.3, we allow for heterogeneity in all parameters and derive an equilibrium for the two-player case.

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., $\mathrm{E}(\tau^i) < \infty$. Second, a player has to stop in case of bankruptcy. More formally, we require $\tau^i \leq \inf\{t \in \mathrm{R}_+ : X^i_t = 0\}$ 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 $\tau^i(r^i)$.

2.2. Payoffs

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

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 $k = |\{i \in N : X^i_{\tau^i} = \max_{j\in N} X^j_{\tau^j}\}|$. Hence, the game is a constant sum game. All agents maximize their expected payoff, i.e., the probability of winning the contest. This optimization is independent of their risk attitude.

2.2. 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.

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

3. Equilibrium Analysis

In this section, we first derive the unique candidate equilibrium distribution. Second, we prove equilibrium existence—this is not trivial as the game has discontinuous payoffs and infinite strategy spaces. Our proof shows that there exists a stopping time inducing the candidate equilibrium distribution. We close the section with an extension to asymmetric parameters.

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

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. The next lemma proves otherwise.

LEMMA 3.1 (Continuity). *In equilibrium, for all i, $F^i(x)$ is continuous. At least one player i has no mass at zero, i.e., $F^i(0) = 0$.*

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; a similar argument shows that not all players can have a mass point at zero. However, some agents can have a mass point at zero: a slightly higher x does not entail a raise of one's win chance by a boundedly positive probability if at least one player places no mass at zero.

Lemma 3.1 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 $u_i(x) : \mathrm{R}_+ \to [0,1]$. As there are no mass points away from zero, we can express $u_i(x)$ in terms of the other agents' distributions, i.e., for all $x > 0$,

$$u_i(x) = \mathrm{P}(x > \max_{j\neq i} X^j_{\tau_j}) + \frac{1}{k}\underbrace{\mathrm{P}(x = \max_{j\neq i} X^j_{\tau_j}\})}_{=0}$$

$$(3.1) \qquad = \prod_{j\neq i} \mathrm{P}(X^j_{\tau_j} \le x) = \prod_{j\neq i} F^j(x)\,.$$

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

A well-known result about Brownian motion with drift μ and initial value x_0 facilitates the following derivation. Denote by $\tau_{(a,b)} = \inf\{t : X_t \notin (a,b)\}$ the first leaving time of the set (a,b). The probability of hitting the upper barrier b for Brownian motion with two absorbing barriers is

$$(3.2) \qquad \rho(a,b,x) = \mathrm{P}(X_{\tau_{(a,b)}} = b | x_0 = x) = \frac{\exp(\frac{-2\mu a}{\sigma^2}) - \exp(\frac{-2\mu x}{\sigma^2})}{\exp(\frac{-2\mu a}{\sigma^2}) - \exp(\frac{-2\mu b}{\sigma^2})}\,.$$

To state the next results, we need to introduce some additional notation. We define the support of player i by

$$\operatorname{supp} F^i = \{x : F^i(a) < F^i(b) \text{ for all } a < x < b\}\,.$$

Denote the left endpoint of the support of the distribution of player i by $\underline{x}^i = \inf\{x : F^i(x) > 0\}$, the right endpoint by $\overline{x}^i = \sup\{x : F^i(x) < 1\}$, and $\overline{x} = \max_i \overline{x}^i$. The following results establish necessary conditions for the distribution functions in equilibrium; the proofs are in the appendix.

Lemma 3.2 (Left Endpoint). *The support of each players' distribution starts at zero, i.e., $\underline{x}^i = 0$ for all $i \in N$.*

Lemma 3.3 (Symmetry). *The distributions are atomless and symmetric.*

Lemma 3.4 (Connected Support). *The distributions have a connected support.*

Henceforth, we use the symmetry property to suppress subscripts and superscripts i. As no player has a mass point at zero (Lemma 3.3), $u(x) = F(x)^{n-1}$ for all $x \in [0, \overline{x}]$. Therefore, $u(x)$, a product of continuous functions, is continuous. Since 0 and $\overline{x}$ are in the support of F, it is weakly better to continue for any point $x \in (0, \overline{x})$. By continuity, it is always optimal for a player to play the continuation strategy $\tau_{(0,\overline{x})}$. As the support is connected, it is also optimal to stop for any point $x \in (0, \overline{x})$. From formula (3.2), we obtain

$$u(x) = \rho(0, \overline{x}, x)1 + (1 - \rho(0, \overline{x}, x))0 = \rho(0, \overline{x}, x)$$

for any $x \in (0, \overline{x})$. Intuitively, this equation illustrates the trade-off between a higher stopping value (here $\overline{x}$) and a higher risk of bankruptcy.

By symmetry and optimality of stopping at x_0, we get

$$u(x_0) = \frac{1}{n} = \rho(0, \overline{x}, x)$$

which uniquely fixes $\overline{x}$. Note that $\overline{x} < \infty$ if and only if 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 $\frac{1}{n}$.

Inserting $\overline{x}$ and $\rho(0, \overline{x}, x)$ yields $u(x)$ as

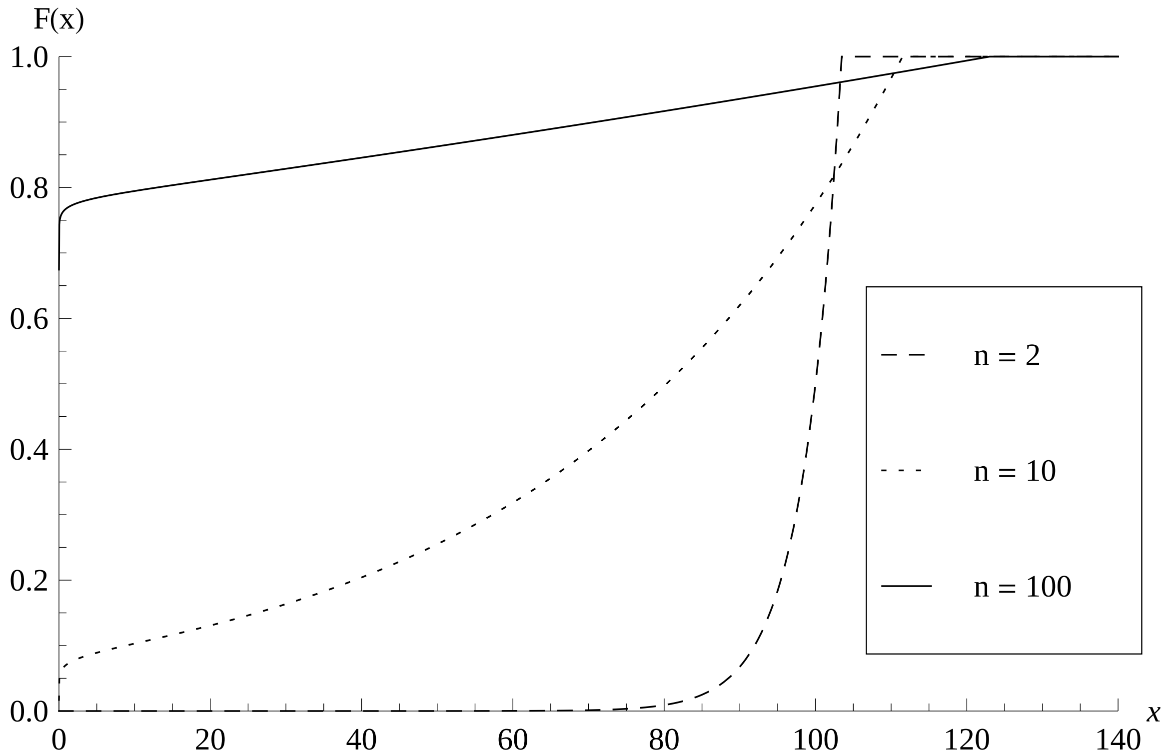

FIGURE 1. An example ($\mu = -0.1$, $x_0 = 100$, $\sigma = 1$) of the equilibrium cdf's for different sizes of players n.

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

Using symmetry, we get $u(x) = F(x)^{n-1}$. Hence, we characterize the unique candidate for an equilibrium distribution as follows (for an illustration, see Figure 1):

PROPOSITION 3.1. *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 is symmetric and induces the distribution F which has finite endpoints. To complete the proof, we need to show that no deviation gives a player a winning probability greater than $\frac{1}{n}$. Note that, by construction of the function $F(\cdot)$, the process $(u(X_t))_{t\in\mathbb{R}_+}$ 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), $\mathrm{E}(u(X_{\min\{\tau,n\}})) \leq u(X_0)$. As $u(X_t) \in [0,1]$ is bounded, we can apply the dominated convergence theorem to get

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

□

To complete the analysis, we scrutinize the special case in which X_t is a martingale, i.e., $\mu = 0$. 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 3.1, using the approximation $e^A = 1 + A + O(A^2)$ for small A.

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. In the next step, we explicitly derive such a strategy for the two-player case to convey the main intuition. The construction uses a mixture of deterministic threshold strategies to induce the final distribution. To formalize this intuition, we introduce the martingale transformation $\phi : \mathrm{R}_+ \to \mathrm{R}_+$, 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 $(\phi(X_t))_{t\in\mathrm{R}_+}$ is a martingale. In this case, $F(x) = \phi(x)/2$.

Proposition 3.2 (Equilibrium Strategy for Two Players). *If an agent randomly selects a number $\alpha \in (0, 1]$ from a uniform distribution and stops if*

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

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

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

$$\mathrm{P}(\phi(X_\tau) = 1 - \alpha) = \mathrm{P}(\phi(X_\tau) = 1 + \alpha) = \frac{1}{2}.$$

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

$$\mathrm{P}(X_\tau \leq x) = \mathrm{P}(\phi(X_\tau) \leq \phi(x)) = \frac{\phi(x)}{2} = F(x).$$

□

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 3.5 (Pedersen and Peskir, 2001, Theorem 2.1.). *Let (X_t) be a non-singular diffusion on* R *starting at zero, let $S(\cdot)$ denote its scale function satisfying $S(0) = 0$, and let ν be a probability measure on* R *satisfying $|S(x)|\nu(\mathrm{d}x) < \infty$. Set $m = \int_{\mathrm{R}} S(x)\nu(\mathrm{d}x)$. 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:*

itemsep=0pt $S(-\infty) = -\infty$ *and* $S(\infty) = \infty$ *;*
iitemsep=0pt $S(-\infty) = -\infty, S(\infty) < \infty$ *and* $m \geq 0$ *;*
iiitemsep=0pt $S(-\infty) > -\infty, S(\infty) = \infty$ *and* $m \leq 0$ *;*
ivtemsep=0pt $S(-\infty) > -\infty, S(\infty) < \infty$ *and* $m = 0$ *.*

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

Proposition 3.3 (Feasibility of the Equilibrium Distribution). *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

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

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

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

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

Proposition 3.1 and 3.3 combined yield F as the unique equilibrium distribution of the game.

3.3. An Extension to the Asymmetric Two-Player Case. In this section, we extend our equilibrium construction to the case of two players with asymmetric parameters x_0^i, μ_i, σ_i, $i \in \{1, 2\}$, i.e.,

$$X_t^i = x_0^i + \mu_i t + \sigma_i^2 B_t^i.$$

Similar to Assumption 1, we impose a sufficient condition for finiteness of both endpoints:

Assumption 2. *The drift of both players is negative, i.e.,* $\mu_i < 0$ *and* $\mu_j < 0$*.*

As in the symmetric case, we first show a lemma which is important to prove uniqueness.

Lemma 3.6. F_i *and* F_j *are continuous with* $F^i(0) = 0$ *for at least one player* i*. Moreover, they have the same connected support* $[0, \overline{x}]$*.*

Define the scale function $\phi_i(x) = \exp(-\frac{2\mu_i x}{\sigma_i^2})$.

Lemma 3.7 (Feasibility). *For both players $i \in \{1,2\}$ the equilibrium distribution satisfies* $\mathrm{E}(\phi_i(X^i_{\tau^i})) = \phi_i(X^i_0)$.

Proof. As $\mu_i < 0$, the argument in the proof of Lemma 8 implies the result. □

As the support of both players is connected, each player is indifferent between any stopping strategy on $(0, \overline{x}]$. We now construct the unique feasible distribution for which both players are indifferent.

$$F_i(x, \alpha_i, \overline{x}^i) = \alpha_i + (1-\alpha_i)\frac{\phi_j(x) - 1}{\phi_j(\overline{x}^i) - 1}\,.$$

By Lemma 3.6, at most one agent places a mass point at zero. As we show below, the other player determines the length of the support. Assume agent i places no mass at zero, i.e., $\alpha_i = 0$. The distribution F_i is feasible if it satisfies

$$\phi_i(x^i_0) = \int_0^{\overline{x}^i} f_i(x, 0, \overline{x}^i)\phi_i(x)\mathrm{d}x\,. \tag{3.3}$$

In the appendix, we show that there exists a unique $\overline{x}^i$ which satisfies equation (3.3).

Without loss of generality assume $\overline{x}^i \geq \overline{x}^j$. By definition, $F_i(\cdot, 0, \overline{x}^i)$ is a distribution and feasible for agent i. Agent j places a mass point of size $\alpha_j \geq 0$ at zero such that the distribution $F_j(\cdot, \alpha^j, \overline{x}^i)$ is feasible for agent j, i.e., α_j is the unique solution of the linear equation

$$\phi_j(x^j_0) = F_j(0, \alpha_j, \overline{x}^i)\phi_j(0) + \int_0^{\overline{x}^i} f_j(x, \alpha_j, \overline{x}^i)\phi_j(x)\mathrm{d}x \tag{3.4}$$

$$= \alpha_j + (1-\alpha_j)\int_0^{\overline{x}^i} f_j(x, 0, \overline{x}^i)\phi_j(x)\mathrm{d}x\,. \tag{3.5}$$

It follows from $\overline{x}^i \geq \overline{x}^j$ that $\int_0^{\overline{x}^i} f_j(x, 0, \overline{x}^i)\phi_j(x)\mathrm{d}x \geq \phi_j(x^j_0)$. As $\phi_j(x^j_0) > 1$ equation 3.4 always has a unique solution $\alpha_j \in [0,1]$.

Proposition 3.4. *The above construction derives the unique Nash equilibrium distributions.*

Proof. Existence: By construction, both distributions satisfy the feasibility condition from Proposition 3. Moreover, each player is indifferent between stopping and continuing on his support. Hence, the logic from Proposition 1 is applicable.

□

For heterogeneity in the starting values only—without loss of generality $x^1_0 > x^2_0$—this solution procedure yields the following result:

Proposition 3.5. *In equilibrium, the cdf of the first player is*

$$F^1(x) = \min\left\{1, \frac{1}{2}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x^1_0}{\sigma^2}) - 1}\right\}\,.$$

The cdf of the second player is

$$F^2(x) = \min\left\{1, \rho + (1-\rho)\frac{1}{2}\frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x^1_0}{\sigma^2}) - 1}\right\}\,.$$

In this example, player 2 first plays until $X_t^2 \notin (0, x_0^1)$; then he uses the same stopping strategy as player 1 if he reaches x_0^1. This induces the above cdf, where the constant ρ—probability of absorption in 0—fulfills

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

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 compensate his initial disadvantage.

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 3.1:

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

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:

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

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

Proposition 4.1. $\mathrm{E}(X_\tau)$ *is quasi-convex, falling, then rising in* μ. *If* $\mu < 0$, $\mathrm{E}(X_\tau)$ *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.

From an economic point of view, Proposition 4.1 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 avoid this drawback.

A Comparison to Related Models. In the static two-player contest of Lazear and Rosen (1981), contest success depends on the effort choice and the realization of a random variable. In their framework, contests are suitable to induce the optimal amount of effort. If, in our two-player model, agents had to

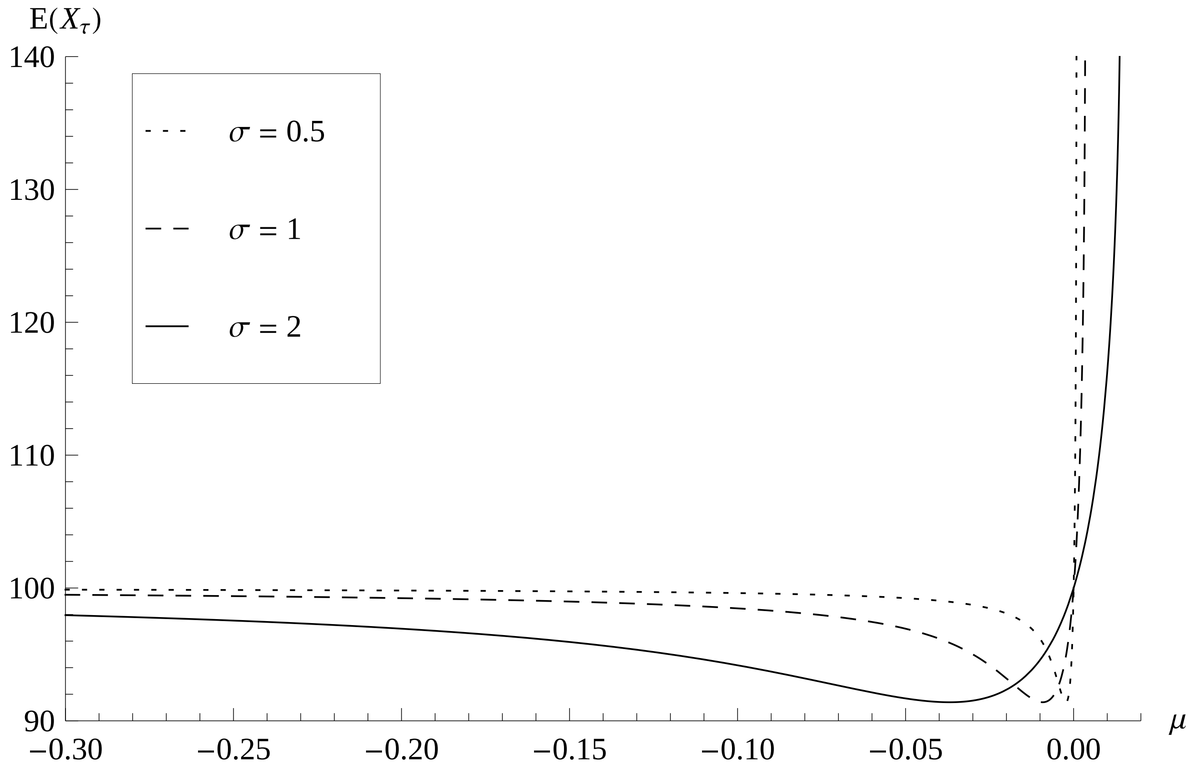

FIGURE 2. An example ($n = 2$, $x_0 = 100$) of the expected value of the stopped processes $\mathrm{E}(X_\tau)$ depending on the drift μ for different values of variance σ.

specify a fixed date at which they stop, they would stop immediately for negative values of the drift. Hence, to obtain our results, we need a dynamic decision problem for each player.

The equilibrium distributions in the present chapter are similar to those of all-pay auctions with complete information (e.g., Hillman and Samet, 1987, or Baye et al., 1996).[3] In both settings, the joint equilibrium distribution of the other players makes each player indifferent. The trade-off between a higher risk and a higher chance to win the prize thus serves as an implicit cost. In contrast to the all-pay auction, all players participate actively in the contest in any equilibrium.

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 4.1, 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.

[3]Complete information about valuations in the all-pay auction corresponds to complete information about starting values in this chapter.

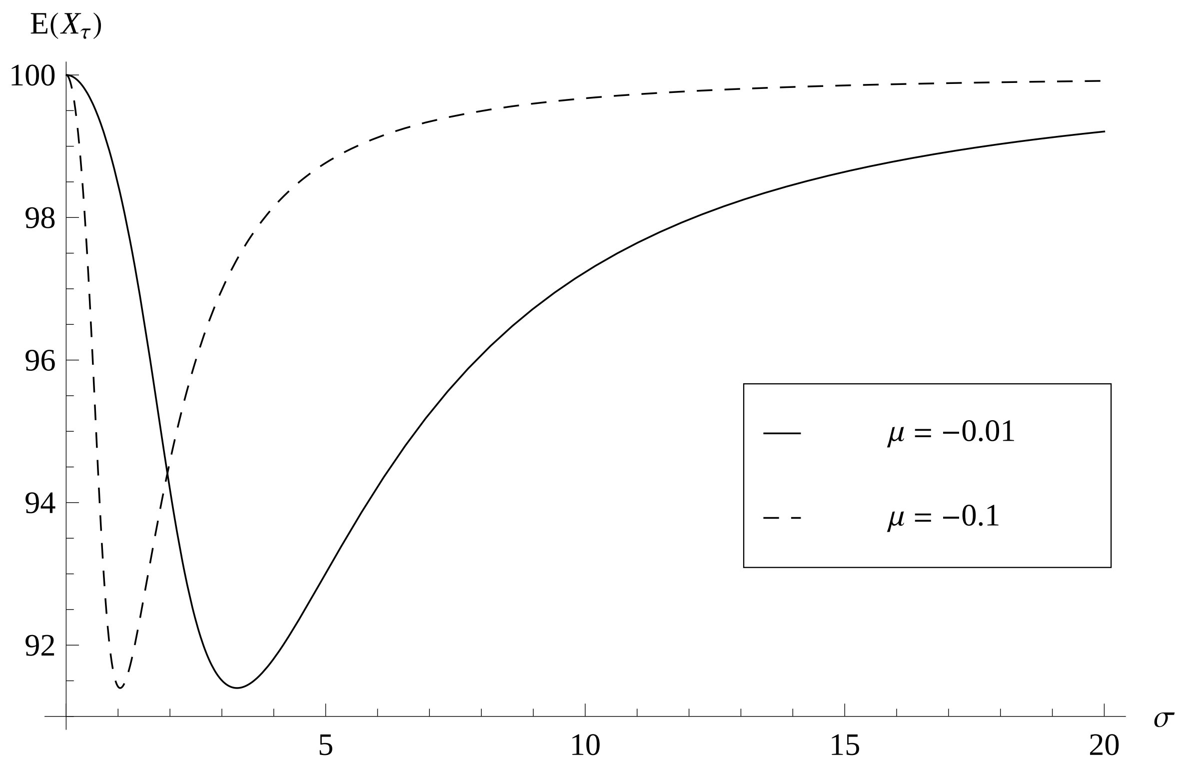

FIGURE 3. An example ($n = 2$, $x_0 = 100$) of the expected value of the stopped processes $\mathrm{E}(X_\tau)$ depending on the variance σ for different values of drift μ.

6. Appendix

We first show two auxiliary results which follow almost immediately from the definition of $u_i(x)$:

COROLLARY 6.1. *$F^i(x) \leq F^j(x)$ if and only if $u_j(x) \leq u_i(x)$ for all $x \in \mathrm{R}_+$.*

PROOF. By definition, $u_i(x)F^i(x) = u_j(x)F^j(x)$. Hence, $\frac{u_i(x)}{u_j(x)} = \frac{F^j(x)}{F^i(x)}$. □

COROLLARY 6.2. *$u_i(0) \geq u_j(0)$ implies that $u_j(\cdot)$ is continuous.*

PROOF. By Lemma 3.1, at least one agent places no mass point at zero, i.e., $u_k(0) > 0$ for at most one agent $k \in N$. As $u_i(0) \geq u_j(0)$, we obtain $u_j(0) = 0$. This implies $u_j(x) = \prod_{i \neq j} F^i(x)$ for all $x \in [0, \overline{x}]$. By Lemma 3.1, F^i is continuous for all i. As the product of continuous functions is continuous, u_j is continuous. □

The next result is useful to prove Lemma 3.2:

LEMMA 6.1. *If $a < b$ and $u_i(a) = u_i(b) < 1$, then $x \notin supp\, F^i$ for all $x \in (a, b]$.*

PROOF. Since u_i is weakly monotone increasing, $u_i(x) = u_i(a)$ for any point $x \in (a, b]$. At x, consider the continuation strategy $\tau_{(a,\overline{x})}$. From equation (3.2), we obtain

$$\mathrm{E}(u_i(X^i_{\tau_{(a,\overline{x})}}) \,|\, X_t = x) = \rho(a, \overline{x}, x)u_i(\overline{x}) + (1 - \rho(a, \overline{x}, x))u_i(a) > u_i(x)\,.$$

Hence, playing the continuation strategy $\tau_{(a,\overline{x})}$ is strictly better for player i than stopping at x. By continuity of u_i, the result remains true for an ϵ-ball around x and thus $x \notin \operatorname{supp} F^i$. □

Proof of Lemma 3.2: Assume there exists a player j such that $\underline{x}^j > 0$. Then $u_i(x) = 0$ for all $x \in [0, \underline{x}^j]$ and $i \neq j$. By Lemma 6.1, $x \notin \operatorname{supp} F^i$ for all $i \neq j$ for all $x \in (0, \underline{x}^j]$. Thus, $u_j(x) = \prod_{i \neq j} F^i(x)$ is constant on $(0, \underline{x}^j]$. By Lemma 6.1, $\underline{x}^j \notin \operatorname{supp} F^j$. Hence, there exists a $z > \underline{x}^j$ such that $F(z) = F(\underline{x}^j)$. As $\underline{x}^j > 0$, continuity of F^j implies that $F^j(\underline{x}^j) = 0 = F^j(z)$ which contradicts the definition of $\underline{x}^j$. □

We derive some additional results which we need to prove Lemma 6.3.

Lemma 6.2. *$u_i(0) \geq u_j(0)$ implies $F^i(x) \leq F^j(x)$ for all $x \in \operatorname{supp} F^i$.*

Proof. For $x = 0$ the result follows directly from Corollary 6.1. Let $x > 0$ and $x \in \operatorname{supp} F^i$. By continuity of u_i at $x > 0$, player i weakly prefers stopping at x to the continuation strategy $\tau_{(0,\overline{x}^j)}$

$$\rho(0, \overline{x}^j, x) u_i(\overline{x}^j) + (1 - \rho(0, \overline{x}^j, x)) u_i(0) \leq u_i(x) \,.$$

As 0 and $\overline{x}^j$ are in the support of player j, it is weakly optimal not to stop at any point $x \in (0, \overline{x}^j)$. By continuity of u_j (Corollary 6.2), the continuation strategy $\tau_{(0,\overline{x}^j)}$ is optimal for every point $x \in (0, \overline{x}^j)$, which yields

$$\rho(0, \overline{x}^j, x) u_j(\overline{x}^j) + (1 - \rho(0, \overline{x}^j, x)) u_j(0) \geq u_j(x) \,.$$

By assumption, $u_i(0) \geq u_j(0)$. Moreover, $F^i(\overline{x}^j) \leq 1 = F^j(\overline{x}^j)$ which by Corollary 6.1 yields $u_i(\overline{x}^j) \geq u_j(\overline{x}^j)$. Hence, the above equations imply $u_j(x) \leq u_i(x)$. By Corollary 6.1, $F^i(x) \leq F^j(x)$. □

Corollary 6.3. *$u_i(0) \geq u_j(0)$ implies $F^i(x) \leq F^j(x)$ for all x.*

Proof. For a given $x \in (0, \overline{x})$, define $y = \inf\{z : F^i(z) = F^i(x)\}$. As $F^i(w) < F^i(y)$ for all $w < y$ it follows that $y \in \operatorname{supp} F^i$. By Lemma 6.2, $F^i(y) \leq F^j(y)$. The monotonicity of F^j yields $F^j(y) \leq F^j(x)$. By definition, $F^i(x) = F^i(y)$. Thus, $F^i(x) = F^i(y) \leq F^j(y) \leq F^j(x)$. □

Define the martingale transformation $\phi(x)$ by $\phi(x) = \exp(\frac{-2\mu x}{\sigma^2})$.

Lemma 6.3. *In equilibrium, $\mathrm{E}(\phi(X^i_{\tau^i})) = 1$.*

Proof. In the first step, we show that $|\phi(X^i_s)|$ is bounded for all $s \leq \tau^i$.
Case 1: Let $\mu \leq 0$. We prove that $\overline{x}^i < \infty$. By the argument from Lemma 6.2, $\tau_{(0,\overline{x}^i)}$ is an optimal strategy for agent i. As $X_{\tau_{(0,\infty)}} = 0$ almost surely, we obtain $\overline{x}^i < \infty$. Therefore, $\phi(X_s) \in [0, \phi(\overline{x}^i)]$ for all $s \leq \tau^i$.
Case 2: For $\mu > 0$, we get $\phi(x) \leq 1$. Thus, $\phi(X^i_s) \in [0, 1]$ for all $s \leq \tau^i$.

We are now ready to prove the lemma. Note that $(\phi(X^i_s))_{s \in \mathrm{R}_+}$ is a martingale. Consider the sequence of bounded stopping times $\min\{\tau^i, n\}$ for $n \in \mathrm{N}$. By Doob's optional stopping theorem (Revuz and Yor, 2005, p.70), $\mathrm{E}(\phi(X_{\min\{\tau^i,n\}})) = \phi(X_0)$. As $\phi(X^i_t)$ is bounded, we can apply the dominated convergence theorem to get

$$\mathrm{E}(\phi(X^i_\tau)) = \mathrm{E}(\lim_{n\to\infty} \phi(X^i_{\min\{\tau^i,n\}})) = \lim_{n\to\infty} \mathrm{E}(\phi(X^i_{\min\{\tau^i,n\}})) = \phi(X^i_0) = 1 \,.$$

□

Corollary 6.4. *In equilibrium, F^i cannot first-order stochastically dominate F^j.*

Proof. If F^i stochastically dominates F^j, we get

$$\mathrm{E}(\phi(X^j_{\tau^j})) = \int_{\mathrm{R}_+} (1 - F^j(x))\phi'(x)\mathrm{d}x < \int_{\mathrm{R}_+} (1 - F^i(x))\phi'(x)\mathrm{d}x = \mathrm{E}(\phi(X^i_{\tau^i}))\,,$$

which violates Lemma 6.3.

□

Proof of Lemma 3.3: Without loss of generality, $u_i(0) \geq u_j(0)$. By Corollary 6.3, $F^i(x) \leq F^j(x)$ for all $x \in [0, \overline{x}]$. By Corollary 6.4, $F^i(x) = F^j(x)$. Symmetry implies that equilibrium distributions are atomless by Lemma 3.1. □

Proof of Lemma 3.4: Assume $\operatorname{supp} F^i$ is not connected. Then, by symmetry, $u_i(x)$ is constant for some interval $I \subset (0, \overline{x})$. Denote $y = \sup\{z : F^i(z) = F^i(x)\}$. By definition, $y \in \operatorname{supp} F^i$. On the other hand, $y \notin \operatorname{supp} F^i$ by Lemma 6.1, a contradiction. □

Proof of Lemma 3.6. The proofs of Lemma 1,2 and 6 do not rely on symmetry assumptions. Hence, F^i and F^j are continuous and their support starts at zero. It remains to prove connectedness and identical support length. Assume the support of an equilibrium distribution F^i is not connected. Then $u^j(x)$ is constant for some interval $I \subset (0, \overline{x})$. The rest of the argument for connectedness is identical to that in the proof of Lemma 4. The support length is identical, because at $\overline{x}^j$, it is the unique optimal strategy for player i to stop and vice versa. □

Lemma 6.4. *Equation 3.3 has a unique solution.*

Proof.

$$\begin{aligned}
\phi_i(x_0^i) &= \int_0^{\overline{x}^i} f_i(x, 0, \overline{x}^i)\phi_i(x)\mathrm{d}x \\
&= \int_0^{\overline{x}^i} \frac{\phi_j'(x)}{\phi_j(\overline{x}^i) - 1}\phi_i(x)\mathrm{d}x \\
\Rightarrow \phi_i(x_0^i)(\phi_j(\overline{x}^i) - 1) &= \int_0^{\overline{x}^i} \phi_j'(x)\phi_i(x)\mathrm{d}x \\
\Rightarrow 0 &= \int_0^{\overline{x}^i} \phi_j'(x)\phi_i(x) - \phi_i(x_0^i)\phi_j'(x)\mathrm{d}x \\
&= \int_0^{\overline{x}^i} \phi_j'(x)(\phi_i(x) - \phi_i(x_0^i))\mathrm{d}x
\end{aligned}$$

□

Formula for the Expected Value in the n-Player Case:
Let *Hyp* denote the Gauss hypergeometric function.

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

Proof of Proposition 4.1. We apply the monotone transformation $y = \exp(\frac{2\mu x_0}{\sigma^2})$ to $\mathrm{E}(X_\tau)$ to get

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

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

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

with the continuous extension $\frac{\partial^2 \mathrm{E}(X_\tau)/x_0}{\partial y^2} = \frac{1}{6}$ 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, $\mathrm{E}(X^i_{\tau^i})$ is quasi-convex in μ. As y is also monotone increasing (decreasing) in σ for $\mu < 0$ ($\mu > 0$), $\mathrm{E}(X^i_{\tau^i})$ is quasi-convex (quasi-concave) in σ if $\mu < 0$ ($\mu > 0$).

It remains to show that $\mathrm{E}(X_\tau)$ is first decreasing, then increasing. For $\mu \to -\infty$ and $\mu \to 0$, $\mathrm{E}(X_\tau) \to x_0$. 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, $\mathrm{E}(X_\tau)$ has to be first decreasing, then increasing.

□

CHAPTER 2

Continuous Time Contests

1. Introduction

Two types of informational assumptions are predominant in the literature on contests, races, and tournaments. Either there is no learning about the performance measure or standings throughout the competition at all or each player can observe the performance of all players 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).

However, there are many applications in which players are unable to observe their rivals, but can base their (effort) decision on their own progress. For instance, in a procurement contest, each participant is well-informed about his own progress, but lacks knowledge about research progress and behavior of his competitors. Another example is a competition for grants. Here, applicants can continuously choose how much time they invest in writing a proposal depending on their previous progress. However, they do not see the proposals of their rivals. Similarly, in a job promotion contest within a large firm, contestants can decide on their effort level. While they take their past success into account, they might not be able to observe effort and success of their competitors. In an appropriate model for these applications, contestants should be allowed to base their research decision over time only on their own progress. We propose such a model in the present paper.

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

Hence, in contrast to a *silent timing game* (Karlin, 1953; Park and Smith, 2008) in which the sole determinant of contest success is the time at which players stop, here, contest success and players' strategies additionally depend on their luck during the game. For $\sigma = 0$, our model reduces to a silent timing game with one prize for the player who stops latest. Thus, our model presents a first step towards a more general class of random timing games.

We characterize the unique Nash equilibrium outcome of the game without any restriction on the stopping time. Moreover, this result remains true for stopping strategies that are bounded by a real number T. Hence, the equilibrium construction is also valid for a contest with a sufficiently high, but fixed deadline. As the noise level approaches zero, the equilibrium converges to, and thus selects, the symmetric equilibrium of an all-pay contest. For two players and constant costs, each participant's profits increase if productivity (drift) of both players decreases, volatility increases, or costs increase. Hence, participants prefer a contest design which impedes progress.

The formal analysis proceeds as follows. Proposition 3.1 and Theorem 3.1 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 (e.g., Skorokhod, 1961, 1965; for a survey, see Obłój, 2004) to show that there exists a stopping strategy which induces this distribution.

The equilibrium distribution is unique for stopping strategies with finite expectation and bounded time stopping strategies, i.e., strategies that have to stop almost surely before a deadline $T < \infty$. As an immediate consequence of the latter result, our equilibrium construction remains valid in a contest with a sufficiently high deadline. This is economically important, since most applications have a fixed deadline—for example, a procurement contest for a jet fighter is usually issued with a deadline in a few years. From a technical point of view, the bounded time result is a major contribution of the present paper, since none of the previously mentioned literature could obtain a similar result.

In the next step, we 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 4.1. On the one hand, the model offers a microfoundation for the use of all-pay auctions to scrutinize environments containing little uncertainty; on the other hand, it gives an equilibrium selection criterion for the equilibria of the symmetric all-pay auction analyzed in Baye, Kovenock, and de Vries (1996a). Furthermore, this result serves as a benchmark to discuss how our predictions differ from all-pay models if volatility is strictly positive.

In most of the remaining analysis, we restrict attention to the case of two players and constant costs. This allows us to derive a tractable closed-form solution for the profits of each player. More precisely, by Proposition 4.4, these profits depend only on the ratio $\frac{2\mu^2}{c\sigma^2}$. In particular, profits increase as costs c increase, volatility σ^2 increases, or drift μ decreases (Theorem 4.2). Hence, contestants prefer worse technologies for both players.

This economically novel comparative statics result has an intuitive explanation: an increase in the productivity of each player makes patience a more important factor for contest success compared to chance. Hence, in equilibrium, expected stopping times and expected costs increase. Since the equilibrium is symmetric, the winning probability of each player remains constant. Summing up, we obtain lower expected profits.

The solution method we develop is applicable beyond the current setting. It provides a new approach to analyze more general variants of other timing games as the ones discussed in Park and Smith (2008), and, possibly, other stochastic games without observability. Moreover, as we show in an extension, the construction is not restricted to Brownian motion with drift, but can be applied to other stochastic processes.

Related Literature. In more recent literature, a few other papers model private information in related settings. Hopenhayn and Squintani (2011) scrutinize a preemption game in which private information arrives according to a Poisson process. They prove existence and derive a closed-form solution for a class of equilibria. In these equilibria, information disclosure occurs later than in the corresponding preemption game with public information.

In a companion paper (Seel and Strack, 2009), we show how relative performance pay might induce gambling behavior, i.e., investments in gambles with negative expectation. For this purpose, we analyze a contest model without any costs, but with a (usually negative) drift and a bankruptcy constraint. The driving forces of both models differ substantially. As in most of the contest literature, here, contestants trade off higher costs versus a higher winning probability, whereas in Seel and Strack (2009) the trade-off is between winning probability and risk.

Taylor (1995) also analyzes a contest model with private information. However, in his T-period model, only the highest draw in a single period determines the winner. In equilibrium, a player stops whenever she has a draw above a deterministic, time-independent threshold value.

We proceed as follows. Section 2 sets up the model. In Section 3, we prove that an equilibrium exists and has a unique distribution. Section 4 discusses the relation to all-pay contests and derives the main comparative statics results. Section 5 concludes. Most proofs are relegated to the appendix.

2. The Model

There are $n \geq 2$ agents indexed by $i \in \{1, 2, \ldots, n\} = N$ who face a stopping problem in continuous time. At each point in time $t \in \mathrm{R}_+$, agent i privately observes the realization of a stochastic process $(X_t^i)_{t \in \mathrm{R}_+}$ 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 $\mu \in \mathrm{R}_+$ is the common expected change of each process X_t^i per time, i.e., $\mathrm{E}(X_{t+\Delta}^i - X_t^i) = \mu \Delta$. The random terms σB_t^i are independent Brownian motions $(B_t^i)_{i \in N}$ scaled by $\sigma \in \mathrm{R}_+$.

2.1. Strategies. A pure strategy of player i is a stopping time τ^i. As each player only observes his own process, the decision whether to stop at time t can only depend on the past realizations of his process $(X_s^i)_{s \leq t}$.[1] 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.

Although the unique equilibrium outcome of this paper can be obtained in pure strategies, we incorporate mixing to make the results more general. To do so, we allow for random stopping decisions. More precisely, each agent i can choose an $(\mathcal{F}_t^i)$-adapted increasing process $(\kappa_t^i)_{t \in \mathrm{R}_+}, \kappa_t^i \in [0,1]$ such that, for every t, $\mathrm{P}(\tau^i \leq t \,|\, \mathcal{F}_t^i) = \kappa_t^i$. A pure strategy is equivalent to a process κ_t^i that equals zero for all $t < \tau$ and one otherwise.

In the remainder of the paper, we require stopping times to be bounded by a real number $T < \infty$ such that $\tau^i < T$ almost surely. This restriction on the strategy space makes it harder to derive the equilibrium—Proposition 3.2 shows that all results go through without it. We still think it is important to impose it, as most of the applications we have in mind, e.g., procurement contests and contests for grants have fixed deadlines. Hence, provided the deadline in the application is long enough—we give a condition in Lemma 3.10—the equilibrium we derive remains to hold for a contest with a deadline.

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 $c : \mathrm{R} \to \mathrm{R}_{++}$ which depend on the current value of the process X_t, but not on the time t. The payoff π^i is thus

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

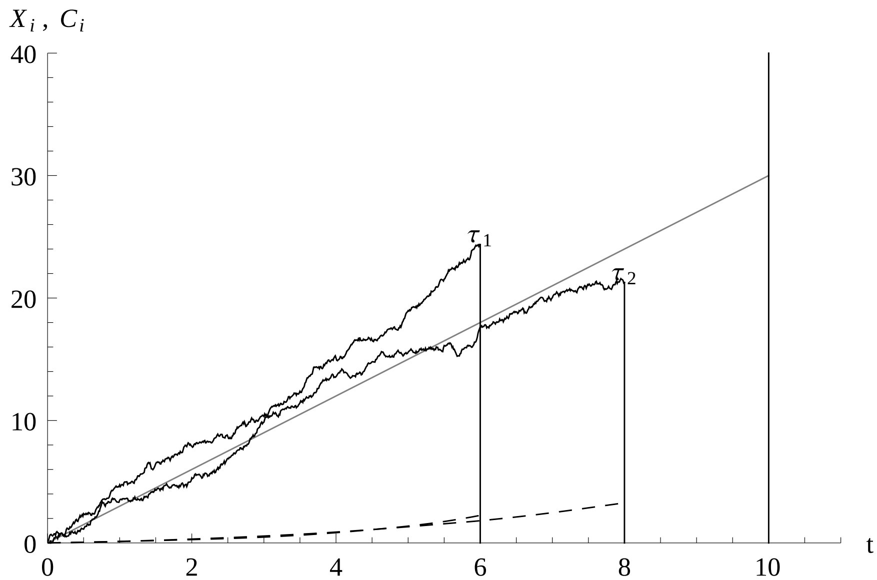

FIGURE 1. An example for the game with two agents $i \in \{1, 2\}$. Time is depicted on the x-axis, while the value X_t^i (solid line) and total costs (dotted line) are depicted on the y-axis. The parameters are $c(X_t^i) = \frac{1}{10}\int_0^t \exp(\frac{1}{10}X_s^i)\mathrm{d}t$, $\mu = 3$, $\sigma = 1$, $T = 10$.

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

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 $\mathrm{E}(\pi^i)$. 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 3. *For every $x \in \mathrm{R}$, the cost function $c : \mathrm{R} \to \mathrm{R}_{++}$ is continuous and bounded away from zero on $[x, \infty)$, i.e., there exists a $\underline{c} > 0$ such that, for all x, $c(x) \geq \underline{c}$.*

Figure 1 illustrates processes, stopping decisions, and corresponding costs in an example. Note that for the Brownian motion specification research progress need not be strictly increasing. This non-monotonicity allows us to capture organizational forgetting in the procurement example. The empirical literature shows that progress in firms is often non-monotone, since "organizations can forget the know-how gained through learning-by-doing due to labor turnover, periods of inactivity, and failure to institutionalize tacit knowledge" (Besanko, Doraszelski, Kryukov, and Satterthwaite, 2010, p.453); see, e.g., Benkard (2000) and Thompson (2001) for empirical evidence in the aircraft and shipbuilding industries. In the example of the grant competition, it also seems reasonable to include a small probability that the current quality of a proposal decreases, e.g., because a laptop is stolen or data are lost. Alternatively, an individual might just forget a detail which he had not yet written down; see Argote, Beckman, and Epple (1990) for related literature.

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) $F^i : \mathrm{R} \to [0,1]$ of his stopped process $F^i(x) = \mathrm{P}(X^i_{\tau^i} \leq x)$. 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 3.1. *At least two players stop with positive probability on every interval $I = (a,b) \subset [\underline{x}, \overline{x}]$.*

Lemma 3.2. *No player places a mass point in the interior of the state space, i.e., for all i, for all $x > \underline{x}$: $\mathrm{P}(X^i_{\tau^i} = x) = 0$. At least one player has no mass at the left endpoint, i.e., $F^i(\underline{x}) = 0$, for at least one player i.*

We omit the proof of Lemma 3.2, since it is simply 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 3.1.

Lemma 3.2 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

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

Lemma 3.3. *All players have the same right endpoint, $\overline{x}^i = \overline{x}$, for all i.*

Lemma 3.4. *All players have the same expected profit in equilibrium. Moreover, each player loses with certainty at $\underline{x}$, i.e., $u^i(\underline{x}) = 0$, for all i.*

Lemma 3.5. *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 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)

$$c(x) = \mu u'(x) + \frac{\sigma^2}{2} u''(x)\,. \tag{3.1}$$

As (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 3.4. We determine the other one in the following lemma:

Lemma 3.6. *In equilibrium, $u'(\underline{x}) = 0$.*

Imposing the two boundary conditions, the solution to equation (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, 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

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

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

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

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 3.7. *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,

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

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

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

By Assumption 1, $\underline{c} > 0$. Continuity of F implies that there exists a unique point $\overline{x} > \underline{x}$ such that $F(\overline{x}) = 1$. □

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

Lemma 3.8. *If $\tau \leq T < \infty$ is a bounded stopping time that induces the continuous distribution $F(\cdot)$, i.e., $F(z) = \mathrm{P}(X_\tau \leq z)$, then $1 = \int_{\underline{x}}^{\overline{x}} \phi(x)F'(x)\mathrm{d}x$.*

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

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

□

We use the necessary condition from Lemma 3.8 to prove that the equilibrium distribution is unique.

Proposition 3.1. *There exists a unique pair* $(\underline{x}, \overline{x}) \in \mathrm{R}^2$ *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. **Uniqueness:** As F is absolutely continuous, the right endpoint $\overline{x}$ satisfies $(a): 1 = \int_{\underline{x}}^{\overline{x}} F'(x; \underline{x}, \overline{x})\mathrm{d}x$. By Lemma 3.7, there exists a unique $\overline{x}$ for every $\underline{x}$ such that (a) is satisfied. Lemma 3.8 states that any feasible distribution satisfies $(b): 1 = \int_{\underline{x}}^{\overline{x}} F'(x; \underline{x})\phi(x)\mathrm{d}x$. We show that intersection of the set of solutions to equation (a) and to equation (b) consists of a single point. 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, the set of solutions to (a) satisfies

$$\frac{\partial \overline{x}}{\partial \underline{x}} = -\frac{-\overbrace{F'(\underline{x};\underline{x})}^{=0} + \int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x;\underline{x})\mathrm{d}x}{F'(\overline{x};\underline{x})} = -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x;\underline{x})\mathrm{d}x}{F'(\overline{x};\underline{x})}. \tag{3.2}$$

Applying the implicit function theorem to (b) gives us

$$\begin{aligned} \frac{\partial \overline{x}}{\partial \underline{x}} &= -\frac{-\overbrace{F'(\underline{x};\underline{x})}^{=0} + \int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x;\underline{x})\phi(x)\mathrm{d}x}{F'(\overline{x};\underline{x})\phi(\overline{x})} \\ &= -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x;\underline{x})\overbrace{\phi(x-\overline{x})}^{>1}\mathrm{d}x}{F'(\overline{x};\underline{x})} \\ &< -\frac{\int_{\underline{x}}^{\overline{x}} \frac{\partial}{\partial \underline{x}} F'(x;\underline{x})\mathrm{d}x}{F'(\overline{x};\underline{x})}. \end{aligned} \tag{3.3}$$

The last inequality follows from $\frac{\partial}{\partial \underline{x}} F'(x; \underline{x}) \geq 0$. As (3.2) > (3.3), the solution sets to equation (a) and (b) intersect at most once. Thus, in equilibrium, both the left and the right endpoint are unique.

Existence: We have shown in Lemma 3.7 that, for every $\underline{x}$, there exists a unique $\overline{x}$ such that $F(\cdot; \underline{x}, \overline{x})$ is a distribution function. Furthermore, rewriting the last inequality in the proof of Lemma 3.7, we get $\overline{x} \leq \underline{x} + \frac{\mu}{\underline{c}}\frac{\sigma^2}{2\mu}$ and thus $\underline{x} \to -\infty \Rightarrow \overline{x} \to -\infty$. Consider a left endpoint $\underline{x} \geq 0$ and a right endpoint $\overline{x}$ such that (a) is satisfied. Then, $\phi(x) < 1$ for all $x \in [\underline{x}, \overline{x}]$ and hence $\mathrm{E}(\phi(x)) < 1$ for $x \sim F(\cdot, \underline{x}, \overline{x})$. Now consider a left endpoint $\underline{x} < 0$ and a right endpoint $\overline{x} < 0$ such that (a) is satisfied. Then, $\phi(x) > 1$ for all $x \in [\underline{x}, \overline{x}]$ and hence $\mathrm{E}(\phi(x)) > 1$ for $x \sim F(\cdot, \underline{x}, \overline{x})$. Equation (b) is equivalent to $\mathrm{E}(\phi(x)) = 1$. By continuity and the intermediate value theorem, there exists an $\underline{x}, \overline{x}$ such (a) and (b) are satisfied. □

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

Lemma 3.9. *Every strategy that induces the unique distribution F from Proposition 3.1 is an equilibrium strategy.*

Proof. Define $\Psi(\cdot)$ as the unique solution to (3.1) with the boundary conditions $\Psi(\underline{x}) = 0$ and $\Psi'(\underline{x}) = 0$. By construction, the process $\Psi(X_t^i) - \int_0^t c(X_s^i)\mathrm{d}s$ 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

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

The last equality results, as each agent is indifferent to stop immediately with the expected payoff $u(X_0)$ or to play the equilibrium strategy with the expected payoff $\mathrm{E}(u(X_\tau))$. The inequality implies that any possible deviation strategy obtains a weakly lower payoff than the equilibrium candidate. □

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., $F(z) = \mathrm{P}(X_\tau \leq z)$. To show that the game has a Nash equilibrium, the existence of a bounded stopping time inducing $F(\cdot)$ remains to be established. 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 for a given $T \in \mathrm{R}_+$ the existence of a stopping time τ that is *bounded* by T and embeds a given distribution in Brownian motion with drift. They provide an analytical construction for such a stopping time (see p.221).[2]

In addition to the assumption stated in the next lemma, Ankirchner and Strack (2011) assume that the condition in Lemma 3.8 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 3.10 (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 lemma enables us to prove the main result of this section:

Theorem 3.1. *The game has a Nash equilibrium. In equilibrium, the strategy of each player induces the distribution F from Proposition 3.1.*

The proof in the appendix verifies Lipschitz continuity of the function g, which makes Lemma 3.10 applicable. Thus, a Nash equilibrium in bounded time stopping strategies exists and, by Proposition 3.1, the equilibrium distribution F is unique.

In the next proposition, we show that existence and uniqueness of the equilibrium distribution F continue to hold if we do not restrict agents to bounded stopping times. For this purpose, define the payoff if an agent never stops to be $-\infty$.

[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, Heyne, and Imkeller (2008).

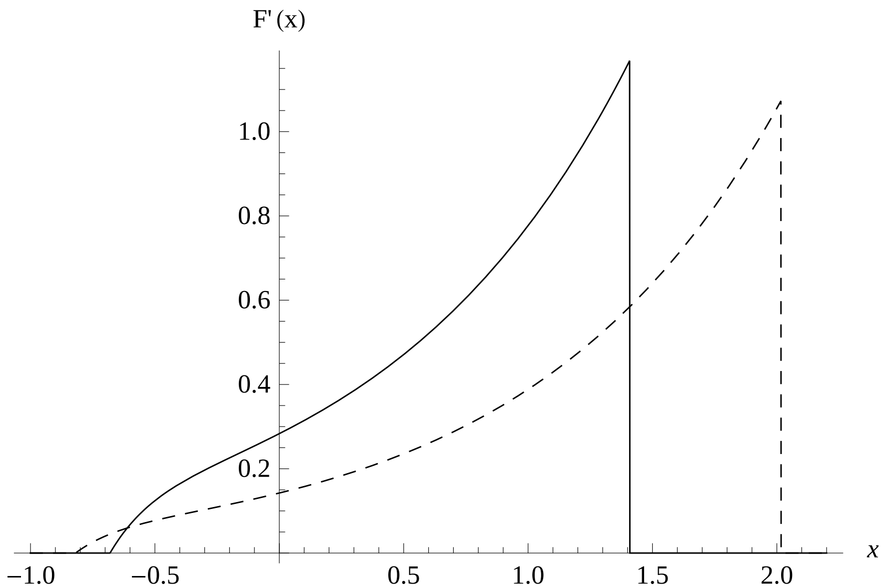

FIGURE 2. The density function $F'(\cdot)$ for the parameters $n = 2$, $\mu = 3$, $\sigma = 1$ and the cost-functions $c(x) = \exp(x)$ solid line and $c(x) = \frac{1}{2}\exp(x)$ dashed line.

PROPOSITION 3.2. *The distribution F is also the unique equilibrium distribution if agents are not restricted to bounded stopping times.*

In summary, we have proven existence and uniqueness of the equilibrium distribution for both finite and bounded stopping times. The uniqueness result differs from related models such as symmetric all-pay auctions or symmetric silent timing games, in which there are usually multiple equilibria. In some of these equilibria only a subset of players is active, i.e., submits a positive bid in the auction or does not stop directly in the silent timing game. Here, the variance leads all players to be active in the game. Hence, we obtain a unique equilibrium distribution.

REMARK 3.1. *While the equilibrium distribution is unique, the equilibrium strategy is not. In fact, there are uncountably many finite stopping times that solve the Skorokhod embedding problem; see Proposition 4.1 in Obłój (2004).*

A major conceptual innovation compared to the literature is the bounded time requirement $\tau < T$, which makes the equilibrium derivation applicable to contests with a fixed deadline. To see that this result is not trivial to obtain, note that 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 simple construction through a mixture over cutoff strategies of the form

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

cannot be used to implement F.

In tug-of-war models with full observability (Harris and Vickers, 1987; Moscarini and Smith, 2007; Gul and Pesendorfer, 2011), it is not possible to obtain the same equilibrium for finite time strategies and bounded time strategies. 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.

4. Equilibrium Analysis

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

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

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

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 $\int_{\underline{x}}^{0} F'(x)\phi(x)\mathrm{d}x \leq \int_{\underline{x}}^{\overline{x}} F'(x)\phi(x)\mathrm{d}x = 1$. □

Taking the limit $\sigma \to 0$, the 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 symmetric 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 4.1. *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 variance is negligible. Figure 3 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, Kovenock, and de Vries (1996a). This chapter offers an equilibrium selection criterion in favor of the symmetric equilibrium. 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 3.2.

4.2. Comparative Statics and Rent Dispersion. Theorem 4.1 has linked all-pay contests with complete information to our model for the case of vanishing noise. In the following, we scrutinize how

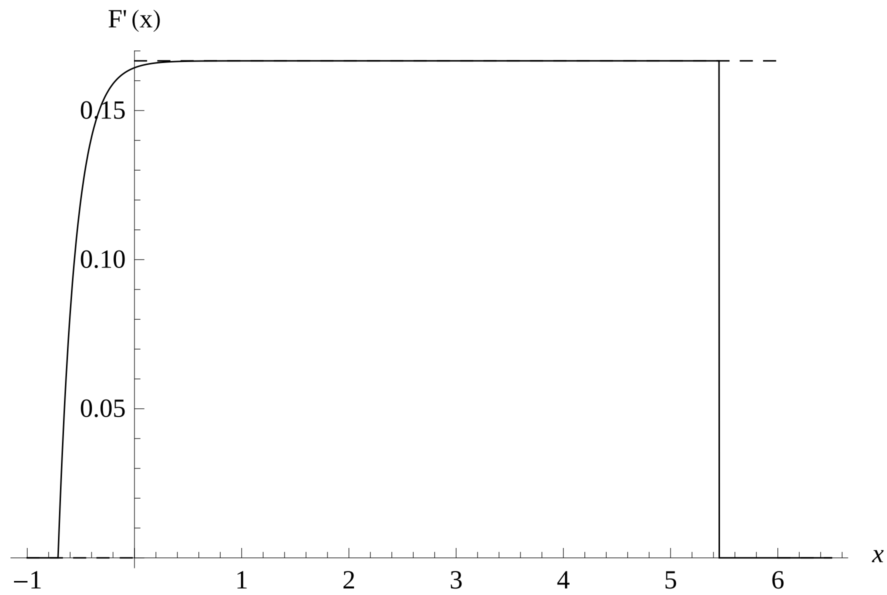

FIGURE 3. This figure shows the density function $F'(\cdot)$ with support $[-0.71, 5.45]$ for the parameters $n = 2$, $\mu = 3$, $\sigma = 1$ and the cost-functions $c(x) = \frac{1}{2}$ (solid line) and for the same parameters the equilibrium density of the all-pay auction with support $[0, 6]$ (dashed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

predictions differ for positive noise. In a symmetric all-pay contest 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 4.1. *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 $\underline{x} < 0$. □

Intuitively, agents generate informational rents through the private information about their 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.

We now derive comparative statics in the number of players for constant costs. Define the support length as $\Delta = \overline{x} - \underline{x}$.

LEMMA 4.2. *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 4.2. *If the number of agents n increases and $c(x) = c$, the expected profit of each agent decreases.*

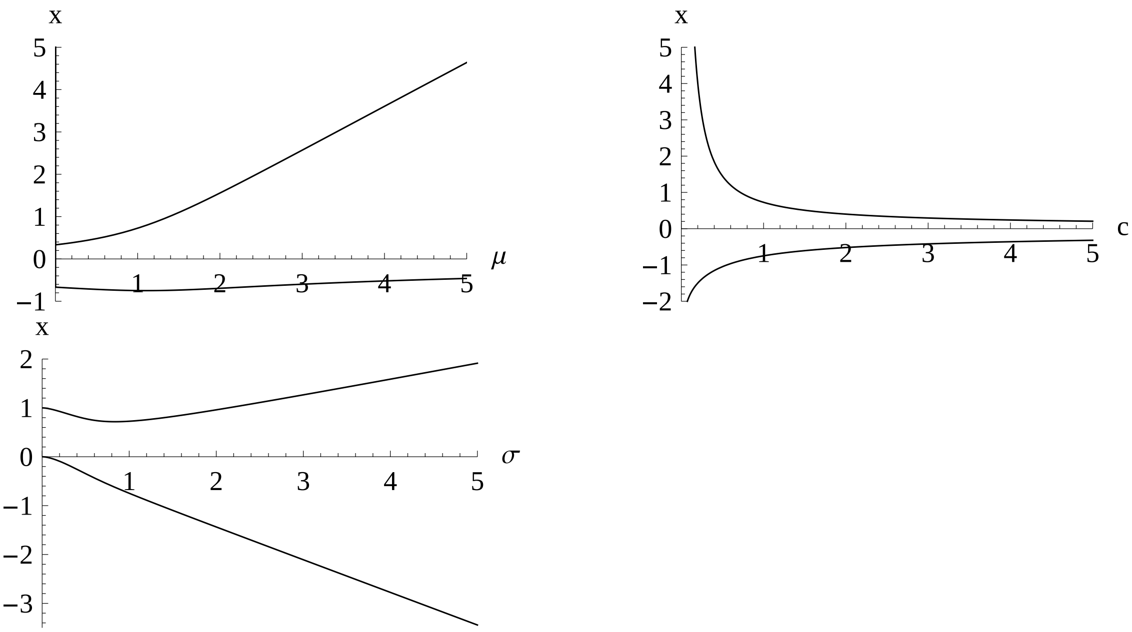

FIGURE 4. This figure shows the left endpoint $\underline{x}$ and the right endpoint $\overline{x}$ for $n = 2, \sigma = 1, \mu = 1$ and constant cost $c = 1$ varying the productivity μ in the first figure, the costs c in the second figure and the variance σ in the third.

PROOF. The function $u(x)$ depends only on $x - \underline{x}$. As n increases, $\underline{x}$ increases by Lemma 4.2. 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.

4.3. The Special Case of Two Players and Constant Costs. We henceforth 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 $W_0 : [-\frac{1}{e}, \infty) \to \mathrm{R}_+$ 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 $h : \mathrm{R}_+ \to [0, 1]$ by

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

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

PROPOSITION 4.3. *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 4. The next proposition derives a closed-form solution of the profits π of each player.

Proposition 4.4. *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 static result of this paper.

Theorem 4.2. *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 $\frac{\mu}{c}$ is a deterministic measure of productivity per time. On the other hand, the second term $\frac{2\mu}{\sigma^2}$ measures the impact of randomness (σ^2) on the final outcome. If both firms increase their productivity, patience becomes more critical to winning than chance. Hence, the expected stopping time in equilibrium increases. As each firms' winning probability remains $\frac{1}{2}$ by symmetry, profits decrease. In summary, participants prefer to have worse—more costly, more random, or less productive—technologies.

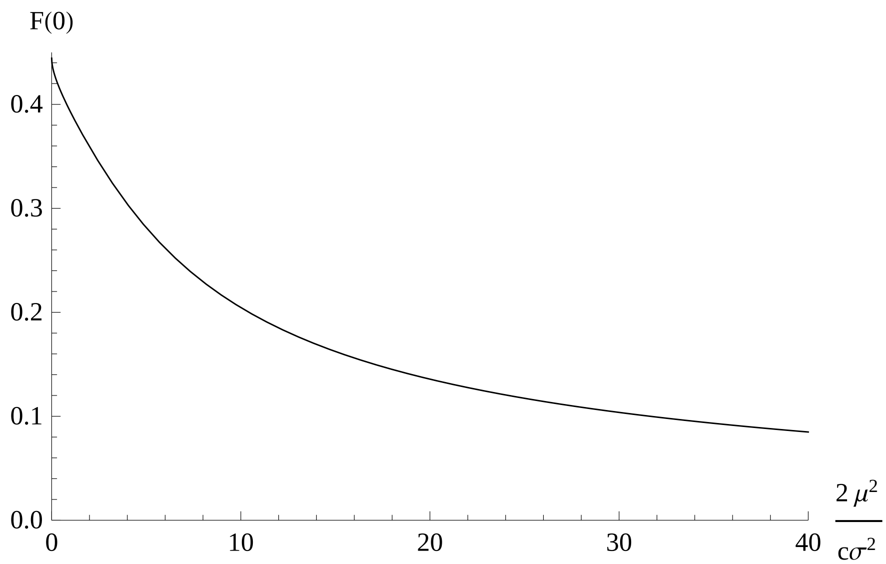

Figure 5. This figure shows the equilibrium profit $F(0)$ of the agents on the y-Axes for $n = 2$ constant cost-functions $c(x) = c \in \mathrm{R}_+$ with $y = \frac{2\mu^2}{c\sigma^2}$ on the x-Axes.

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

Proposition 4.5. *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 4.2. Hence, profits are bounded from above by $\lim_{y\to 0} u(0)$. By l'Hôpital's rule,

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

□

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

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

4.4. Extension to Other Stochastic Processes. Up to now, we have derived our results for Brownian motion with drift. However, it is straightforward to extend them to any process Z_t such that Z_t is the result of a strictly monotone transformation $\rho : \mathrm{R} \to \mathrm{R}$ of Brownian motion with drift X_t, i.e., $Z_t = \rho(X_t)$. First, observe that the success process Z_t is only payoff relevant as an ordinal variable, as agent i receives a prize if

$$Z^i_{\tau^i} > \max_{j\neq i} Z^j_{\tau^j} \Leftrightarrow \rho^{-1}(Z^i_{\tau^i}) > \max_{j\neq i} \rho^{-1}(Z^j_{\tau^j}) \Leftrightarrow X^i_{\tau^i} > \max_{j\neq i} X^j_{\tau^j}\,.$$

As costs $c(\cdot)$ depend on the value of the process Z_t, we define the cost function $\tilde{c}(x) = c(\rho(x))$ for the process X_t. Let $\tilde{F}$ be the equilibrium distribution for the game with a Brownian motion with drift X_t and cost $\tilde{c}$, i.e., $X^i_{\tau^i} \sim \tilde{F}$. The equilibrium distribution of the original game is now given by $F(x) = \tilde{F}(\rho^{-1}(x))$, i.e., $Z^i_{\tau^i} \sim F$. For example, this generalization includes geometric Brownian motion with drift ($Z_t = \exp(X_t)$), which is relevant in many financial applications.

5. Conclusion

In this paper, we have introduced a model of contests in continuous time. Our informational assumptions, which differ from most of the contest literature, seem appropriate to model applications such as procurement contests and contests for grants. Under mild assumptions on the cost function, a Nash equilibrium outcome exists and is unique for both bounded and finite time stopping strategies. If the research progress contains little uncertainty, the equilibrium is close to the outcome of the symmetric equilibrium of a static all-pay auction. Thus, our model provides an equilibrium selection result for the symmetric all-pay auction. If the research outcome is uncertain, each player prefers higher research costs, worse technologies, and higher uncertainty for all players. This economically novel feature results from chance becoming more important for contest success than patience.

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 lower bound on the time horizon that ensures the existence of such equilibria. These methodological contributions are economically relevant, since most real world applications have a fixed deadline. They should prove useful to future research on random timing games and other stochastic games without observability.

6. Appendix

Proof of Lemma 3.1. 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 $\overline{x}^i > \overline{x}^{-i} + \epsilon$. □

Remark 6.1. *We write* $\tau^i_{(a,b)}(x)$ *shorthand for the continuation strategy* $\inf\{t : X^i_t \notin (a,b) | X^i_s = x\}$ *in the next three proofs. Clearly,* $\tau^i_{(a,b)}(x)$ *is not a bounded time continuation strategy, but we use it to bound the payoffs. Moreover, 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 3.3. Assume $\overline{x}^j > \overline{x}^i$. For at least two players j, j', the payoff from continuation with $\tau^j_{(\underline{x}^j,\overline{x}^j)}(\overline{x}^i)$ is weakly higher than from stopping at $X^j_t = \overline{x}^i$ by Lemma 3.1. By Lemma 3.2, 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 non-negative continuation payoffs,

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

On the other hand, optimality of stopping at $\overline{x}^i$ for player i implies

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

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

Proof of Lemma 3.4. 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 $\pi^i \leq \mathrm{P}(X^i_{\tau^i} = \overline{x} | \tau^i_{(\underline{x}^i,\overline{x})}(0)) - \mathrm{E}(c(\tau^i_{(\underline{x}^i,\overline{x})}(0)))$. If player j deviates to the strategy $\min\{\tau^j_{(\underline{x}^i,\overline{x})}(0), T\}$, he 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 3.1, there exists another player j who also stops in the interval. In particular, for $x \in I$, for any $\epsilon > 0$, we get

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

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

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 3.5. 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 continuation 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, for all i, j. This implies $F^i(x) = F$, for all i.

□

Proof of Lemma 3.6. 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 $u : \mathrm{R} \to \mathrm{R}_+$, let $\Psi : \mathrm{R} \to \mathrm{R}$ be the unique function that satisfies the second order ordinary differential equation $c(x) = \mu\Psi'(x) + \frac{\sigma^2}{2}\Psi''(x)$ with boundary conditions $\Psi(\underline{x}) = \partial_+ u(\underline{x})$ and $\Psi'(\underline{x}) = \partial_+ u(\underline{x})$. As $\Psi'(\underline{x}) > 0$, there exists a point $\hat{x} < \underline{x}$ such that $\Psi(\hat{x}) < 0 = u(\hat{x})$. Consider the strategy S that stops when either the point $\hat{x}$ or $\overline{x}$ is reached or at 1,

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

As $u(\hat{x}) > \Psi(\hat{x})$, it follows that $\mathrm{E}(u(X_S)) > \mathrm{E}(\Psi(X_S))$. Thus,

$$\mathrm{E}(u(X_S) - \int_0^S c(X_t^i)\mathrm{d}t) > \mathrm{E}(\Psi(X_S) - \int_0^S c(X_t^i)\mathrm{d}t) .$$

Note that, by Itô's lemma, the process $\Psi(X_t^i) - \int_0^t c(X_s^i)\mathrm{d}s$ 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.,

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

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. □

Proof of Theorem 3.1. 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 $n = 2$ 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

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

□

Proof of Proposition 3.2. The existence proof proceeds in three steps. First, we show that any optimal strategy has finite expectation. In the second step, we prove that the support of any equilibrium distribution induced by a finite stopping time is bounded. These conditions allow us to apply Doob's theorem in Step 3. Thus, Lemma 3.8 and 3.9, which guarantee equilibrium existence, extend to arbitrary stopping times.

Step 1: As c is bounded from below $c(\cdot) \geq \underline{c}$, the payoff of agent i if he does not stop before time t is bounded from above by

$$\mathbf{1}_{\{X^i_{\tau^i}\geq \max_{j\neq i} X^j_{\tau^j}\}} - \int_0^t c(X^i_s)\mathrm{d}s \leq 1 - t\underline{c}\,.$$

Furthermore,

$$\mathrm{E}(\mathbf{1}_{\{X^i_{\tau^i}\geq \max_{j\neq i} X^j_{\tau^j}\}} - \int_0^{\tau^i} c(X^i_s)\mathrm{d}s) \leq 1 - \mathrm{E}(\tau^i)\underline{c}\,.$$

Thus, to satisfy $1-\mathrm{E}(\tau^i)\underline{c} > 0$, any optimal strategy must have finite expectation.

Step 2: In the next step, we show that the support of any equilibrium distribution induced by a finite stopping time is bounded. Note that $\overline{x} \neq \infty$, since for positive drift ($\mu > 0$), the expected stopping time would be infinite otherwise. Hence, it remains to show that $\underline{x} \neq -\infty$. As the optimal strategy is Markovian, for any point x in the support, it is optimal to continue until $X^i_\tau \in \{\underline{x}, \overline{x}\}$. Suppose the left endpoint of the equilibrium distribution of an agent equals minus infinity, i.e. $\underline{x} = -\infty$. The expected stopping time for the strategy that stops only if the right endpoint is reached equals

$$\mathrm{E}(\tau_{(-\infty,\overline{x})}(x) - t) = \frac{\overline{x} - x}{\mu}\,.$$

As $\tau_{(-\infty,\overline{x})}$ is an optimal continuation strategy, $\underline{c}\mathrm{E}(\tau) \leq 1$. Hence,

$$1 \geq \underline{c}\mathrm{E}(\tau) = \frac{\underline{c}}{\mu}(\overline{x} - x) \Leftrightarrow x \geq \overline{x} - \frac{\mu}{\underline{c}}\,.$$

Consequently, $\tau_{(\infty,\overline{x})}$ is not an optimal continuation strategy for $x < \overline{x} - \frac{\mu}{\underline{c}}$. This contradicts the assumption $\underline{x} = -\infty$.

Step 3: As the support of the equilibrium distribution is bounded, for any equilibrium stopping time τ, the process $X_{\min\{t,\tau\}}$ is uniformly integrable. Thus, the conditions for Doob's optional sampling theorem ($\mathrm{P}(\tau < \infty) = 1$, $\mathrm{E}[|X_\tau|] < \infty$, and uniform integrability) hold. Hence, Lemma 3.8 and 3.9 follow by the same argument as in the text.

The proof for uniqueness of the equilibrium distribution for finite expected stopping times works similar as in bounded time case, with slight modifications in the proofs of Lemma 3.1, 3.2, and 3.5. The details are available upon request.

□

Proof of Proposition 4.3 and 4.4. 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 $W_0 : [-\frac{1}{e}, \infty) \to \mathrm{R}_+$ 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 3.8 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 $g:\mathrm{R}_+ \to [0,1]$ 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 4.3.

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 4.4. □

Proof of Theorem 4.2. By Proposition 4.4, it suffices to show that the profit $F(0)$ is increasing in y. Consider the following 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

Poisson Contests

1. Introduction

There is a large amount of literature on all-pay auctions that are often motivated as reduced-form models of a dynamic problem in which contestants do not observe each other's decisions.[1] Applications include contests – the bid in the auction is a proxy for the effort each participant incurs in the contest; think, for example, of an R&D, job promotion or lobbying contest. The aim of this chapter is to explore formally the relationship between all-pay auctions and such contests. In particular, we consider a stochastic contest in continuous time and show that it is equivalent in an appropriate sense to the all-pay auctions previously studied. This justifies the implicitly assumed equivalence in the previous literature.

In a seminal paper, Taylor (1995) analyzes a stochastic T-period contest, in which each player can take a costly draw from an i.i.d. distribution function in any period. As we do below, he assumes that players cannot observe actions or outcomes of their rivals throughout the contest. At time T, only the player with the highest overall draw receives a prize. In equilibrium, players stop whenever their draw is above a constant threshold value and the induced probability distributions over final values do not resemble those of all-pay auctions. In contrast to Taylor's research function that depends only on the highest draw, we model research in a cumulative way, i.e., the probability of reaching a certain value in the next period depends on the previous contest success.

As an example illustrating the different approaches, consider a university department that wants to hire a new researcher. Under Taylor's description of the game, the applicant with the single best publication would be hired, while the criterion of this chapter chooses the applicant who has more (major) publications. 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 prototype of a fighter jet or tank. Here, the government usually awards the prize depending on a number of attributes.

Motivated by the aforementioned applications, we consider a model in which successes, e.g., major publications, arrive randomly over time as long as a player exerts costly effort. Formally, each player decides when to stop a privately observed Poisson process and bears costs, which increase in the realization of her stopping time. At time T, the player with the most successes wins a prize; ties are broken randomly.

If the contest lasts infinitely long, it is possible to replicate any distribution of the all-pay auction by a corresponding stopping strategy. Hence, the choice of a stopping time in the infinite contest is equivalent to the choice of a distribution. If the designer adjusts the prize to the Poisson arrival rate, this is a pure

[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.

coalescing of moves that has no effect on the set of Nash equilibria.[2] An equivalence result for an infinite contest, however, does not seem sufficient as real-world contests end within bounded time.

The main question that this chapter addresses is whether an equilibrium equivalence holds even if the game ends at a fixed date $T < \infty$. Here, coalescing of moves does not establish the equivalence as the contest and the all-pay auction have different normal forms. Hence, we analyze if there is a strategy that stops almost surely before time T and replicates the equilibrium distribution of the all-pay auction. This is not trivial, as, for example, a distribution that puts all probability mass on one success cannot be replicated in bounded time.

Nevertheless, if T is large enough, any symmetric Nash equilibrium distribution of the discrete all-pay auction can be implemented in the contest. Moreover, we provide time bounds suggesting that the critical level T for the equilibrium distributions to be identical is moderate. Hence, despite the stochastic research outcomes in the contest equilibrium distributions coincide, which offers a microfoundation for the all-pay auction.

The contest resembles a silent timing game - as first explored by Karlin (1953), and, most recently by González-Díaz, Borm, and Norde (2007) and Park and Smith (2008) - because actions of the other players are unobservable. Yet, in this literature, the sole determinant of success is the amount of time players stay in the game; in this chapter research success depends on a privately observed stochastic process and players decide accordingly.

We proceed as follows. Section 2 discusses the discrete all-pay auction. In Section 3, we describe the stochastic contest and derive the main equivalence results. Section 4 concludes. We relegate most proofs to the appendix.

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 $x^i \in \mathrm{N}_0$, entailing costs cx^i which he has to pay independently of the contest outcome. We henceforth normalize $c = 1$. The agent with the highest bid wins a prize p. Ties are broken randomly. Hence, the utility of player i is

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

A mixed strategy of each player is a probability measure $f^i \colon \mathrm{N}_0 \mapsto [0, 1]$ which specifies the probability of bidding x^i. Denote the associated cumulative distribution function by $F^i(z) = \mathrm{P}(x^i \leq z) = \sum_{y=0}^{z} f^i(y)$. We henceforth impose the following assumption:

Assumption 4. *If $n = 2$, then $\frac{p}{2} \notin \mathrm{N}$.*

This assumption rules out non-generic parameter settings for which infinitely many equilibria exist; for an analysis of this case, see Example 2 in Cohen and Sela (2007). In the next lemma, we derive a global indifference property for the set of equilibrium strategies.

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

Lemma 2.1. *In every symmetric Nash equilibrium, all players are indifferent between the pure strategies* $0, 1, \dots, \tilde{x}$ *with* $\tilde{x} = \min_{x \in \mathrm{N}_0} F(x) = 1$. *Moreover, any symmetric equilibrium distribution has a connected support including* 0.

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

Proposition 2.1. *Assume* $n = 2$. *The unique Nash equilibrium of the all-pay auction is*

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

where $\tilde{x} = \max\{x \in \mathrm{N}_0 \colon x \text{ even and } \frac{x}{p} < 1\}$.

Differing from a previous analysis of the two-player case in Baye, Kovenock and de Vries (1994), we get a unique equilibrium. The different result comes from Assumption 4, which rules the parameter setting of Baye, Kovenock and de Vries (1994) in which the prize has to be an integer.

The probability measure $f(x)$ alternates due to the tie-breaking rule. If a specific bidder always receives the prize in case of a tie, all equilibrium distributions are uniform on $\{1, \dots, \tilde{x} - 1\}$ with some mass on 0 and $\tilde{x}$. For the appropriate choices of the prize – $p = y$ for any odd integer y – the above expression becomes a uniform distribution. Furthermore, as the uniform distribution is also the unique Nash equilibrium in the continuous bid case, it can be seen as the most natural equilibrium of the two-player all-pay auction.

3. The Stochastic Contest

In the contest, players can do research until time T. At every point in time $t \leq T \in \mathrm{R}_+$, each player $i \in N$ privately observes his research results given by a time-homogeneous Poisson process X_t^i with intensity $\lambda \in \mathrm{R}_+$ and jump size 1. A strategy of player i is a stopping time τ^i with respect to the natural filtration $\mathcal{F}_t^i$ generated by the process X_t^i. Hence, in accordance with the contest literature, each player's strategy is independent of the stopping decision and realization of any other player. The strategy, however, might depend on the previous path of realizations of her own process.[4]

Stopping at time t entails costs of $\bar{c}t$.[5] The player who stops his process at the highest value wins a prize $\bar{p}$ at time T. Thus, each player's profit is

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

Formally, the number of bidders in a tie is $m = |\{j \in N | X^j_{\tau^j} = \max_{i \in N} X^i_{\tau^i}\}|$. Denote the probability distribution, induced by the stopping time τ^i, by $F^i \colon \mathrm{N}_0 \to [0, 1]$, where $F^i(x) := \mathrm{P}(X^i_{\tau^i} \leq x)$. To determine

[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.

[4]If the stopping decision is reversible, the set of equilibrium distributions remains the same; no agent has an incentive to delay research as no new information arrives over time.

[5]The main results of this chapter are also valid in a stochastic control problem, in which players can influence the value of the Poisson arrival rate λ subject to a linear cost function $c(\lambda)$. A possible interpretation is a firm hiring more employees.

his optimal strategy, each player cares only about the final distributions of the other players, as no new information about their successes arrives over time. Hence, we use Nash equilibrium as the solution concept.

Define the payoff process $(\Pi^i_t)_{t\in\mathrm{R}_+}$ of player i as his expected payoff π^i of stopping immediately

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

Hence, it is the price times a weighted probability of winning it alone or sharing it given the stopping time minus the expected cost. In a next step, we consider whether for a given distribution F, a process can be stopped such that the distribution over X_τ equals F.

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

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

Lemma 3.1. *For any distribution F with a connected support including 0, there exists a time bound $T' \in \mathrm{R}_+$ such that for all $T \ge T'$, F is feasible.*

In particular, it is possible to implement all symmetric equilibria of the all-pay auction if the contest lasts long enough. The following proposition establishes the main result.

Proposition 3.1. *Let F be a symmetric equilibrium distribution of the all-pay auction, T' as defined in Lemma 3.1 and $\frac{\bar{c}}{\bar{p}} = \frac{1}{p}\lambda$. Then, for all $T \ge T'$, F is also a Nash equilibrium distribution of the stochastic contest.*

Proof. By Lemma 3.1, there exists a strategy in the stochastic contest that induces the Nash equilibrium distribution of the all-pay auction. To show that this is indeed an equilibrium distribution for the stochastic contest, we verify that no player has a profitable deviation. Define $f^{-i}(x) = \mathrm{P}(\max_{j\neq i} X^j_{\tau^j} = x)$ and

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

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

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

This formula can be interpreted as follows: λ is the probability of a jump per time; $f^{-i}(X_t)(1-\mathrm{E}(\frac{1}{m^i(X_t)}))+ f^{-i}(X_t+1)\mathrm{E}(\frac{1}{m^i(X_t+1)})$ is the increase in winning probability given a jump occurs and $\bar{c}$ is the cost per time. The formula for the infinitesimal generator omits multiple jumps, because they occur with probability 0. As the indifference property proven in Lemma 2.1 is mathematically equivalent to

$$0 = (f^{-i}(X_t)(1-\mathrm{E}(\frac{1}{m^i(X_t)})) + f^{-i}(X_t+1)\mathrm{E}(\frac{1}{m^i(X_t+1)}))p - 1,$$

it follows that $\frac{\mathrm{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t} = 0$ for all $X_t < \tilde{x}$ with $\tilde{x} = \max\{x \in \mathrm{N}_0\colon x \text{ even and } \frac{x}{p} < 1\}$. By optimality of a bid $\tilde{x}$ compared to $\tilde{x}+1$ in the all-pay auction, we have $\frac{\mathrm{E}(\mathrm{d}\Pi^i_t)}{\mathrm{d}t} \le 0$ at $\tilde{x}$. For all points $x > \tilde{x}$ the winning probability does not increase after a success and consequently $\frac{\mathrm{E}(\mathrm{d}\Pi_t^i)}{\mathrm{d}t} = -\bar{c} < 0$. Hence, $\frac{\mathrm{E}(\mathrm{d}\Pi_t^i)}{\mathrm{d}t} \leq 0$ for all x and the process Π_t^i is a supermartingale. Denote by τ^i the stopping time inducing the equilibrium distribution F and by τ_0^i the strategy to stop immediately. As every possible strategy $\hat{\tau}^i$ of player i has to stop before $T < \infty$, we can apply Doob's optional stopping theorem (Rogers and Williams, 2000, p.189) to obtain $\mathrm{E}(\Pi_{\hat{\tau}^i}^i) \leq \mathrm{E}(\Pi_{\tau_0^i}^i)\,\forall \hat{\tau}^i$. Lemma 2.1 implies indifference and $\mathrm{E}(\Pi_{\tau_0^i}^i) = \mathrm{E}(\Pi_{\tau^i}^i)$. Thus, there is no profitable deviation. □

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. Players use the stochasticity to replicate the equilibrium distribution of the all-pay auction. As a major difference, players in the dynamic 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 3.2. *Assume $\frac{\bar{c}}{\bar{p}} = \frac{1}{p}\lambda$. There exists a time bound $T'' \in \mathrm{R}_+$ 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. It is possible to calculate the minimal implementation time numerically for any distribution $F(x)$ with the approach given in the proof of Lemma 3.1 in the appendix. However, as there is no closed-form solution, we estimate 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 3.3. *The uniform distribution on $\{0, 1, \ldots, \tilde{x}\}$ is feasible, if the contest lasts at least $T = \frac{\tilde{x}\ln(\tilde{x}+1)}{\lambda}$ periods.*

With the construction in the proof it is possible to give a time bound for other equilibrium distributions. For more than two players, the local martingale property implies that in any symmetric equilibrium, players stop earlier than in the two-player case, as they need to win against more competitors. Hence, the two-player time bound is also valid for $n > 2$ players.

4. Conclusion

In this chapter, we have provided a microfoundation for using static all-pay auction models to scrutinize contests in which the decision problem of each player is dynamic and stochastic. If the contest lasts long enough − a moderate contest length is sufficient − and has a cumulative structure, the 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.

5. Appendix

Proof of Lemma 2.1. By contradiction, assume

$$\exists\, x \in \mathrm{N}_0 : x < \max supp\{f\},$$

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 $x' := \min\{\tilde{x} > x\colon$ player i is indifferent between $\tilde{x}$ and any $\hat{x} \in supp\{f\}\}$ as the lowest bid above x for which the indifference holds. By definition $x' - 1$ is not played and thus $f^i(x'-1) = 0$ for all players implies that $\mathrm{P}(\max_{j\neq i} x^j = x' - 1) = 0$. Yet by optimality of x' compared to $x' - 1$ in equilibrium we get

$$\mathrm{P}(\max_{j\neq i} x^j = x')\mathrm{E}\frac{1}{m^i(x')} > \frac{p}{} \tag{5.1}$$

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

$$\left(1 - \mathrm{E}\frac{1}{m^i(x')}\right)\mathrm{P}(\max_{j\neq i} x^j = x') + \mathrm{P}(\max_{j\neq i} x^j = x' + 1)\mathrm{E}\frac{1}{m^i(x'+1)} \leq \frac{p}{.} \tag{5.2}$$

Due to $\mathrm{E}(\frac{1}{m^i(x')}) \leq \frac{1}{2}$, the left hand side of this equation is at least as high as $\mathrm{P}(\max_{j\neq i} x^j = x')\mathrm{E}(\frac{1}{m^i(x')})$. Thus, (5.1) and (5.2) yield a contradiction.

If $n > 2$, then $\mathrm{E}(\frac{1}{m^i(x')}) < \frac{1}{2}$ for all points in $\{0, 1, \ldots, \tilde{x}\}$ for which $f(x) > 0$. It follows that the support is connected, as the contradiction is valid even if (5.1) holds only with equality. For two players, we establish the connectedness result in the proof of Proposition 2.1. □

Proof of Proposition 2.1. In any equilibrium in which both players are indifferent between all strategies $\{0, 1, \ldots, \tilde{x}\}$, we have

$$\frac{f^i(x)}{2} + \frac{f^i(x+1)}{2} = \frac{1}{p} \;\; \forall x < \tilde{x} - 1, \; \forall i \text{ and}$$

$$\sum_{0}^{\tilde{x}} f^i(x) = 1. \tag{5.3}$$

The unique solution to this system of equations is given in Proposition 2.1, if $\frac{p}{2} \notin \mathrm{N}$. Hence, by the indifference argument in Lemma 2.1, the equilibrium is the unique symmetric equilibrium. This also implies connectedness of the support, the missing step in the proof of Lemma 2.1.

In the next lines, we show that the game does not have an asymmetric equilibrium. First, we prove that $f(x+2) > 0 \implies f(x+1) > 0$ or $f(x) > 0$ for all x. Assume to the contrary

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

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 to $x+2$ otherwise. Yet, in this case player j prefers, at least weakly, the point $x+2$ compared to x, which means

$$2\frac{1}{p} \leq \frac{f^i(x+2)}{2}. \tag{5.4}$$

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 (5.4). Thus, player j strictly prefers to bid $x+3$ compared to $x+2$. This yields the required contradiction to optimality of bidding $x+2$ and proves that in equilibrium for any two subsequent points smaller than $\tilde{x}$ there is at most one point without mass.

In the next step, we show that players are indifferent on $\{0,\dots,\tilde{x}\}$. Assume to the contrary that player j finds it strictly worse to play at point $x<\tilde{x}$. 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

$$\frac{f^j(x-1)+f^j(x+1)}{2}\geq\frac{2}{p}. \tag{5.5}$$

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

$$\frac{f^j(x)+f^j(x+1)}{2}\geq\frac{1}{p}\quad\text{and}\quad\frac{f^j(x+1)+f^j(x+2)}{2}\leq\frac{1}{p}. \tag{5.6}$$

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 (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 (5.3) as $\frac{p}{2}\notin\mathrm{N}$ by Assumption 4.

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 $f^i(0)=f^j(0)=0$, it has to be that $f(1)\leq\frac{2}{p}$ to make it weakly optimal to bid 1 compared to 2. 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 $f^j(x)>0$. In this case, however, her expected profits are zero and her expected profit 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 (5.3).

Thus, in any equilibrium, players have to be indifferent on $\{0,\dots\tilde{x}\}$. In addition, the probability measure condition (5.3) must hold. By the first part of the proof, this determines the equilibrium distributions uniquely. □

Proof of Lemma 3.1. Rost (1976) constructs a stopping time τ for general right continuous Markov processes $(X_t)_{t\in\mathrm{R}_+}$ that minimizes the residual expectation $\mathrm{E}(\int_{\tau^t}^{\tau}\mathrm{d}s)$ for all $t\in\mathrm{R}_+$ 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 $g\colon\mathrm{N}_0\to\mathrm{R}$, it holds

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

In the following, we show that this condition is fulfilled

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

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

$$\tau = \inf\{t \in \mathrm{R}_+ | (t, X_t) \in A\} \, . \tag{5.7}$$

We can equivalently formulate (5.7) as

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

with some function $H\colon \mathrm{N}_0 \to \mathrm{R}_+$. 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 $H^\star = \min_{0 \leq x \leq \tilde{x}} H(x)$. This stopping time embeds F in a minimal time $H^\star$ with $\tau \leq H^\star$ almost surely. □

Proof of Proposition 3.2. By Proposition 3.1 any 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 equilibria in the stochastic contest. To be feasible, an equilibrium distribution in the stochastic contest requires a connected support including 0 and a lower bound for the probability measure $f(x) \geq \epsilon(x, T, F) \; \forall x \in supp\{x\}$.

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 playing 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 about continuing at $\tilde{x}+1$, contradict each other for $n > 2$,

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

As $f^i(\tilde{x})$ approaches zero as T increases and $\mathrm{E}(\frac{1}{m^i(x)}) < 1 - \mathrm{E}(\frac{1}{m^i(x)})$ 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 ϵ. We first show that the following inequalities have to hold

$$\frac{2}{p} - \tilde{x}\epsilon \leq \sum_{k=0}^{1} f^i(x+k) \leq \frac{2}{p} + \epsilon \; \forall x \in \{0, 1, \ldots, \tilde{x}-1\}.$$

The mass at $\tilde{x}+1$ can at most be $\frac{2}{p}$ in equilibrium, as the player weakly prefers to stop there. Consequently, the expected gain of continuing at $\tilde{x}$ is at most ϵp – the expected gain of continuation until the next success, if the contest length is infinite. Hence, to make continuation at least weakly optimal at $\tilde{x}-1$ for some $t < T$,

$f^i(\tilde{x}-1) \in [\frac{2}{p} - 2\epsilon, \frac{2}{p} - \epsilon]$. But then, the mass at point $\tilde{x}-2$ can be at most 2ϵ, as the player would always continue to play there otherwise. This argument extends to the whole support.

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

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

$\hat{\epsilon}$ goes to zero as T increases and by Assumption 4 $\frac{p}{2} \neq \mathrm{N}$. Consequently, there exists a time T, such that $\sum_{i=0}^{\tilde{x}} f^i(k) \neq 1$ which contradicts the probability measure condition (5.3). □

Proof of Proposition 3.3. We construct a strategy which waits at some points in time, calculate its implementation time and then argue that it is slower than the optimal strategy.

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

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

CHAPTER 4

Strategic Experimentation with Private Payoffs

1. Introduction

In many real-life situations economic agents face a tradeoff between exploring new options and exploiting their knowledge about which option is likely to be best. A stylized model capturing this feature is a two-arm bandit problem in which a gambler repeatedly decides which of two different slot machines to play with the ultimate goal of maximizing his monetary reward. The consecutive payoffs of the arms are identical and independently distributed random variables whose underlying distribution is unknown. When playing a given arm, the agent learns more about its underlying distribution—knowledge that is useful for future choices. Starting with Rothschild (1974), variants of the multi-arm bandit problem have been applied to a wide variety of economic settings (see Bergemann and Välimäki 2008 for a recent overview).

Often, however, it is natural to presume that an agent can learn not only from her own exploration but also from the experiences of others. Experimental consumption is a case in point. As a stylized example, suppose there are two agents with common tastes. The agents choose between two items on a given menu and both know the quality of one regular item that has been on the menu for a long time. There is also a new item of unknown quality on the menu whose quality critically depends on how well it is cooked. The restaurant's chef is believed to be either good or very good. A very good chef is able to sometimes prepare the item nearly perfectly, while a good chef is only able to prepare a dish of average quality. Each agent can now learn through three channels: he may experiment himself and try the new item, he may learn from observing what the other agent chooses, and finally he may ask the other agent about whether the item was well-prepared whenever she tried it.

While we are mainly interested in the problem of strategic experimentation with private payoffs, we begin by analyzing a benchmark case in which agents' payoffs are publicly observable. We consider a game in which two agents face identical two-armed bandit problems with a common underlying state. The scenario where each player can observe all other players' actions and signals (without any communication) is well studied in the literature (Bolton and Harris 1999, 2000, Keller, Rady and Cripps 2005, Keller and Rady 2010). As a consequence of payoffs being observable, all players share a common belief about the state of the world. The literature focuses on continuous-time setups and shows that, if the agents can only condition on this common belief, i.e. use Markov perfect strategies, it is impossible to realize the socially optimal outcome. Furthermore, if a single success is fully revealing as in our model, then in any Markow perfect equilibrium agents stop experimenting once the common belief reaches the single-agent cutoff. We introduce a discrete-time setup and considerably strengthen these existing results by illustrating that in any subgame perfect equilibrium, agents stop experimenting once the common belief falls beneath the single-agent cutoff. Furthermore, we prove that in any subgame perfect equilibrium outcome, agents engage in either exactly the single-agent amount of experimentation or in one additional experiment.

We then turn to the main focus of this chapter: strategic experimentation under private information. Agents observe each others' behavior but not the realized payoffs. Furthermore we allow agents to communicate via cheap talk.[1] In such a context, agents have three sources of information: their own signals, their observation of other agents, and the cheap talk messages. We begin by noting that cheap talk in this environment is very effective. In particular, we show that for every subgame perfect equilibrium with publicly observable payoffs, there exists a perfect Bayesian equilibrium with privately observable payoffs that yields the same outcome. The key intuition lies in the fact that truthful communication is easy to sustain. Following a success, a player is certain about the underlying state of the world and is willing to truthfully communicate her success. Furthermore, if a player believes that the other player is communicating truthfully, she believes that the state of the world is good with probability 1 upon hearing that her fellow player had a success. In this case, she is willing to play the risky arm forever without communicating her future payoffs. But then no player will want to wrongly announce a success because this will make it impossible to learn anything from one's fellow player in the future. Truthful revelation of payoffs is always incentive compatible, therefore, and any outcome with observable payoffs can be replicated when payoffs are only privately observable.

More surprisingly, we also prove that if the initial belief is sufficiently optimistic, then the socially optimal symmetric strategy profile can be supported as a perfect Bayesian equilibrium outcome.[2] To see the logic behind the equilibrium construction, suppose that the optimal symmetric solution of the social planner requires $\tau > 1$ periods of experimentation. Since the initial belief is above the myopic one, both players are willing to experiment in the first period. Thereafter, player i may prefer to wait while her fellow player j engages in costly experimentation. Following such a deviation of not experimenting by player i, however, player j believes that i had a prior success and, hence, j is willing to experiment forever without communicating.[3] This implies that player i cannot learn anything from player j once she deviates and waits, and thereby makes free-riding impossible.

While these beliefs are consistent in the sense of sequential equilibrium, one may find them not fully convincing. This leads us to consider an ad-hoc refinement of pessimistic beliefs in which a player who observes a deviation to the safe action believes that the deviating player's previous experiments all failed. We prove that the socially optimal symmetric strategy profile can be supported as a perfect Bayesian equilibrium with pessimistic beliefs if the player's initial belief is sufficiently optimistic.

The remainder of the chapter proceeds as follows. Section 2 sets up the model. Section 3 solves the social planner's problem. Section 4 studies strategic experimentation, first with public payoffs, then with private payoffs. Section 5 concludes and discusses possible extensions.

2. The Model

There is an infinite number of periods $t = 0, 1, \ldots$ and there are two players who in each period choose between a safe and a risky action. Before making this choice, they can costlessly communicate with each other.

[1] It is very natural to let agents communicate in a strategic environment in which there is only an information externality between agents. A similar approach has been taken in collusion models with private monitoring—where communication however is illegal in contrast to the current setting—to facilitate the analysis; see e.g. Compte (1998) and Kandori and Matsushima (1998).

[2] Focussing on the socially optimal *symmetric* profile entails only a minor efficiency loss. In fact, we show that the "unrestricted" social optimum can be achieved by a strategy profile such that there is at most one period in which a single player experiments.

[3] This "punishment via beliefs" construction is similar in spirit to the one sustaining collusion in Blume and Heidhues (2006).

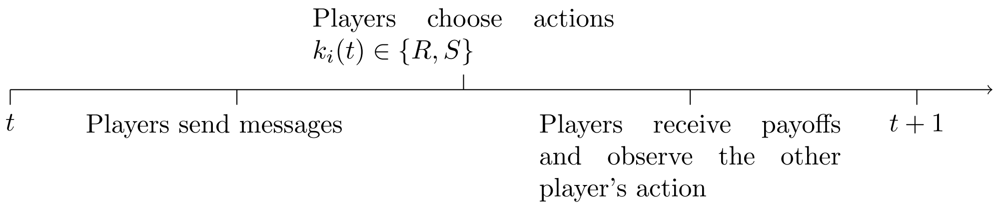

Figure 1. Timeline of the game.

More precisely, at the beginning of period t, each player i chooses a cheap-talk message $m_i(t) \in [0,1]$. Upon having observed the other player's cheap-talk message, each player i then chooses an action $k_i(t) \in \{R, S\}$. If $k_i(t) = S$, the player receives a safe payoff normalized to 0; if $k_i(t) = R$, the player receives a risky payoff $X_i(t)$ that is either low (X_L) or high (X_H), where $X_L < 0 < X_H$.

The distribution of the risky payoff depends on an unknown state of the world, which is either good ($\theta = 1$) or bad ($\theta = 0$). Conditional on the state of the world, payoffs are drawn independently across players and periods. In the good state of the world, the probability of the high payoff $\mathrm{P}(X_H|\theta = 1) = \pi \in [0, \frac{1}{2}]$; in the bad state, it is $\mathrm{P}(X_H|\theta = 0) = 0$. Thus, a single successful experiment proves that the state of the world is good. This makes our model a discrete-time analog of the model analyzed in Keller, Rady and Cripps (2005). We write E_θ for the conditional expectation $\mathrm{E}[X_i(t)|\theta]$ of the risky payoff in any given period, and assume that $E_0 < 0 < E_1$.

Our primary interest below is in analyzing the game in which the realizations of the payoffs $X_i(t)$ are private information but the choice whether or not to experiment is observable. When considering this game, we partition the set of private histories for each player into those after which he has to send a message and those after which he has to choose an action. Define the set of possible observations in any given period by

$$O_i = \underbrace{\{(R, X_L), (R, X_H), (S, 0)\}}_{\text{own action \& payoff}} \times \underbrace{\{R, S\}}_{\text{observed action}} .$$

Formally, the set of *private message histories* of player i at time t is

$$H^m_{i,t} = \left(\underbrace{[0,1]^2}_{\text{messages}} \times \underbrace{O_i}_{\text{actions \& payoff}} \right)^t .$$

Define the set of all private message histories $H^m_i = \bigcup_{t=0}^{\infty} H^m_{i,t}$. Similarly, the set of all *private action histories* at time t is defined by

$$H^a_{i,t} = \left(\underbrace{[0,1]^2}_{\text{messages}} \right)^{t+1} \times \left(\underbrace{O_i}_{\text{actions \& payoff}} \right)^t ,$$

and the set of all private action histories by $H^a_i = \bigcup_{t=0}^{\infty} H^a_{i,t}$.

A *pure strategy* is thus a measurable mapping that assigns to each private message history $h^m_i \in H^m_i$ a message $m_i(h^m_i) \in [0,1]$ and to each private action history $h^a_i \in H^a_i$ an action $k_i(h^a_i) \in \{R, S\}$. Mixed strategies are defined in the usual way.

Given a probability $p_i(0) = p$ that player i initially assigns to the good state of the world, his expected payoff from a pure-strategy profile is

$$(1-\delta)\,\mathrm{E}_p\left[\sum_{t=0}^{\infty}\delta^t \mathbf{1}_{\{k_i(h(t))=R\}} X_i(t)\right],$$

where the factor $1-\delta$ serves to express the overall payoff in per-period units. Note that player j's strategy only enters through the expectation operator—there is just an informational externality at play here.

We will assume that players start with a common prior $p = p_1(0) = p_2(0)$ and solve for *perfect Bayesian equilibria* (PBE) of the game. We call an equilibrium perfect Bayesian if all players act optimally after every history given their beliefs and the other players' strategies, and if the players' beliefs are updated according to Bayes' rule whenever possible.

As a benchmark, we will also study the game with observable payoffs and no communication. Starting from a common prior, this is a stochastic game in which the posterior belief $p(t)$ evolves according to Bayes' rule from one period to the next. For this game, we shall use the solution concept of *subgame perfect equilibrium* (SPE).

Whether payoffs are public information or not, we shall say that a player *experiments* if she chooses the risky action while still being uncertain about the true state of the world. For future reference, we define

$$p^m = \frac{|E_0|}{|E_0| + E_1}\,.$$

This is the belief at which the expected current payoff from the risky option just equals zero, i.e. the safe payoff. A myopic player chooses the risky arm if and only if his posterior belief exceeds p^m. We therefore call p^m the myopic cutoff belief.

3. The Planner's Problem

In this section, we discuss the problem of a social planner who chooses a strategy profile to maximize the average of the two players' objective functions. The optimal outcome in this situation will serve as a benchmark for the case that individual players pursue their goals independently.

Let $k = (k_1, k_2)$ denote a pure strategy profile. Then the expected average payoff, expressed in per-period units, is

$$U(p,k) = (1-\delta)\,\mathrm{E}_p\left[\tfrac{1}{2}\sum_{i=1}^{2}\sum_{t=0}^{\infty}\delta^t \mathbf{1}_{\{k_i(h(t))=R\}} X_i(t)\right],$$

where p denotes the probability that the planner initially assigns to the good state.

For the social planner, it can never be strictly beneficial to have players hide information from each other because he can always choose a strategy profile that ignores unwanted information. Hence, when discussing the planner's problem, we will focus on strategy profiles in which all players truthfully communicate their past payoffs via their cheap-talk messages. The planner's problem then becomes a Markovian decision problem with the posterior belief as the single state variable.

A single success in the past fully reveals that the state of the world is good. It is then a dominant choice for the planner to have both players take the risky action in all following periods. Using this fact and restricting attention to symmetric strategy profiles, we can think of the planner as choosing a period τ such that both agents experiment in periods $t \leq \tau - 1$ and, in case all these experiments were unsuccessful, no agent experiments in periods $t \geq \tau$. For any such τ, expected average payoffs are

$$\begin{aligned} U(p,\tau) &= (1-\delta)\left\{(1-p)\sum_{t=0}^{\tau-1}\delta^t E_0 + p\sum_{t=0}^{\tau-1}\delta^t E_1 + p[1-(1-\pi)^{2\tau}]\sum_{t=\tau}^{\infty}\delta^t E_1\right\} \\ &= (1-\delta^\tau)E_p + \delta^\tau p[1-(1-\pi)^{2\tau}]E_1 \end{aligned}$$

with $E_p = pE_1 + (1-p)E_0$. Since the difference

$$U(p,\tau+1) - U(p,\tau) = \delta^\tau\left\{(\delta-1)(1-p)|E_0| - (1-\pi)^{2\tau}pE_1(\delta(1-\pi)^2-1)\right\}$$

is positive if and only if

$$\tau < \frac{1}{2\ln(1-\pi)}\left(\ln\frac{\delta-1}{\delta(1-\pi)^2-1} + \ln\frac{1-p}{p} + \ln\frac{|E_0|}{E_1}\right),$$

the value function of the social planner is $v^{sc}(p) = U(p,\tau^{sc}(p))$ with[4]

$$\tau^{sc}(p) = \left\lfloor\frac{1}{2\ln(1-\pi)}\left(\ln\frac{\delta-1}{\delta(1-\pi)^2-1} + \ln\frac{1-p}{p} + \ln\frac{|E_0|}{E_1}\right)\right\rfloor.$$

As $\tau^{sc}(p) = 0$ if and only if p lies below the cutoff

$$p^{sc} = \frac{(1-\delta)|E_0|}{(1-\delta)(E_1+|E_0|) + \delta(2-\pi)\pi E_1},$$

we have the following result.

Proposition 3.1 (**Optimal symmetric pure-strategy profile**). *Among all pure symmetric strategy profiles, the following maximizes the players' expected average payoff: both players always communicate their payoffs truthfully, both choose the risky arm when $p(t) \geq p^{sc}$, and both choose the safe arm otherwise.*

If the social planner can use asymmetric strategy profiles, it is optimal for him to experiment beyond the belief p^{sc}. In fact, the expected discounted payoff from letting one player experiment and stopping all experimentation thereafter, $(1-\delta)\frac{1}{2}E_p + \delta p\pi E_1$, is positive above the cutoff

$$p^c = \frac{(1-\delta)|E_0|}{(1-\delta)(E_1+|E_0|) + \delta 2\pi E_1} < p^{sc}.$$

For the full description of the social optimum, we will also need the following cutoff:

$$\tilde{p}^c = \frac{(1-\delta)|E_0|}{(1-\delta)(E_1+|E_0|) + \delta 2(1-\pi)\pi E_1}.$$

It is straightforward to see that $\tilde{p}^c > p^{sc}$ and that starting at $p = \tilde{p}^c$, one failed experiment takes the posterior belief below p^c.

Proposition 3.2 (**Optimal strategy profile**). *There exists a socially optimal strategy profile in which both players always communicate their payoffs truthfully, both choose the risky arm when $p(t) \geq \tilde{p}^c$, one player chooses the risky arm when $\tilde{p}^c > p(t) \geq p^c$, and both choose the safe arm otherwise. Furthermore, there is at most one period in which just one agent experiments.*

Proof. See the appendix. □

[4]The superscript "sc" indicates the "symmetric cooperative" solution. For any real number x, the *floor* $\lfloor x \rfloor$ is the largest integer not greater than x.

Restricting the planner to symmetric strategy profiles thus entails only a small loss in expected average payoffs.

4. Strategic Experimentation

We now turn to the analysis of strategic experimentation, first with public payoffs, then with private payoffs. While the planner's solution identified in the previous section constitutes an upper bound on the average of the two players' equilibrium payoffs, the solution to the single-agent bandit problem constitutes a lower bound on each player's individual equilibrium payoff. In fact, each player always has the option to ignore the information contained in the opponent's actions and payoffs (if the latter are observable).

A simpler version of the arguments leading up to Proposition 3.1 establishes that the single-agent (or "autarky") solution is given by the cutoff belief

$$p^a = \frac{(1-\delta)|E_0|}{(1-\delta)(E_1+|E_0|)+\delta\pi E_1} < p^{sc}\,,$$

where—as in the planner's solution—we adopt the convention that the agent experiments when she is indifferent, that is, when her belief equals p^a. We shall show that with publicly observable payoffs, equilibrium experimentation cannot go beyond this cutoff. With privately observed payoffs and cheap-talk communication, by contrast, more efficient equilibria can be played.

4.1. Public Payoffs. As both players choose the risky arm after any history $h(t)$ in which a success has been observed, we can restrict our attention to histories with no prior success. We begin our analysis of equilibrium behavior with the observation that in every SPE of the game with public payoffs, both players choose the safe arm after any history that takes their common belief below the social planner's cutoff p^c. To see this, suppose to the contrary that there exists an SPE in which a player experiments at some belief $p < p^c$. From the analysis of the planner's solution, we know that the average of the players' objective functions at p is negative. Consequently there needs to be at least one player who receives a negative expected payoff. By deviating and always choosing the safe arm this player can increase her payoffs.

To get a first intuition for why equilibrium experimentation cannot go beyond the single-agent cutoff, consider pure-strategy SPE first. Since players do not experiment below p^c, there are only finitely many periods in which a pure-strategy equilibrium can require players to experiment in the absence of a prior success. In the last period in which a player is meant to experiment, the player knows that if she fails, no player will experiment in future. Hence, she will only be willing to experiment if this is individually optimal, that is, if the current belief is at least p^a. When both players are meant to experiment, the belief must be above p^a because the value of experimenting in this last period is lower if one's fellow player also experiments. Hence, in any pure-strategy equilibrium there can be at most one more experiment than in the single-agent solution.

Conversely, it cannot be the case that both players permanently stop experimenting at a belief above the single-agent cutoff. The reason is simply that each player—believing that the other player stopped experimenting—would then face the single-agent tradeoff. The following proposition exploits this logic and extends it to mixed-strategy equilibria.

We call the number of times a player chooses the risky arm on the path of play when every experiment is unsuccessful the *amount of experimentation* performed by that player. The *total amount of experimentation* by both players is simply the sum of the individual amounts. The total amount will typically depend on the initial belief and, with mixed strategies, may be a random variable.

PROPOSITION 4.1. *Given an initial belief, let the optimal amount of experimenation in the single-agent problem be K. In any SPE of the experimentation game with public payoffs, the total amount of experimentation is K or $K+1$.*

PROOF. First, consider any history of length t for which $p^a \leq p(t) < 1$. Since $p(t) < 1$, no prior experiment has been successful, and since $p(t) \geq p^a$, the total amount of experimentation by both players is less than K. We now argue that players experiment with probability 1 at least one more time following any such history. Let $v^a(p(t)) > 0$ be the value of the single-agent problem at the belief $p(t)$. Let τ be the smallest integer such that $\delta^\tau E_1 < v^a(p(t))$, and ψ the probability that there will be at least one experiment in the periods $t, t-1, \ldots, t+\tau-1$. The period t continuation value of each player is then trivially bounded above by $(1-\psi)\delta^\tau E_1 + \psi E_1$. So ψ cannot be smaller than $\bar{\psi} = [v^a(p(t)) - \delta^\tau E_1]/[(1-\delta^\tau)E_1]$, because otherwise it would be profitable for each player to deviate to the single-agent solution. For $n = 1, 2, \ldots$, the probability that no player experiments in the next $n\tau$ periods is bounded above by $(1-\bar{\psi})^n$. Letting $n \to \infty$, we see that there will almost surely be another experiment on the path of play.

Next, consider a history of length t for which $p(t) = p < p^a$. Let ϕ be the probability with which player j experiments at time t. Suppose that the SPE requires player i to experiment with positive probability. Then she can do no better by switching to the strategy of playing safe now and, in case player j experiments and is unsuccessful, continuing to play safe forever. This implies

$$\delta\phi p\pi E_1 \quad \leq \quad (1-\delta)E_p + \delta\left\{p[\pi + \phi\pi - \phi\pi^2]E_1 + \left(1 - p[\pi + \phi\pi - \phi\pi^2]\right)v\right\},$$

where $\pi + \phi\pi - \phi\pi^2$ is the probability of at least one success, and v player i's continuation value after a double failure, that is, a payoff realization $X_1(t) = X_2(t) = X_L$. As $0 \leq v \leq E_1$, this in turn requires that

$$0 \leq (1-\delta)E_p + \delta\left\{p\pi E_1 + (1-p\pi)v\right\}.$$

As $p < p^a$, we have $(1-\delta)E_p + \delta p\pi E_1 < 0$, and hence $v > 0$. So some player must experiment with positive probability in round $t+1$ or later. Repeating this step until a time $t+\tau$ at which $p(t+\tau) < p^c$ in the absence of a success, we obtain a contradiction because no player can experiment below p^c in equilibrium. □

Keller, Rady and Cripps (2005) establish in a continuous-time setup that with fully revealing successes on the risky arm, the amount of experimentation is limited by the single-agent amount in any Markov perfect equilibrium. Since we have a discrete-time setup, we can relax the assumption of Markov strategies and, with a minor qualification due to discrete rather than continuous time, establish that their finding extends to all subgame perfect equilibria.

For future reference, we note

COROLLARY 4.1. *Whenever the total amount of experimentation in the planner's optimal (or optimal symmetric) strategy profile exceeds the single-agent amount by more than 1, it cannot be implemented in a subgame perfect equilibrium of the experimentation game with observable payoffs.*

Above, we have fully characterized subgame perfect equilibria in terms of the total amount of experimentation that is carried out on the path of play. These, equilibria, however, may differ in other dimensions such as when agents experiment. For example, just above p^a agents may engage in a war of attrition as to who has to carry out the final experiment. Thus, there may be periods in which no agent experiments. Furthermore, agents may use their communication to coordinate on whether a given agent is meant to experiment in a given period. Below, we nevertheless show that the entire set of equilibrium paths of experimentation with observable payoffs can be replicated in perfect Bayesian equilibria of the game with unobservable payoffs. Moreover, we show that under certain conditions, higher amounts of experimentation can be supported when payoffs are unobservable.

For future reference, it is useful to establish the existence of a symmetric subgame perfect equilibrium. Since in all subgame perfect equilibria the total number of experiments is bounded by $K+1$, we can restrict attention to finitely many beliefs when constructing equilibria. Using these finitely many beliefs as states (and thinking of the lowest belief as an absorbing state), it follows that there exist only finitely many pure Markov perfect strategies. We construct an auxiliary game in which players' finite action set corresponds to the set of pure Markov perfect strategies and payoffs are defined as in the original game. This is a symmetric finite game and the existence of a symmetric mixed Nash equilibrium in this auxiliary game follows from Nash (1951). Because there always exists a Markov perfect best response to a Markov perfect strategy, the equilibrium in the auxiliary game corresponds to a Markov perfect equilibrium in the original game, which trivially is also subgame perfect.

Below, we use the following observation proven in the Appendix:

Lemma 4.1. *The payoff in any symmetric subgame perfect equilibrium is lower than the payoff in the optimal symmetric pure-strategy profile of Proposition 3.1.*

4.2. Private Payoffs. We begin by noting that truthful communication can easily be sustained with privately observed, fully revealing payoffs. Suppose that after every period in which a player experimented, she announces a first success by sending the message $m_i(t) = 1$ and randomizes uniformly over all other messages otherwise. Furthermore, after the first success has been announced and both players know that the state of the world is good, suppose there are no meaningful messages any more, that is, both players always randomize uniformly over the interval $[0, 1]$. Similarly, a player who did not experiment randomizes uniformly over $[0, 1]$, i.e. babbles. Finally, on and off the equilibrium path, players believe that past communication by the other player was truthful. Intuitively, we are then back in the case of public payoffs.[5]

The key observation now is that if players anticipate this communication strategy, truthful communication is incentive compatible. Following a first success on player i's risky arm, player i knows the state of the world and hence is indifferent as to what player j believes. So truthfully announcing a success is optimal for player i following an own success. After such an announcement, player j believes with certainty that the state of world is good, and hence will play risky in all future periods independent of what player i does after the announcement. If player i incorrectly announces a success, she cannot infer anything from player j's future behavior, so she is weakly better off telling the truth. We thus have the following result.

Proposition 4.2. *For every SPE of the game with public payoffs, there exists an outcome-equivalent PBE of the game with private payoffs.*

So far, we have established that private information does not hurt players who play subgame perfect equilibria. Our next result establishes that players can often do better. We require players to perform the

[5] One caveat here is that players may use the messages to create a controlled joint lottery to coordinate continuation play in the SPE with observable payoffs. But then we can encode the original message by using the odd digits only, while using the even digits to send the messages constructed above.

optimal symmetric amount of experimentation whereafter, on the path of play, they once communicate and announce whether they had a prior success. If so, both players keep experimenting forever; otherwise, both stop experimenting. We punish early deviations (after the initial period) through beliefs: if a player refrains from experimenting at a time when the socially optimal symmetric strategy profile requires her to experiment, then the other player reacts to this out-of-equilibrium event by assigning probability 1 to the good state of the world. Formally, our equilibrium concept would allow us to assign the same optimistic beliefs to a player who observes a deviation at $t = 0$. Such beliefs, however, are clearly implausible: a player deviating in $t = 0$ cannot have seen a prior success, and so we will—in the spirit of sequential equilibrium—require the other player not to update her beliefs in response to the deviation.[6]

Proposition 4.3. *There exists a threshold $p^\dagger < p^m$ such that for all initial beliefs $p \geq p^\dagger$, the experimentation game with private payoffs admits a PBE that is outcome-equivalent to the socially optimal symmetric pure-strategy profile.*

Proof. First, we specify the players' strategies, beginning with the behavior after all message histories and then turning to action histories. Play following any history in which at least one player did not experiment in period 0 is specified separately below; following such a history, we will prescribe players to play a PBE with truthful communication. For brevity, we do not explicitly specify behavior following observable simultaneous deviations as this is irrelevant for the incentives on the path of play.[7]

If both players experimented in all periods $t \leq \tau^{sc} - 1$, they both report truthfully in period τ^{sc}; following any other message history in which both players experimented in period 0, they send babbling messages. Both players experiment in any period $t \leq \tau^{sc} - 1$ as long as both experimented in all prior periods. If both players experimented in all periods $t \leq \tau^{sc} - 1$ and player j announces a success in period τ^{sc}, player i experiments in all periods $t \geq \tau^{sc}$. If both players experimented in all periods $t \leq \tau^{sc} - 1$ and player j announces no success in period τ^{sc}, player i experiments in any period $t' \geq \tau^{sc}$ if and only if he had a success himself in a period $t'' \leq t' - 1$. We are left to specify behavior for all action histories in which at least one player did not experiment in any period $t \leq \tau^{sc} - 1$. If the first single deviation occurs in a period $t' \in \{1, \ldots, \tau^{sc} - 1\}$, and player j is the one who deviates, then in all periods $t \geq t' + 1$, both players babble, player i experiments and player j plays the autarky strategy (conditioning her behavior on her own signals only). Recall that the case of a first deviation in period $t = 0$ is handled separately below.

Second, we specify the players' beliefs if there was no deviation in period $t = 0$. It is convenient to specify the players' beliefs about whether the state of the world is good (rather than the usual beliefs about nodes in an information set).[8] Both players use Bayesian updating on the path of play; in particular, if both players experimented in all periods $t \leq \tau^{sc} - 1$, both update under the assumption that the opponent reports truthfully in period τ^{sc}. If the first single deviation occurs in a period $t' \in \{1, \ldots, \tau^{sc} - 1\}$, and player j is the one who deviates, then player i switches irrevocably to the belief that the state of the world

[6] We could avoid this problem by letting the players observe one draw from the distribution of risky payoffs before the game starts. The following proposition would then hold with $p^\dagger = p^{sc}$.

[7] One possible specification is to simply ignore a simultaneous failure to experiment in $t = 0$ by "restarting" the game, and to require both players to babble and experiment forever following a history in which both players simultaneously failed to experiment for the first time in a period $t \in \{1, ..., \tau^{sc} - 1\}$.

[8] A probability distribution about possible nodes in a player's information set (where a node can be identified through whether and when the other player had a success when experimenting in addition to one's observations), can be constructed in the obvious way from the probability that the state of the world is good together with how often the other player experimented; this probability distribution is unique but for the fact that we can arbitrarily prescribe when another player had a success following an out-of-equilibrium observation.

is good, while player j continues to apply Bayes' rule, conditioning on her own observations only. If the first deviation occurs in a period $t \geq \tau^{sc}$ by player i, player i ignores any future behavior of player j and continues to use Bayes' rule to update his beliefs conditioning only on his own experimentation results. Player j's belief is unaffected by player i's deviation or behavior thereafter, and player j updates using only his own experimentation results from thereon.

Third, we prove sequential rationality if there was no deviation in period $t = 0$. Any player who had a success in a period $t \leq \tau^{sc} - 1$ is willing to announce it truthfully in round τ^{sc}; given that the other player believes the announcement to be truthful (i.e. has the belief that the state of the world is good with probability 1), it is optimal for her to choose the risky action forever. If the first single deviation occurs in a period $t' \in \{1, \ldots, \tau^{sc} - 1\}$, and player j is the one who deviates, then player i holds the belief that the state of the world is good and hence it is optimal for him to experiment in all subsequent periods. As a consequence, player j ceases to learn anything from observing i's future behavior, hence finds herself in an autarky situation. Consequently it is optimal for her to play the autarky strategy and beliefs are consistent with this. If player j does not deviate, she obtains the value of the symmetric cooperative solution. Since the latter is always weakly larger than the autarky value, the deviation is unprofitable. We are left to consider behavior following histories in which the first deviation occurred in a period $t \geq \tau^{sc}$. If both players experimented in all periods $t \leq \tau^{sc} - 1$ and either player i had a success or player j announced a success, it is obviously optimal for player i to always experiment since he believes that the state of the world is good. If both players experimented in all periods $t \leq \tau^{sc} - 1$ and player j announced no success, player i's belief is below p^{sc} if he has no prior success himself, and hence it is optimal for him not experiment.

Fourth, we construct a continuation equilibrium that deters deviations at $t = 0$. For $n \in \mathrm{N}$, let

$$B(n,p) = \frac{p(1-\pi)^n}{p(1-\pi)^n + 1 - p}\,.$$

This is the posterior probability assigned to the good state of the world after n failed experiments. Fix a symmetric SPE of the game with observable payoffs starting with the common prior $B(1,p)$. Let v be the equilibrium value in this SPE, and note that $v \leq v^{sc}(B(1,p))$. Consider a PBE of the game with private payoffs whose experimentation path coincides with that of the selected SPE; such a PBE exists by Proposition 4.2. Following a single deviation by player j in $t = 0$, we require player i to communicate truthfully whether he had a success; by the argument underlying Proposition 4.2, this is incentive compatible. If he announces a failure, both players' continuation play corresponds to the selected PBE.

So, if player j experiments at $t = 0$, her expected overall payoff is

$$v^{sc}(p) = (1-\delta)E_p + \delta \underbrace{\left\{p[1-(1-\pi)^2]E_1 + \left(1 - p[1-(1-\pi)^2]\right) v^{sc}(B(2,p))\right\}}_{(I)};$$

if she deviates, this payoff is no more than

$$\delta \underbrace{\left\{p\pi E_1 + (1-p\pi)\, v^{sc}(B(1,p))\right\}}_{(II)}\,.$$

Being the upper envelope of linear functions, v^{sc} is convex. This implies that $v^{sc}(B(1,p)) \leq B(1,p)\pi E_1 + (1 - B(1,p)\pi)\, v^{sc}(B(2,p))$. As the three points $(B(1,p), v^{sc}(B(1,p)))$, $(B(2,p), v^{sc}(B(2,p)))$ and $(1, E_1)$ do not lie on a line, this inequality is in fact strict. Using this and the fact that $(1-p\pi)B(1,p) = p(1-\pi)$, one has $(I) > (II)$. As $E_p \geq 0$ for $p \geq p^m$, this completes the proof. $\square$

In our equilibrium construction, deviations in the initial period are punished through a symmetric continuation equilibrium. Furthermore, in the above determination of the critical initial belief $p^\dagger$ we bound the payoff of such a symmetric continuation equilibrium from above by that of the optimal symmetric strategy profile a social planner would choose. In general, we can of course find harsher punishments for a player who deviates in the first period. For example, we could search for the subgame perfect equilibrium in the game with observable payoff that minimizes the payoff of a given player i. By playing a continuation equilibrium that implements this path of play following an initial deviation, we would punish such deviations more severely and hence increase the range of initial beliefs for which the symmetric social optimal strategy profile can be implemented. In the above proposition as well in Proposition 4.4 below, we refrain from doing so for ease of exposition.

The belief following a deviation in an early (but not the initial) round may seem somewhat unusual in the above equilibrium. Intuitively, upon observing such a deviation, player i reasons as follows: "Clearly, player j was not careful and made a mistake. She must already know that the state of the world is good to be so careless." While this reasoning is compatible with the logic of sequential equilibrium, the equilibrium construction hinges crucially on this particular choice of off-equilibrium beliefs. Our next aim is therefore to show that private payoffs can lead to a more efficient outcome even under the stringent requirement that whenever a player observes a deviation to the safe action, her beliefs become as pessimistic as possible.

Definition 2 (Pessimistic Beliefs). *We say that a perfect Bayesian equilibrium has pessimistic beliefs if a player who observes a deviation to the safe action and does not yet assign probability 1 to the good state of the world, believes that the deviating player had only failures before the deviation.*

Recall from Proposition 4.2 that for any subgame perfect equilibrium with publically observable payoffs, we can find a perfect Baysian equilibrium that implements the same experimentation path. In our construction, players truthfully communicate after every period in which they experimented until a first success is announced.

Lemma 4.2. *The perfect Baysian equilibrium constructed in Proposition 4.2 has pessimistic beliefs.*

Proof. If player i had a success, he believes that the state of the world is good with probability 1. Similarly, once player j has announced a success, the equilibrium is such that player i believes with probability 1 that there has been a success, and assigns probability 1 to the good state of world from then on. Consider a history, therefore, in which neither player j announced a success nor player i observed a success himself. As player i believes that player j communicated truthfully in the past, he believes that all of player j's experiments have been failures. So his beliefs are indeed pessimistic. $\square$

Our next proposition shows that even if we restrict ourselves to pessimistic beliefs, we can find perfect Bayesian equilibria that implement the optimal symmetric strategy profile for sufficiently optimistic starting beliefs. To see the intuition, recall that in the symmetric optimum, the planner updates his beliefs on the basis of two experiments in every period until the belief falls below p^{sc}. Absent meaningful communication, however, each player updates her belief using only the result of her own experimentation, and hence beliefs decrease at a slower rate. The key observation is that for high enough starting beliefs, this slower learning implies that the social planner reaches p^{sc} before players who do not communicate reach p^m. Above p^m, players myopically prefer to experiment, and in the equilibrium that we will construct, they do so up to the period in which a social planner who has not observed a success would cease experimentation. At that point in time, players truthfully communicate and continue experimenting only if at least one player had a prior success. A player who deviates by not experimenting prior to this point in time reduces her myopic payoff and induces a symmetric continuation equilibrium in which payoffs, of course, are weakly lower than in the optimal symmetric strategy profile; thus, such a deviation is unprofitable.

Proposition 4.4. *There exists a threshold $p^\ddagger < 1$ such that for all initial beliefs $p \geq p^\ddagger$ there exists a PBE with pessimistic beliefs that is outcome-equivalent to the socially optimal symmetric pure-strategy profile.*

Proof. Recall that there exists a symmetric Markov perfect equilibrium in the case of observable payoffs. Choose such a symmetric Markov perfect equilibrium for any starting belief p. By Proposition 4.2 there exists a PBE that implements the same mapping from payoff histories to action profiles. Denote this PBE by $\sigma(p)$.

Recall the definition of $B(n, p)$ from the proof of Proposition 4.3. Let $m = \min\{n \in \mathrm{N} : B(n, p^m) < p^{sc}\}$, and define $p^\ddagger$ implicitly by $p^m = B(m, p^\ddagger)$. Note that $B(2m, p^\ddagger) < p^{sc} \leq B(2m-1, p^\ddagger)$.

We are now ready to specify strategies.[9] In all periods $t \leq m-1$, both players babble and experiment provided no player chose the safe arm in a previous period $t' \leq t-1$. If a single player deviated and chose the safe arm in a period $\tau \leq m-1$, the players truthfully communicate at the beginning of period $\tau+1$. If player i communicated truthfully in period $\tau+1$, he plays the strategy prescribed by $\sigma(q)$ in periods $t \geq \tau+1$, where $q = B(2\tau-1, p)$ if no player announced a success, and $q = 1$ otherwise. If player i did not communicate truthfully in period $\tau+1$, he either had a success he did not announce or he incorrectly announced a success he did not have. In the former case, he uses the risky arm in periods $t \geq \tau+1$; in the latter case, he plays the autarky strategy.

We are left to specify strategies in case both players experimented in all periods $t \leq m-1$. In this case, both players truthfully announce at the beginning of period m whether they had a success in any of the previous rounds. If player j announces a success, or if player i observed a success herself in any prior period, player i babbles and plays the risky arm in all periods $t \geq m$. If player i observed no success herself, and player j does not announce a prior success, player i babbles and plays the safe arm in all periods $t \geq m$.

Next, we specify beliefs about the state of the world. Any player who had a prior success believes the state of the world is good with probability 1. In any period $t \leq m-1$ such that both players experimented in all periods $t' \leq t$, the belief of agent i is $p_i(t) = B(t-1, p)$ if all his experiments failed. If a player deviated by not experimenting in a period $\tau \leq m-1$, she believes the state of the world is good with probability 1 if the other player announces a success in period $\tau+1$ (or she had a prior success herself), otherwise her belief is $q = B(2\tau-1, p)$. Thereafter, beliefs are as in the PBE $\sigma(q)$.

Now suppose that both players experimented in all periods $t \leq m-1$. Each player believes the other's message in period m to be truthful, and all subsequent messages to be uninformative. If player i deviates in period m by incorrectly announcing a success he did not have, and player j does not announce a success in period m, then player i's belief in period $\tau \geq m$ equals 1 if he experiences a success in one of the rounds $t = m, \ldots, \tau-1$, and equals $B(2m+n, p)$ if he carries out n experiments in periods $t = m, \ldots, \tau-1$ and they all fail.

As the beliefs that we have specified follow Bayes' rule whenever possible, it remains to show sequential rationality. Each player experiments whenever he assigns probability 1 to the good state of the world, which

[9]We follow the usual convention again and ignore simultaneous deviations.

is clearly optimal. If a single player deviated by choosing the safe arm in a period $\tau \le m-1$, it is optimal for both players to communicate truthfully at the beginning of period $\tau+1$ by the argument underlying Proposition 4.2. If player i communicates truthfully in period $\tau+1$, he believes with probability 1 that player j plays according to $\sigma(q)$ with q specified above. As $\sigma(q)$ constitutes a PBE, it is optimal for player i to play according to $\sigma(q)$ as well. If player i announces a success in period $\tau+1$ that he did not have, player j experiments forever, thus player 1 finds himself in an autarky situation and optimally plays the autarky strategy. If player i does not announce a success he had, it is clearly optimal for him to take the risky arm forever. Finally, if player i deviates in a period $\tau \le m-1$, his belief about the state of the world is weakly above p^m; by the same argument as the one used in the proof of Proposition 4.3 for deviations in period $t=0$, such a deviation is unprofitable.

We are left to rule out deviations by player i in a period $t \ge m$ following a history in which both players experimented in all periods $t \le m-1$, player i had no success, and player j did not announce a success in period m. In this case player i's belief is below p^{sc}. As $p^{sc} < p^a$ and player i does not expect player j to experiment in future periods, it is optimal for player i to play the safe arm. □

Starting from the definition of $p^\ddagger$ in the proof, some straightforward computations show that

$$p^\ddagger \ge \frac{(1-\pi)[1-\delta+\delta(2-\pi)\pi]|E_0|}{(1-\pi)[1-\delta+\delta(2-\pi)\pi]|E_0|+(1-\delta)E_1} > p^m$$

and

$$p^\ddagger < \frac{[1-\delta+\delta(2-\pi)\pi]|E_0|}{[1-\delta+\delta(2-\pi)\pi]|E_0|+(1-\delta)E_1}.$$

For the reasons mentioned after the proof of Proposition 4.3, however, the interval $[p^\ddagger, 1]$ constitutes just part of the set of initial beliefs for which the outcome of the socially optimal symmetric pure-strategy profile can be obtained in a PBE with pessimistic beliefs.

5. A Continuous-Time Limit

We now take our discrete-time framework to a continuous-time limit that coincides (up to a normalization of the safe payoff to zero) with the two-player version of the setup studied by Keller, Rady and Cripps (2005).

Let time be continuous and suppose that operating the risky arm comes at a flow cost of $s > 0$ per unit of time. In the good state ($\theta = 1$), the risky arm yields lump-sum payoffs which arrive at the jump times of a Poisson process with intensity $\lambda > 0$. These lump-sums are independent draws from a time-invariant distribution with known mean $h > 0$, and the Poisson processes in question are independent across the two players. In the bad state ($\theta = 0$), the risky arm never generates a lump-sum payoff. The safe arm does not produce any such payoffs either, but is free to use.

Given the common discount rate $r > 0$, a player's payoff increment from using a bad risky arm for a length of time $\Delta > 0$ is

$$\int_0^\Delta r\, e^{-r\,t}\,(-s)\,dt = \left(1-e^{-r\,\Delta}\right)(-s).$$

The expected discounted payoff increment from a good risky arm is

$$\mathrm{E}\left[\int_0^\Delta r\, e^{-r\,t}\,(h\,dN_t - s\,dt)\right] = \int_0^\Delta r\, e^{-r\,t}\,(\lambda h - s)\,dt = \left(1-e^{-r\,\Delta}\right)(\lambda h - s);$$

here N_t is a standard Poisson process with intensity λ, and the first equality follows from the fact that $N_t - \lambda t$ is a martingale. We assume $\lambda h > s$ so that a good risky arm dominates the safe arm. Finally, the probability of observing at least one lump-sum on a good risky arm during a time interval of length Δ is $1 - e^{-\lambda\Delta}$.

If we let the players adjust their actions only at the times $t = 0, \Delta, 2\Delta, \ldots$ for some fixed $\Delta > 0$, we are back in our discrete-time framework with $\pi = 1 - e^{-\lambda\Delta}$, $E_0 = -s$, $E_1 = \lambda h - s$, and $\delta = e^{-r\,\Delta}$. Letting $\Delta \to 0$, we can thus study the impact of vanishing time lags on the results we derived above.

6. Conclusion

We have analyzed a discrete-time experimentation game with two-armed bandits. For publicly observable payoffs, the free-rider problem is so severe that in any subgame perfect equilibrium, both players together perform at most one experiment more than a single agent would. Our result generalizes the free-rider result obtained in continuous time in Keller, Rady and Cripps (2005) for the set of Markov perfect equilibria to all subgame perfect equilibria.

In the second part of the chapter we have analyze the model with privately observed payoffs. Any subgame perfect equilibrium with publicly observed payoffs can be implemented as a perfect Bayesian equilibrium with truthful communication in the game where payoffs are privately observed. Consequently private information does not harm the players if communication is possible.

Furthermore privately observed payoffs can mitigate the free-rider problem to the point where for sufficiently optimistic prior beliefs, it becomes possible to sustain the socially optimal symmetric pure-strategy profile as a perfect Bayesian equilibrium with communication.

Appendix

Proof of Proposition 3.2. As $p^c \leq p^{sc}$ it follows that below p^c the social planner prefers not experimenting over one or two experiments. We calculate the payoff of performing one experiment and then stopping as

$$(1-\delta)\tfrac{1}{2}E_p + \delta p \pi E_1,$$

and the payoff of performing two experiments and then stopping as

$$(1-\delta)E_p + \delta p[1-(1-\pi)^2]E_1\,.$$

Subtracting the latter payoff from the former, we get

$$-(1-\delta)\tfrac{1}{2}E_p + \delta p \pi (1-\pi) E_1,$$

which is negative above $\tilde{p}^c$ and positive below. Because it is suboptimal to perform more than two experiments below $\tilde{p}^c$, it follows that between p^c and $\tilde{p}^c$, the planner prefers one experiment to two experiments.

Next, we establish that above $\tilde{p}^c$, the planner prefers the sequence 2–1 (two experiments this round followed by one experiment next round) to the sequence 1–2. Suppose it is optimal for the social planner to let one player experiment in one round, two players in the subsequent round. We will prove that the planner can make himself strictly better off by first letting two players experiment, then one player, and afterwards using the same strategy as before. In the two rounds under consideration, the expected payoff of the social planner from 2–1 is

$$(1-\delta)\left\{E_p + \delta\left(\tfrac{1}{2}E_p + p[1-(1-\pi)^2]\tfrac{1}{2}E_1\right)\right\}\,.$$

The expected payoff from 1–2 is

$$(1-\delta)(\tfrac{1}{2}E_p + \delta E_p)\,.$$

Subtracting the latter payoff from the former, we get

$$(1-\delta)\tfrac{1}{2}\left\{\left(E_p+\delta p[1-(1-\pi)^2]E_1\right)-\delta E_p\right\}.$$

The part in parentheses is positive above p^{sc} by definition, and $-\delta E_p$ is positive above p^{sc} because $p^{sc}\leq p^m$, so the sequence 1–2 will never be used by the social planner above p^{sc}.

This means that if the planner ever switched to 1 above $\tilde{p}^c$, he would have to continue with 1 until p^c is reached. In a last step, we rule this out by showing that the planner engages in at most one round of experimentation by a single agent before stopping.

If the planner finds it optimal at some belief p to engage in two rounds of experimentation by a single agent and then stop, his expected discounted payoff is

$$\begin{aligned}&(1-\delta)\tfrac{1}{2}E_p+\delta p\pi E_1\\ &\quad+\delta(1-p\pi)\left\{(1-\delta)\tfrac{1}{2}\left[\frac{p(1-\pi)}{1-p\pi}E_1+\frac{1-p}{1-p\pi}E_0\right]+\delta\frac{p(1-\pi)}{1-p\pi}\pi E_1\right\},\end{aligned}$$

which must be non-negative. The expression in braces must be non-negative as well or else it would not be optimal to perform one final experiment after a failure. The expected discounted payoff from performing two experiments at once and then stopping is

$$(1-\delta)E_p+\delta p[1-(1-\pi)^2]E_1.$$

Subtracting the former payoff from the latter, we obtain

$$\begin{aligned}&(1-\delta)\tfrac{1}{2}\left[p(1-\pi)E_1+(1-p)E_0\right]+(1-\delta)\tfrac{1}{2}p\pi E_1+\delta p(1-\pi)\pi E_1\\ &\quad-\delta\left\{(1-\delta)\tfrac{1}{2}\left[p(1-\pi)E_1+(1-p)E_0\right]+\delta p(1-\pi)\pi E_1\right\}\\ =\;&(1-\delta)\left\{(1-\delta)\tfrac{1}{2}\left[p(1-\pi)E_1+(1-p)E_0\right]+\delta p(1-\pi)\pi E_1\right\}+(1-\delta)\tfrac{1}{2}p\pi E_1,\end{aligned}$$

which is positive since the term in braces is nonnegative and $E_1>0$. □

Proof of Lemma 4.1. Suppose we are in a mixed strategy symmetric Markov perfect equilibrium at the lowest belief p in which the agents experiment with strictly positive probability $q\in(0,1]$ than

$$\begin{aligned}\delta\pi qpE_1 &\leq (1-\delta)E_p+\delta pE_1\left(q\left(1-(1-\pi)^2\right)+(1-q)\pi\right)\\ \Leftrightarrow\; 0 &\leq (1-\delta)E_p+\delta pE_1\left(\pi\left(1-q\pi\right)\right)\\ &< (1-\delta)E_p+\delta\pi pE_1\,.\end{aligned}$$

Consequentially in any MPE the agent experiment only with positive probability for beliefs $p\geq p^a$. As $p\leq\frac{1}{2}$ it follows that $\tilde{p}^c\leq p^a$ and thus that the social planner prefers two experiments to one if he stops experimenting afterwards for all beliefs $p\geq p^a$. Hence the socially optimal strategy if we restrict the social planner to strategies that stop experimenting above the belief p^a is pure and prescribes both agents to always experiment. Consequently there exists a pure symmetric strategy that gives both agents higher payoffs than they get in any symmetric MPE.

CHAPTER 5

On Prospect Theory in the Dynamic Context

1. Introduction

While expected utility theory (EUT, Bernoulli (1738/1954), von Neumann-Morgenstern (1944)) is the leading normative theory of decision making under risk, cumulative prospect theory (CPT, Kahneman and Tversky (1979), Tversky and Kahneman (1992)) is the most prominent positive theory. EUT is well-studied in both static and dynamic settings, ranging from game theory over investment problems to institutional economics. In contrast, for CPT most research so far has focused on the static case. In this chapter, we investigate CPT's predictions in the dynamic context and point out fundamental properties of CPT. We illustrate immediate consequences for typical dynamic decision problems such as irreversible investment, casino gambling, or the disposition effect.

Usually, CPT is characterized by four features: First, outcomes are evaluated by a utility function relative to some reference point which separates all outcomes into gains and losses. Second, utility is S-shaped, i.e., convex for losses and concave for gains. Third, probabilities are distorted by inverse-S-shaped probability weighting functions (one for gains and one for losses). Therefore, probabilities close to zero or one are overweighted while moderate probabilities are underweighted. Fourth, losses loom larger than gains, which is referred to as loss aversion.

A small amount of probability weighting is the single driving source this chapter's results. "Small" is relative to the amount of loss aversion and will be precisely defined. Notably, the curvatures of the value and weighting functions are immaterial to our results, and so is the choice of the reference point. If some likelihood insensitivity is considered to be a fundamental element of prospect theory, then so are the results in this chapter.

Our dynamic results can be traced back to a seemingly innocuous result that we call *skewness preference in the small*. At any wealth level, a CPT agent wants to take a sufficiently right-skewed binary risk which is arbitrarily small, even if it has negative expectation. We call such a risk *attractive* to the CPT agent. Therefore, a CPT agent can always be lured into gambling by offering an attractive risk. We show that such a risk *may* be small. However, depending on the value function, the risk may in fact be quite large. For the original parametrization of Tversky and Kahneman (1992) and most empirical estimates we show that there exists an attractive risk of *any* size.

A theory-free definition of risk aversion (risk seeking) at wealth level x is that any zero-mean risk is unattractive (attractive) to the agent. Therefore, skewness preference in the small implies that a CPT agent is not risk averse at any wealth level, and a symmetric result says that CPT is also risk seeking nowhere. We are not aware of a formal proof of this result, not even for the original version of CPT by Tversky and Kahneman (1992). In particular, our result implies that a small amount probability weighting eventually dominates any curvature effects on risk aversion that concave and convex parts of the value function may have: CPT does never imply risk aversion over gains and risk seeking over losses.

In a seminal paper, Barberis (2012) has revealed how the probability weighting component of prospect theory induces a time inconsistency. He provides the general intuition, and illustrates the mechanics for when gambling a 50-50 bet up to five periods. Barberis explains why naive agents (who are unaware of their time inconsistency) typically plan to follow a stop-loss strategy when entering the casino, but end up playing a gain-exit strategy. Such behavior reminds of the disposition effect pointed out by Shefrin and Statman (1985): individual investors are more inclined to sell stocks that have gained in value (winners) rather than stocks that have declined in value (losers). We will show that in more general settings, i.e., for less specific stochastic processes, prospect theory does not predict such behavior for naive agents.

In Xu and Zhou (forthcoming) the optimal strategy of sophisticated agents who have acess to a commitment device is derived in the context of optimal stopping in continuous time. Due to the availability of commitment there is no issue of time inconsistency in this context.

In contrast to this our main result is a Theorem stating that a naive CPT agent will never stop a stochastic process representing his wealth. This holds for our general version of CPT, and for a large class of stochastic processes including geometric and arithmetic Brownian motion. In particular, the process may have arbitrary (positive or negative) drift. By planning to follow a proper stopping strategy (which is stop-loss) the agent can always implement a binary gamble which is attractive because of skewness preference in the small. Intuitively, at any point in time the agent thinks "If I loose just a little bit more, I will stop. And if I gain, I will continue." But once a loss *or* a gain has occurred, a new attractive stopping strategy will come to his mind and thus he will continue to gamble.

Differently from Barberis' five-period example with 50-50 bets, our agent will also continue to gamble if he makes a gain. In Barberis (2012), the combination of symmetric gambles and finite (very short) time horizon implies that the "casino dries out of skewness." For example, in the last period there is only a 50-50 gamble available in the casino so that skewness preference in the small does not apply. The availability of sufficiently skewed gambles is the crucial input to our limit result. A continuous, infinite time horizon setup is sufficient for it. However, we show that the result is robust to finite and/ or discrete time spaces, as long as the stochastic process allows for a rich set of possible gambling strategies. In complete markets, for example, our result will hold irrespective of the time space.

In Section 2 we will define our general version of CPT. In Section 3 we present our static result that CPT implies skewness preference in the small and point out several implications. Section 4 presents the main result that a naive agent never stops a stochastic process that represents his wealth. Section 5 discusses the implications for CPT models of casino gambling, optimal time to invest, and the disposition effect. Section 6 discusses the robustness of our result towards discrete and finite time spaces. Section 7 concludes.

2. Prospect Theory Preferences

We consider an agent with CPT preferences over random variables $X \in \mathrm{R}$. A CPT agent evaluates the risk X as

$$CPT(X) = \int_{\mathrm{R}_+} w^+(\mathrm{P}(U(X) > y))\mathrm{d}y - \int_{\mathrm{R}_-} w^-(\mathrm{P}(U(X) < y))\mathrm{d}y \tag{2.1}$$

with non-decreasing weighting functions $w^-, w^+ : [0,1] \to [0,1]$ with $w^+(0) = w^-(0) = 0$ and $w^+(1) = w^-(1) = 1$ and a value function $U : \mathrm{R} \to \mathrm{R}$ that satisfy the assumptions explained in the following.

Assumption 5 (Value function). *The value function is absolutely continuous and strictly increasing. Further,* $\lambda = \sup_{x \in \mathrm{R}} \frac{\partial_- U(x)}{\partial_+ U(x)} < \infty$ *exists, where* $\partial_- U(x)$ *and* $\partial_+ U(x)$ *denote the left and right derivative of* U*, respectively. W.l.o.g.* $U(0) = 0$.

Under CPT, preferences are defined on changes relative to some reference point, which is denoted by $r \in \mathrm{R}$. Typical choices for r are the status quo or some other benchmark. For example, when investing in a risky asset, r could be the return of a risk-free investment. Realizations x of X with $x < r$ are referred to as losses, and realizations $x \geq r$ are called gains. Our results hold for any $r \in \mathrm{R}$. In other words, the choice of r is immaterial to our findings. In many specifications of prospect theory, the additional assumption is made that U is differentiable everywhere, except at the reference point such that $\lambda = \frac{\partial_- U(r)}{\partial_+ U(r)}$. It is further assumed that $\lambda > 1$ and that the reflection property

$$U(x) = \begin{cases} u(x - r), & \text{if } x \geq r \\ -\lambda u\left(-(x - r)\right), & \text{if } x < r \end{cases} \tag{2.2}$$

holds for some function u. $\lambda > 1$ then implies that losses loom larger than gains to the CPT agent; see Köbberling and Wakker (2005) for an analysis of the loss aversion index $\frac{\partial_- U(r)}{\partial_+ U(r)}$. We allow for non-differentiable utility because it allows to model preference over assets with non-differentiable payoffs such as option contracts. The original choice for u by Tversky and Kahneman (1992) was power utility. Since several caveats have been pointed out for the power utility parametrization, exponential utility has become another popular choice (de Georgi and Hens (2006)).[1] The final important feature of CPT is that the agent distorts the decumulative probabilities associated with gains and the cumulative probabilities associated with losses by means of respective weighting functions w^+ and w^-.

Assumption 6 (Likelihood-Insensitive Weighting Functions). *The weighting functions* w^+ *and* w^- *satisfy*

(1) $\limsup_{p \to 0} \frac{w^+(p)}{p} > \lambda$.
(2) $\limsup_{p \to 0} \frac{1 - w^-(1-p)}{p} > \lambda$.

Note that if w^+ and w^- are differentiable at 0 and 1, then the conditions in Assumption 6 simplify to ${w^+}'(0) > \lambda$ and ${w^-}'(1) > \lambda$. If these derivatives do not exist because they approach infinity, the limit superior is infinite which is also consistent with Assumption 6.
We do not require the weighting functions to be inverse-S-shaped, i.e., to start out concave and to turn convex at some point. Our assumption is weaker and much more in the spirit of Wakker and Tversky (1995), who define a connected likelihood insensitivity region bounded away from both 0 and 1. Actually, here we only require overweighting of small probabilities when associated with high outcomes; see Wakker (2010, pp. 222-233) for a comprehensive discussion of likelihood insensitivity and inverse-S-shape.

Observation 1. *Assumption 6 is satisfied by the commonly used weighting functions of Kahneman and Tversky (1979), Goldstein and Einhorn (1987), Prelec (1998), and the neo-additive weighting function. Whether the weighting function of Rieger and Wang (2006) suffices Assumption 6 depends on parameter estimates.*

[1] Exponential utility satisfies Assumption 5. Power utility does not suffice Assumption 5, because both partial derivatives are infinite at the reference point and thus the Kbberling-Wakker index of loss aversion is not well defined. We will treat power utility separately in Subsection 3.4. It will be seen that, for this case, even stronger results may be obtained than those based solely on Assumption 5.

Proof. See appendix. □

Rieger and Wang (2006) farsightedly proposed a weighting function with finite derivatives at 0 and 1, and showed that this ensures non-occurrence of the St. Petersburg paradox under CPT. However, it remains to be investigated whether other desirable predictions of CPT will be maintained for such weighting functions. In any case, while existing, finite derivatives at 0 and 1 are a cure to the St. Petersburg paradox, such a version of CPT may still succumb to our results.[2]

Moreover, many of our results allow to relax Assumption 6 so that they apply to any weighting function with derivatives at 0 and 1 strictly larger than one, which corresponds to a minimal departure from EUT's linear processing of probabilities.

If not noted otherwise, in the following our sole assumptions on the CPT preference functional (2.1) are Assumptions 5 and Assumption 6. Our point is that a small amount of probability weighting *alone* is sufficient for a fundamental property of CPT in the static case, which in turn has drastic implications for CPT in a dynamic setting.

Many of our results use binary lotteries $L \equiv L(p, b, a)$ that yield outcome b with probability $p \in (0, 1)$, and $a < b$ otherwise. Thus it is convenient to note that the CPT preference functional (2.1) evaluates binary risks as

$$CPT(L) = \begin{cases} w^+(p)U(b) + (1 - w^+(p))U(a), & \text{if } r \leq a \\ w^-(1-p)U(a) + w^+(p)U(b), & \text{if } a < r \leq b \\ (1 - w^-(1-p))U(b) + w^-(1-p)U(a) & \text{if } b < r. \end{cases} \tag{2.3}$$

3. Static Results

3.1. Prospect Theory's Skewness Preference in the Small. This chapter starts out with a seemingly innocuous result on prospect theory preferences and small, skewed risks. We say that a risk is *attractive* or that an agent *wants to take a risk* if the CPT utility of current wealth plus the risk is *strictly* higher than the CPT utility of current wealth.

THEOREM 3.1 (Prospect Theory's Skewness Preference in the Small). *For every wealth level x and every $\epsilon > 0$ there exists an attractive zero-mean binary lottery $L \equiv L(p, b, a)$ with a, $b \in (-\epsilon, +\epsilon)$, i.e., L may be arbitrarily small.*

Proof. We split the proof into three cases $x > r$, $x < r$, and $x = r$. We prove the equivalent result that for all $x \in \mathrm{R}$ and every $\epsilon > 0$ there exists a binary lottery $L \equiv L(p, b, a)$ with mean x and a, $b \in (x - \epsilon, x + \epsilon)$ such that $CPT(L) > CPT(x)$. L having mean x yields

$$x = (1 - p)a + pb \iff p = \frac{x - a}{b - a}\,.$$

Proof of case 1 ($x > r$). Choose $a > r$ such that both a and b are gains. Then lottery L gives the agent a utility of $CPT(L) = w^+(p)U(b) + (1 - w^+(p))U(a)$. Therefore, the agent prefers L over x if there exist $a < x$

[2] Finite derivatives may be consistent with our Assumption 6. We only require that the derivatives at 0 and 1 are larger than λ. Unfortunately, we are not aware of any parameter estimates of the Rieger and Wang (2006) weighting function.

and $b > x$ such that

$$0 < \left(1 - w^+\left(\frac{x-a}{b-a}\right)\right) U(a) + w^+\left(\frac{x-a}{b-a}\right) U(b) - U(x)$$
$$= (U(b) - U(a)) \left(w^+\left(\frac{x-a}{b-a}\right) - \frac{U(x)-U(a)}{U(b)-U(a)}\right)$$
$$(3.1) \qquad = \underbrace{p\,(U(b)-U(a))}_{\geq 0} \left(\frac{w^+(p)}{p} - \frac{\frac{U(x)-U(a)}{x-a}}{\frac{U(b)-U(a)}{b-a}}\right).$$

Consider sequences $(a_n, b_n)_{n \in \mathrm{N}}$, with $a_n = x - \frac{p}{n}$ and $b_n = x + \frac{1-p}{n}$. Note that by construction

$$\frac{U(b_n)-U(a_n)}{b_n-a_n} = \frac{U(b_n)-U(x)}{b_n-x}\frac{b_n-x}{b_n-a_n} + \frac{U(x)-U(a_n)}{x-a_n}\frac{x-a_n}{b_n-a_n}$$
$$= \frac{U(b_n)-U(x)}{b_n-x}(1-p) + \frac{U(x)-U(a_n)}{x-a_n}p.$$

Therefore, according to equation (3.1), the agent prefers lottery L over x if

$$(3.2) \qquad 0 < \frac{w^+(p)}{p} - \frac{\frac{U(x)-U(a_n)}{x-a_n}}{\frac{U(b_n)-U(x)}{b_n-x}(1-p) + \frac{U(x)-U(a_n)}{x-a_n}p} =: \xi_n(p).$$

First, suppose that $\frac{w^+(p)-w^+(0)}{p-0} = \frac{w^+(p)}{p} \to \infty$ for $p \to 0$. Because the subtracted part in equation (3.2) is bounded for every n, equation(3.2) is fulfilled for sufficiently small p. Moreover, since $(a_n) \nearrow x$ and $(b_n) \searrow x$ we have $a_n, b_n \in (x-\epsilon, x+\epsilon)$ for n sufficiently large. Second, suppose $\lim_{p \to 0} \frac{w(p)}{p} = w'(0) < \infty$ exists. Since $(a_n) \nearrow x$ and $(b_n) \searrow x$, for all $p \in (0,1)$,

$$(3.3) \qquad \lim_{n\to\infty} \xi_n(p) = \frac{w^+(p)}{p} - \frac{\partial_- U(x)}{\partial_+ U(x)(1-p) + \partial_- U(x) p} =: \xi(p) \text{ exists.}$$

By Assumption 6,

$$(3.4) \qquad 0 < w^{+'}(0) - \frac{\partial_- U(x)}{\partial_+ U(x)} = \lim_{p \to 0} \xi(p).$$

Since $\xi(p)$ is continuous (Assumption 5) there exists $\tilde{p} \in (0,1)$ such that also $\xi(\tilde{p}) > 0$, i.e., $\lim_{n\to\infty} \xi_n(\tilde{p}) > 0$. Therefore, equation (3.2), and also $a_n, b_n \in (x-\epsilon, x+\epsilon)$, is fulfilled for $n = n(\tilde{p}, \epsilon)$ sufficiently large, i.e., $L(\tilde{p}, b_{n(\tilde{p},\epsilon)}(\tilde{p}), a_{n(\tilde{p},\epsilon)}(\tilde{p}))$ is preferred over x for sure. The proofs for $x < r$ and $x = r$ are given in the appendix. □

Corollary 3.1 (Unfair Attractive Gambles). *For every wealth level $x \in \mathrm{R}$ there exists an attractive, arbitrarily small binary lottery with negative mean.*

Proof. The claim follows from continuity of the CPT preference functional (Assumption 5). □

It is straightforward to formulate a local version of Theorem 3.1.

Corollary 3.2 (Local Result). *At some given wealth level x there exists an attractive, arbitrarily small zero-mean binary lottery even if Assumption 6 is relaxed by replacing $\lambda := \sup_{x \in \mathrm{R}} \frac{\partial_- U(x)}{\partial_+ U(x)}$ with $\frac{\partial_- U(x)}{\partial_+ U(x)}$. If U is differentiable at x, then Assumption 6 may be further relaxed by replacing λ with* 1.

Proof. The claim is evident from the proof of Theorem 3.1, and since $\frac{\partial_- U(x)}{\partial_+ U(x)} = 1$ if U is differentiable at x. □

The intuition of the proof of Theorem 3.1 is that CPT implies skewness preference. Ebert (2011) illustrates

that, for binary lotteries, skewness—according to both the tails and moments definitions—is exhaustively captured in the probability parameter. Therefore, we can interpret the proof of Theorem 3.1 as the construction of a sufficiently right-skewed fair lottery. By letting p go to zero, the binary lottery becomes more and more right-skewed. At some point the lottery is so much skewed that a CPT agent wants to take it.
Skewness preference has been of major interest in the recent economics and finance literature. Numerous empirical and experimental papers find support for skewness preference (e.g., Kraus and Litzenberger (1976) and Boyer, Mitton, and Vorkink (2010) for asset returns, Golec and Tamarkin (1998) for horse-race bets, and Ebert and Wiesen (2011) in a laboratory experiment). Moreover, various economic behaviors and financial phenomena can be explained by skewness preference, e.g., casino gambling (Barberis (2012)), underdiversification in stock portfolios (Barberis and Huang (2008)), or positive expected first-day returns accompanied by negative medium-run expected returns for initial public offerings (Green and Hwang (forthcoming)). The famous coexistence of lottery and insurance demand under CPT stems from skewness preference. In many of these situations, prospect theory may do such a good job in explaining behavior because, through its probability weighting component, it implies skewness preference. Other papers have argued like this. To best of our knowledge, however, Theorem 3.1 is the first rigorous result that relates CPT to skewness preference.

3.2. Prospect Theory Agents are Risk-Averse and Risk-Seeking Nowhere. A decision-theoretic implication of Theorem 3.1, which is of independent interest, is on how risk aversion manifests in CPT. A theory-free definition of risk aversion (risk-seeking) at wealth level x is that any zero-mean risk is unattractive (attractive) to the agent. Numerous qualitative statements on how the curvature of the value function affects risk aversion in CPT can be found in the literature, but formal results are hard to find. A notable exception is the paper of Schmidt and Zank (2008) who characterize the curvatures of value and weighting function under which CPT exhibits strong risk aversion globally. Kahneman and Tversky (1979, p. 285) themselves noted that "[Our previous analysis] restricts risk seeking in the domain of gains and risk aversion in the domain of losses to small probabilities [...]" Here is a stronger result derived from weaker assumptions, i.e., from some probability weighting only.

Corollary 3.3. *At any wealth level a CPT agent is not risk-averse.*

Proof. The statement is a direct consequence of Corollary 3.2: At any wealth level, there exists an attractive risk. □

As in the Corollary 3.2, Assumptions 5 and 6 may be relaxed to obtain a tighter result locally. It is straightforward to formulate analogous versions of Theorem 3.1 and its corollaries on the unattractiveness of left-skewed gambles and risk-seeking under CPT. To this means, we have to assume that probabilities associated with bad outcomes are overweighted. These assumptions[3] are complementary to our Assumptions 5 and 6 and likewise fulfilled by the specifications in Observation 1. Then we have that, everywhere, there exists an arbitrarily small, left-skewed binary risk which is unattractive, and that a CPT agent is everywhere not risk-seeking. We find it striking that just some probability weighting "dominates" the impact of the curvature of the value function. In particular, our result illustrates that the intuition that the S-shaped value function of prospect theory implies risk aversion for gains and risk-seeking for losses is misleading.

[3]Specifically, $\tilde{\lambda} = \inf_{x \in \mathrm{R}} \frac{\partial_- U(x)}{\partial_+ U(x)} < \infty$ must exist and 1. $\limsup_{p \to 0} \frac{w^-(p)}{p} > \lambda$ and 2. $\limsup_{p \to 0} \frac{1 - w^+(1-p)}{p} > \lambda$. Evidently, these properties are also necessary for the famous inverse-S-shape, and consistent with the likelihood-insensitivity definition of Wakker and Tversky (1995). Under these assumptions one can construct an unattractive, left-skewed binary risk. The proof is similar to that of of Theorem 3.1, with the main difference that one must let $p \to 1$ rather than $p \to 0$ to generate left-skew.

3.3. Large Risks. Next, note that we may construct an attractive risk which is arbitrarily small. However, it must not be misunderstood that the attractive risk has to be small. Striking results on large attractive gambles have been presented by Rieger and Wang (2006) who investigate the occurrence of the St. Petersburg Paradox under CPT, and by Azevedo and Gottlieb (forthcoming) who show that risk-neutral firms can extract unbounded profits from CPT consumers. These authors construct attractive gambles that involve arbitrarily large payoffs, and thus it is intuitive that their results also require assumptions on the value function. Azevedo and Gottlieb (forthcoming) point out that for the power value function and for any attractive binary gamble L the multiple cL ($c > 1$) is also attractive. In combination with our result this then implies that there exist attractive gambles of any size. This we will make precise in the next section.

3.4. The Case of a S-Shaped Power Value Function. In this section we consider a power value function which suffices the reflection property, equation (2.2).

ASSUMPTION 7 (S-Shaped Power Value Function). *The value function is given by*

$$U(x) = \begin{cases} (x-r)^{\alpha}, & \text{if } x \geq r \\ -\hat{\lambda}\left(-(x-r)\right)^{\alpha}, & \text{if } x < r \end{cases} \tag{3.5}$$

with $\alpha \in (0,1)$ *and* $\hat{\lambda} > 1$.

For this very choice, the Kbberling-Wakker index of loss aversion $\frac{\partial_{-}U(r)}{\partial_{+}U(r)}$ is not well-defined (in particular, it is not equal to $\hat{\lambda}$) because the power function has infinite derivative at 0. Therefore, Assumption 5 is not fulfilled, and thus Theorem 3.1 does not apply. However, we can state a similar result under a slightly different assumption on the weighting functions.

ASSUMPTION 8. *The weighting functions* w^+ *and* w^- *satisfy*

(1) $\limsup_{p\to 0} \frac{w^+(p)}{p^{\alpha}} > \hat{\lambda}$.
(2) $\limsup_{p\to 0} \frac{1-w^-(1-p)}{p} > 1$.

Note that condition 1 in Assumption 8 is stronger than condition 1 of Assumption 6. However, it is weaker than the assumption in Azevedo and Gottlieb's Proposition 1 when applied to power utility, which requires that the limit is infinite. This is the case for the weighting functions of Tversky and Kahneman (1992) and Goldstein and Einhorn (1987) under parameter restrictions that are typically fulfilled according to most empirical studies; see Azevedo and Gottlieb (forthcoming) for an elaboration. For the weighting function of Rieger and Wang (2006) our condition is not fulfilled.[4] For the weighting function of Prelec (1998) Azevedo and Gottlieb's assumption is always true.

THEOREM 3.2 (Skewness Preference in the Small for the S-Shaped Power Value Function). *Assume Assumptions 7 and 8 instead of Assumptions 5 and 6. For every wealth level x and every $\epsilon > 0$ there exists an attractive, zero-mean binary lottery $L \equiv L(p,b,a)$ with $a, b \in (-\epsilon, \epsilon)$, i.e., L may be arbitrarily small.*

Proof. Since U is differentiable everywhere except at r, the result for $x \neq r$ follows from Corollary 3.2. The case $x = r$ is proven in the appendix. □

Power utility is differentiable everywhere except at the reference point. Therefore, note that Corollary 3.2,

[4]Note that there is a typo in Azevedo's and Gottlieb's paper, which says that their (stronger) condition (1) is fulfilled. Actually, it is their condition (2) which is met unless the power utility parameter is equal to one. In the latter case, none of their conditions is fulfilled, but our Assumption 5 is.

which assumes just minimal probability weighting, also applies to power utility whenever we are not at the reference point. Therefore, we need Assumption 8 exclusively to cover gambling at the reference point. Finally, let us combine Theorem 3.2 with the result of Azevedo and Gottlieb (forthcoming).

Corollary 3.4 (Skewness Preference in the Small and in the Large for the S-Shaped Power Value Function). *Assume Assumptions 7 and 8 instead of Assumptions 5 and 6. Then there exists an attractive, zero-mean binary lottery of arbitrary size.*

Proof. According to Theorem 3.2 there exists an attractive, arbitrarily small binary risk. According to Azevedo and Gottlieb (forthcoming) it can be scaled up to any size. □

4. Prospect Theory in the Dynamic Context

In this section we investigate the consequences of skewness preference in the small in the dynamic context. Assume that Assumptions 5 and 6 are fulfilled. Alternatively, assume the power utility case, i.e., Assumptions 7 and 8. We now define a stochastic process $(X_t)_{t\in \mathrm{R}_+}$ that could reflect the cumulated returns of an investment project, or the price development of an asset traded in the stock market. It could likewise model an agent's wealth when gambling in a casino. Let $(W_t)_{t\in \mathrm{R}_+}$ be a Brownian motion and $(X_t)_{t\in \mathrm{R}_+}$ a Markov diffusion that satisfies

$$\mathrm{d}X_t = \mu(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}W_t$$

where we assume $\mu : \mathrm{R} \to \mathrm{R}$ and $\sigma : \mathrm{R} \to (0, \infty)$ such that there exists a unique solution with continuous paths.[5] Note that the most frequently considered processes, arithmetic and geometric Brownian motion, are covered by this definition. We denote by $\mathcal{S}$ the set of all stopping times such that the agent bases his stopping decision only on his past observations. Formally, all $\tau \in \mathcal{S}$ are adapted to the natural filtration $(\mathcal{F}_t)_{t\in \mathrm{R}_+}$ of the process $(X_t)_{t\in \mathrm{R}_+}$. At every point in time the naive prospect theory agent faces the problem of finding a stopping time $\tau \in \mathcal{S}$ that maximizes his prospect value $CPT(X_\tau, \mathcal{F}_t)$ given his information $\mathcal{F}_t$ at time t where

$$CPT(X_\tau, \mathcal{F}_t) = \int_{\mathrm{R}_+} w^+(\mathrm{P}(u(X_\tau - r) > t \,|\, \mathcal{F}_t))\mathrm{d}t \\ - \int_{\mathrm{R}_-} w^-(\mathrm{P}(u(X_\tau - r) < t \,|\, \mathcal{F}_t))\mathrm{d}t\,.$$

The agent stops at time t if and only if his prospect value $CPT(X_\tau, \mathcal{F}_t)$ of any stopping time $\tau \in \mathcal{S}$ is less than or equal to what he gets if he stops immediately, which would be $CPT(X_t)$.

The probability weighting of prospect theory induces a time inconsistency. This has been pointed out by Barberis (2012), who very well illustrates the mechanics along the lines of a casino gambling example. While the agent plans to follow a certain strategy τ at the beginning, she might decide for another one once her wealth has changed. A naive agent does not anticipate that later she might deviate from her initial plan. Therefore, at every point in time, the agent looks for a strategy τ that brings her higher CPT utility than stopping immediately. If such a strategy exists, she continues to gamble—irrespective of her initial plan. In the following, we always consider such a naive agent.

[5] $\mu : \mathrm{R} \to \mathrm{R}$ and $\sigma : \mathrm{R} \to (0, \infty)$ are locally Lipschitz continuous Borel functions with linear growth, i.e., there exists a $K > 0$ such that $|\mu(x)|^2 + |\sigma(x)|^2 \leq K(1 + |x|^2)$.

Because no analytical solution is available, Barberis (2012) investigates planned and actual behavior by computing the CPT values of all possible gambling strategies that can be generated by a finite 50-50 binomial tree, for more than 8000 parameter combinations of the CPT parametrization of Tversky and Kahneman (1992). The exercise could then be repeated for other CPT parametrizations and for other stochastic processes. We now first present a general solution for the case of an infinite time horizon which is independent of the CPT specification, of the curvatures of the value and weighting functions, and of the reference point (which my change over time). It holds for the general class of stochastic processes specified above, in particular for processes with zero, positive, and negative drift. In Section 6 we discuss discrete and finite time.

Theorem 4.1. *The naive agent never stops.*

The intuition of the proof (given below) is to construct a stop-loss strategy, i.e., one where the agent plans to stop if the process falls a little bit (arrives at a) and plans to continue until it has risen significantly (up to b). This results in a right-skewed binary risk which the agent prefers to stopping immediately due to Theorem 3.1 or, in the case of a power value function, due to Theorem 3.2.

Proof of Theorem 4.1. Suppose the agent arrives at wealth x at time t, i.e., $X_t = x$. The agent can stop and get a utility of $CPT(x)$, or she may continue to gamble. She continues to gamble if there exists a gambling strategy $\tau \in \mathcal{S}$, i.e., a stopping time such that $CPT(x) < CPT(X_\tau)$. We consider strategies $\tau_{a,b}$ with two absorbing endpoints $a < x < b$ which stop if the process $(X_t)_{t\in \mathrm{R}_+}$ leaves the interval (a, b), i.e.,

$$\tau_{a,b} = \inf\{s \geq t : X_s \notin (a, b)\}\,.$$

Denote with $p = \mathrm{P}(X_{\tau_{a,b}} = b)$ the probability that with strategy $\tau_{a,b}$ the agent will stop at b. Note that strategy $\tau_{a,b}$ results in a binary lottery for the agent. We first prove that the agent never stops if X_t is a martingale. For every stopping time $\tau_{a,b}$ consider the sequence of bounded stopping times $\min\{\tau_{a,b}, n\}$ for $n \in \mathrm{N}$. By Doobs optional stopping theorem (Revuz and Yor, p. 70), $\mathrm{E}(X_{\min\{\tau_{a,b},n\}}) = X_t = x$. By the theorem of dominated convergence it follows that

$$\mathrm{E}\left(X_{\tau_{a,b}}\right) = \mathrm{E}\left(\lim_{n\to\infty} X_{\min\{\tau_{a,b},n\}}\right) = \lim_{n\to\infty} \mathrm{E}\left(X_{\min\{\tau_{a,b},n\}}\right) = x\,.$$

Hence, $X_{\tau_{a,b}}$ implements the binary lottery $L(p, a, b)$ with expectation x. From Theorem 3.1 (Theorem 3.2) it follows that there exist $a, b \in \mathcal{I}$ such that the agent prefers the binary lottery induced by the strategy $\tau_{a,b}$ over the certain outcome x.

In the last step we prove that the naive agent never stops even if $(X_t)_{t\in\mathrm{R}_+}$ is not a martingale. Define the strictly increasing scale function $S : \mathrm{R} \to \mathrm{R}$ by

$$S(x) = \int_0^x \exp\left(-\int_0^y \frac{2\mu(z)}{\sigma^2(z)}\mathrm{d}z\right)\mathrm{d}y\,.$$

Define a new process $\hat{X}_t = S(X_t)$ and a new value function $\hat{U}(x) = (U \circ S^{-1})(x)$. Note that the loss aversion index of the value function $\hat{U}$ equals the loss aversion index of U because

$$\frac{\partial_-\hat{U}(x)}{\partial_+\hat{U}(x)} = \frac{\partial_-\hat{U}(x)S'(x)}{\partial_+\hat{U}(x)S'(x)} = \frac{\partial_-U(x)}{\partial_+U(x)}\,.$$

A CPT agent with the value function $\hat{U}$ facing the process $(\hat{X}_t)_{t\in\mathrm{R}_+}$ evaluates all stopping times exactly as a CPT agent with value function U who faces $(X_t)_{t\in\mathrm{R}_+}$. The process $\hat{X}_t = S(X_t)$ satisfies (Revuz and Yor

(1999, p. 303 ff))

$$\mathrm{P}\left(\hat{X}_{\tau_{a,b}} = S(b)\right) = \mathrm{P}\left(X_{\tau_{a,b}} = b\right) = \frac{S(x) - S(a)}{S(b) - S(a)},$$

and hence it follows from the argument for martingales that the agent never stops. □

5. Applications

5.1. Casino Gambling. Our first example is the continuous, infinite time horizon analogue to the discrete, finite time setting of Barberis (2012). Let $(X_t)_{t \in \mathrm{R}_+}$ be a Brownian motion with negative drift $\mu(x) = \mu < 0$ and constant variance $\sigma(x) = \sigma > 0$, i.e.,

$$\mathrm{d}X_t = \mu \mathrm{d}t + \sigma \mathrm{d}W_t \,.$$

Due to the negative drift the agent loses money in expectation if he does not stop. Further assume that the process absorbs at zero since then the agent goes bankrupt. From Theorem 4.1 it follows that the naive agent gambles until the bitter end, i.e., he will not stop gambling unless he is forced to due to bankruptcy. From standard results in probability theory we know that this will happen almost surely, i.e., $\mathrm{P}(X_\tau = 0) = 1$. We will compare this result to that of Barberis (2012) in Section 6.

5.2. Exercising an American Option. Let $(X_t)_{t \in \mathrm{R}_+}$ be a geometric Brownian motion with drift $\mu \in \mathrm{R}$ and variance $\sigma > 0$, i.e.,

$$\mathrm{d}X_t = X_t(\mu \mathrm{d}t + \sigma \mathrm{d}W_t)$$

The agent holds an American option that pays

$$\pi(X_t) = \max\{e^{-\alpha t}(X_t - K), 0\}$$

if exercised at time t where $\alpha > 0$ denotes the risk-free rate. Here $K \in \mathrm{R}_+$ represents the costs of investment. The American option could be interpreted as an investment opportunity, i.e., a real option (compare Dixit and Pindyck (1994)). We assume $\mu < \alpha$ to ensure that the value of the expected value maximizer is finite. The payoff $\pi(X_t)$ is incorporated into the model by simply replacing the agent's value function $U(\cdot)$ by $\hat{U}(X_t) = U(\pi(X_t))$. Here we benefit from not having assumed differentiability of the value function. The agent is allowed to exercise his option at every point in time $t \geq 0$. From Theorem 4.1 it follows that the agent will never exercise his option, i.e., $\tau = \infty$. As $\lim_{t\to\infty} e^{-\alpha t}(X_t - K) \xrightarrow{P} 0$ the naive prospect theory agent gets a payoff of zero even though he could get a strictly positive payoff by exercising the option immediately whenever $X_0 > K$.

5.3. Prospect Theory Fails to Explain The Disposition Effect. The disposition effect (Shefrin and Statman (1985)) refers to individual investors being more inclined to sell stocks that have gained in value (winners) rather than stocks that have declined in value (losers). Numerous papers have addressed this phenomenon, and some of the most immediate explanations such as transaction costs, tax concerns, or portfolio rebalancing have been formidably ruled out by Odean (1998).

Several papers have investigated whether prospect theory can explain the disposition effect. However, all of them seem to have done so without the consideration of probability weighting (Barberis (2012, footnote 26)). Formal models (without probability weighting) have been put forward just recently by Kyle, Ou-Yang, and Xiong (2006), Kaustia (2010), Barberis and Xiong (2009), and Henderson (forthcoming). The results are mixed. Some find that prospect theory can predict the disposition effect, and others that it cannot, at least not under all relevant circumstances. Barberis (2012) notes that the binomial tree in his paper, which models a casino, may likewise represent the evolution of a stock price over time. Then, naive investors may exhibit a disposition effect, even though they plan to do the opposite of the disposition effect. Our result can be related to the disposition effect in the same spirit.

We have shown that, in general, under probability weighting a naive CPT agent will sell neither losers nor winners at any time. As a consequence, a continuous time model of prospect theory with probability distortion does not predict a disposition effect for naive investors. This is especially striking as Henderson (forthcoming) shows that, in an analogous model without probability distortion, prospect theory can explain the disposition effect.

Note that the continuous time price processes such as geometric Brownian motion that are covered by our setup fit particularly well for financial market models. In any case, in the next section we show that our result also applies to a wide range of continuous or discrete, finite or infinite time horizon processes.

6. Robustness to Discrete and Finite Time Specifications

While it may seem that our results are related to the continuous time setup, they are not. Continuous time ensures that at every point in time the strategy set of the agent is sufficiently rich. To illustrate this point consider a binomial random walk $(X_t)_{t\in\mathrm{N}}$ with jump size one and equal probability for up- and down movements. At every point in time t the agent can choose the stakesize $s_t \in [0,1]$ (as a fraction of his wealth y_t) to bet. The evolution of his wealth is then given by

$$y_{t+1} = y_t + s_t y_t (X_{t+1} - X_t)$$

with initial wealth $y_0 > 0$. The following strategy (of choosing s_t) results in any given fair binary lottery $L(p, b, a)$. Choose s_t maximal such that $y_{t+1} \in [a, b]$, i.e.,

$$\begin{aligned} s_t &= \max\{\tilde{s} \in [0,1] : (1+\tilde{s})y_t \geq a \text{ and } (1+\tilde{s})y_t \leq b\} \\ &= \min\{1 - \frac{a}{y_t}, \frac{b}{y_t} - 1\}\,. \end{aligned}$$

Due to the martingale property it follows from Doob's optional sampling theorem that the probabilities of hitting b and a that are induced by this strategy are fair, i.e., are p and $1-p$, respectively. If $L(p, b, a)$ is attractive according to either Theorem 3.1 or 3.2, then the agent will gamble with this strategy in mind. Since an attractive lottery exists at any wealth level (i.e., at any time t) the agent never stops.

The crucial point of this example is that the *time space* may be discrete if we ensure that the *strategy space* is sufficiently rich. Specifically, a global result like Theorem 3.1 requires that, at any time t, for any state X_t, there is at least one stopping strategy available that results in an attractive gamble. This explains why the gambling behavior documented in Barberis (2012) is different. Barberis considers behavior when gambling a 50-50 bet up to five periods. This *combination* of symmetric gambles and finite (very short) time horizon ensures that the "casino dries out of skewness." That is, at some exogenous point in time, the casino does not allow for gambling strategies any more that result in attractive gambles. The set of possible gambling strategies becomes smaller, or *coarser*, over time.

With this in mind it is immediate that we can also have never stopping with a finite time space. To this means, the casino must be able to offer a sufficiently skewed gamble (which is attractive according to either Theorem 3.1 or 3.2) in a single period, i.e., in the final period. This is just a less subtle way (compared to allowing for an infinite time horizon) to enrich the strategy space. To illustrate this point we give a numerical example in Subsection 8 in the appendix. There we assume the original finite, discrete time setting introduced by Barberis (2012), and simply change the probability of an up-movement in the binomial tree from 1/2 to 1/37. We show that an agent with CPT preferences of Tversky and Kahneman (1992) and the parameter estimates from that paper never stops gambling for any finite or infinite horizon.

If a casino cannot offer one-shot gambles with sufficient skewness, then several periods of gambling might be necessary to construct an attractive gamble. The number of periods will depend both on the maximal skewness of available one-shot gambles, and on the particular CPT specification and parameter choices. Then, we will have an *endgame effect* as in Barberis (2012). The analysis of this effect is extremely insightful to understand the interaction of probability weighting, time-inconsistency, and naivet.

However, our result shows that such effects will vanish if a casino can offer a rich set of gambling strategies. Likewise, endgame effects will vanish in financial markets that offer a variety of products. In complete markets, in particular, we have that a naive CPT agent never stops investing into a risky asset.

These thoughts point to a fundamental challenge for the application of prospect theory in finite time horizon models of gambling or investment. In such models it will be hard to disentangle whether the conclusions stem from the particular process assumed, from the particular prospect theory preference specification—or from other features of the model one might actually be interested in. In particular, the number of periods is important because they influence the richness of gambling strategies. In other words, a tough question is to what extent results are due to an endgame effect, which, according to the results of this chapter, will vanish when enriching the strategy space.

7. Conclusion

We set up a very general version of cumulative prospect theory (CPT) and point out fundamental implications for that theory that stem from probability weighting alone. The reference point and the curvatures of the weighting and value functions are immaterial to our results. We first prove that probability weighting implies skewness preference in the small. At any wealth level, a CPT agent wants to take a sufficiently right-skewed binary risk that is arbitrarily small, even if it has negative expectation. To best of our knowledge, this is first rigorous result that relates CPT to skewness preference. A corollary is that CPT agents are not risk-averse, even if, for example, the value function is concave everywhere. While we prove the existence of small attractive risks, under additional assumptions on the value function we show that attractive risks may, in fact, be quite large. For the power value function of Tversky and Kahneman (1992) we show that for typical parameter estimates there exist attractive risks of arbitrary size.

These static results have consequences for CPT in the dynamic context. We investigate the predictions of probability weighting for a naive agent who is unaware of his time-inconsistency, which is induced by probability weighting. Such a naive agent will never stop a stochastic process that represents his wealth. The implications of this result are very extreme. Naive agents will gamble in a casino until the bitter end, i.e., they will go bankrupt almost surely. They will never exercise an American option, even if it is profitable to do so right from the beginning. CPT does not predict the disposition effect for naive agents. These results are formulated for a continuous, infinite time horizon.

Then we illustrate that the results extend to discrete time, as long as the space of stopping strategies that can be generated from the process is sufficiently rich. Likewise, we may also allow for finite time. If the time space is such that the set of stopping strategies is coarse in the sense that it does not allow to adopt stopping strategies that result in small, skewed gambles, then our result will not apply. However, casinos and even more financial markets allow for a very rich set of gambling and investment strategies. In complete markets, in particular, our never stopping result always applies.

In finite time there will be an endgame effect, which will be particularly pronounced for coarse, discrete processes. Then the set of available stopping strategies decreases with every time step. This may lead to interesting observations on the planned and actual behavior of naive agents as has been illustrated by Barberis (2012). Generally, however, it is hard to disentangle whether the conclusions stem from an endgame effect, i.e., are particular to the process assumed—or whether the conclusions stem from other features of the model, features one might actually be interested in. This may be a drawback for the application of CPT in finite time horizon models of naive behavior. Infinite time, on the other hand, yields very extreme predictions always. Therefore, the results of this chapter are fundamental to any paper that investigates prospect theory's predictions for naive agents in dynamic decision models.

8. Appendix

Proof of Observation 1. Here we show that all commonly used weighting functions exhibit likelihood insensitivity according to our Assumption 6. Most results are not new, but we think that the following collection is convenient, and we are not aware of any citable source.

The weighting function of Kahneman and Tversky (1979) given by

$$w^+(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{\frac{1}{\delta}}}$$

is differentiable on $(0,1)$, and the derivative is given by

$$w'(p) = \delta p^{\delta-1} \left(p^\delta + (1-p)^\delta\right)^{-\frac{1}{\delta}} \left[1 + \frac{p^\delta (1-p)^{\delta-1}}{p^\delta + (1-p)^\delta}\right].$$

For $\delta \geq 0.28$ the function is strictly increasing. For $\delta > 1$ the function is S-shaped while for $\delta = 1$ we have $w(p) = p$. The interesting parameter range thus is $\delta \in (0,1)$ for which w is increasing and likelihood insensitive. For $\delta \in (0,1)$ it is easy to see that

$$\lim_{p \searrow 0} w'(p) = \lim_{p \nearrow 1} w'(p) = +\infty.$$

The weighting function of Prelec (1998) is given by

$$w(p) = (\exp(-(-\ln(p))^a))^b \tag{8.1}$$

with both a and b strictly positive. According to Wakker (2010), p. 207, "Ongoing empirical research suggests that $a = 0.65$ and $b = 1.05$ [...] are good parameter choices for gains." These choice implies a strong likelihood insensitivity. A special case of the function axiomatized by Prelec is for $b = 1$:

$$w(p) = \exp(-(-\ln(p))^a).$$

We show that w is likelihood insensitive if $a < 1$. Similar arguments imply an S-shape of w if $a > 1$. For $a = 1$ the function has power form and is thus either convex ($b > 1$), linear, ($b = 1$), or concave $b < 1$, and also not sufficing Assumption 6. The derivative of the general version, given by equation (8.1), is given by

$$w'(p) = \frac{ab}{p} \left(-\ln(p)\right)^{a-1} \cdot w(p).$$

It is straightforward that $\lim_{p\nearrow 1} w'(p) = +\infty$ for $a < 1$. To compute $\lim_{p\searrow 0} w'(p)$, substitute $x = -\log(p)$, and observe that

$$\begin{aligned} w'(p) &= \frac{ab}{p}(-\log(p))^{a-1}(\exp(-(-\log(p))^a))^b = \frac{ab}{\exp(-x)}x^{a-1}\exp(-b\cdot x^a) \\ &= a\cdot b\cdot x^{a-1}\exp(x - b\cdot x^a) \\ &= a\cdot b\cdot x^{a-1}\exp(x(1-b\cdot x^{a-1})). \end{aligned}$$

Note that $p \to 0$ as $x \to \infty$. The expression $\lim_{x\to\infty} x(1-\frac{b}{x^{1-a}})$ goes to infinity if and only if $a < 1$. Moreover,

$$\frac{\exp(x(1-b\cdot x^{a-1}))}{x^{1-a}}$$

goes to infinity as x^{1-a} is only of polynomial growth while $\exp(x(1-b\cdot x^{a-1}))$ is of exponential growth.

The weighting function of Goldstein and Einhorn (1987) is defined by

$$w(p) = \frac{bp^a}{bp^a + (1-p)^a}$$

for $a > 0$ and $b > 0$. According to Wakker (2010), p. 208, "The choices $a = 0.69$ and $b = 0.77$ fit commonly found data well." These parameter choices imply a mild inverse-S-shape. The interesting parameter range for both a and b is $(0, 2)$. $a = 1$ and $b = .77$ imply linearity. For $b = 0.77$ fixed, $a > 1$ implies S-shape, and for $a < 1$ the function is inverse-S-shaped, with likelihood insensitivity decreasing (i.e., stronger inverse-S-shape) for $a \searrow 0$.

The derivative is

$$\begin{aligned} w'(p) &= \frac{abp^{a-1}\left(bp^a + (1-p)^a\right) - bp^a\left(bap^{a-1} + a(1-p)^{a-1}(-1)\right)}{\left(bp^a + (1-p)^a\right)^2} \\ &= \frac{abp^{a-1}\left[bp^a + (1-p)^a - bp^a + p(1-p)^{a-1}\right]}{\left(bp^a + (1-p)^a\right)^2} \\ &= \frac{abp^{a-1}\left[(1-p)^{a-1}\left((1-p) + p\right)\right]}{\left(bp^a + (1-p)^a\right)^2} \\ &= \frac{abp^{a-1}(1-p)^{a-1}}{\left(bp^a + (1-p)^a\right)^2}. \end{aligned}$$

The following holds for arbitrary $b > 0$. For $0 < a < 1$ we have

$$\lim_{p\searrow 0} w'(p) = \lim_{p\nearrow 1} = +\infty,$$

which indicates a likelihood insensitive weighting function. For $b > 0$ and $a = 1$,

$$\lim_{p\searrow 0} w'(p) = b \text{ and } \lim_{p\nearrow 1} = 1.$$

For $b > 0, a > 1$,

$$\lim_{p\searrow 0} w'(p) = \lim_{p\nearrow 1} = 0,$$

which is consistent with an S-shaped weighting function. The neo-additive weighting function is defined for a, b positive, $a + b \leq 1$, $w(0) = 0$ and $w(1) = 1$, and for $p \in (0, 1)$:

$$w(p) = b + ap.$$

That is, this function is (in general) discontinuous in 0 and 1 and linear on the interior of its domain. Therefore, it is likelihood insensitive according to our Assumption 6 except for $b = 0$ and $a = 1$.
Finally, let us consider the weighting function proposed by Rieger and Wang (2006) which can be calibrated such that Assumption 6 is not fulfilled. For $a, b \in (0,1)$ it is given by

$$w(p) = \frac{3-3b}{a^2-a+1}\left(p^3-(a+1)p^2+ap\right)+p$$

with derivatives at 0 and 1

$$w'(0) = \frac{3-3b}{a^2-a+1}a+1 \text{ and } w'(1) = \frac{3-3b}{a^2-a+1}(1-a)+1.$$

Moreover, it is easy to show that

$$\frac{\partial}{\partial a}w'(0) > 0,\ \frac{\partial}{\partial a}w'(1) < 0, \frac{\partial}{\partial b}w'(0) < 0 \text{ and } \frac{\partial}{\partial b}w'(1) < 0$$

which implies that

$$\sup_{a,b\in(0,1)} w'(0) = \lim_{b\to 0}\lim_{a\to 1} w'(0) = 4 \text{ and } \sup_{a,b\in(0,1)} w'(1) = \lim_{b\to 0}\lim_{a\to 0} w'(1) = 4.$$

It then easily follows that $w'(0)$ and $w'(1)$ may take any value in $(0,4)$. The smaller b, the more pronounced is the inverse-S shape of w and also the steeper are the functions at $0,1$. By construction, $w(a) = a$, and the derivative at 0 (1) is increasing (decreasing) in a. Thus a allows to account for different overweighting of good- and bad-outcome probabilities. Generally, the higher b the more likely our Assumption 6 is fulfilled. Unfortunately, we are not aware of any empirical estimates of the Rieger-Wang weighting functions. Moreover, Azevedo and Gottlieb (forthcoming) show that their unbounded profits paradox emerges for this function in combination with both power and exponential utility.

Full Proof of Theorem 3.1. Proof of case 2 ($x < r$.) Choose $b < r$ such that both a and b are losses. In that case, lottery $L = L(p,b,a)$ secures the agent a utility of

$$CPT(L) = (1-w^-(1-p))U(b) + w^-(1-p)U(a)$$

with $1-p = \frac{b-x}{b-a}$. Therefore, the agent continues to gamble if there exist $a < x$ and $b > x$ such that

$$\begin{aligned}
0 &< \left(1-w^-\left(\frac{b-x}{b-a}\right)\right)U(b) + w^-\left(\frac{b-x}{b-a}\right)U(a) - U(x)\\
&= U(b)-U(a)+U(a)-U(x)-w^-\left(\frac{b-x}{b-a}\right)(U(b)-U(a))\\
&= (U(b)-U(a))\left(1-w^-\left(\frac{b-x}{b-a}\right)+\frac{U(a)-U(x)}{U(b)-U(a)}\right)\\
&= \underbrace{(1-p)\,(U(b)-U(a))}_{\geq 0}\left(\frac{w^-(1)-w^-(1-p)}{1-p} - \frac{\frac{U(x)-U(a)}{x-a}}{\frac{U(b)-U(a)}{b-a}}\right) \qquad (8.2)
\end{aligned}$$

which is the analogue to equation (3.1). Therefore, similar to the proof of case 1, according to equation (8.2), the agent prefers lottery $L(p,b_n,a_n)$ over x if

$$0 < \frac{w^-(1)-w^-(1-p)}{p} - \frac{\frac{U(x)-U(a_n)}{x-a_n}}{\frac{U(b_n)-U(x)}{b_n-x}(1-p)+\frac{U(x)-U(a_n)}{x-a_n}p} =: \xi_n(p), \qquad (8.3)$$

which is the analogue to equation (3.2). The proof continues similar to that of case 1.
Proof of case 3 ($x = r$.) Consider $x = r$ such that a is a loss and b is a gain. In that case, lottery $L(p, b, a)$ secures the agent a utility of

$$CPT(L) = w^-(1-p)U(a) + w^+(p)U(b)\,.$$

Note that, since $x = r$ by definition $U(x) = U(r) = 0$. Therefore, the chooses L over x if there exist $a < x$ and $b > x$ such that

$$\begin{aligned}
0 &< w^-(1-p)U(a) + w^+(p)U(b) - U(x) \\
&= w^+(p)\left(U(b) - U(a)\right) + \left(U(a) - U(x)\right)\left(w^-(1-p) + w^+(p)\right) \\
&= \left(U(b) - U(a)\right)\left(w^+(p) - \frac{U(x) - U(a)}{U(b) - U(a)}\left(w^-(1-p) + w^+(p)\right)\right) \\
&= \underbrace{p\left(U(b) - U(a)\right)}_{\geq 0}\left(\frac{w^+(p)}{p} - \frac{\frac{U(x)-U(a)}{x-a}}{\frac{U(b)-U(a)}{b-a}}\left(w^-(1-p) + w^+(p)\right)\right)\,. \qquad (8.4)
\end{aligned}$$

First suppose $\frac{w^+(p)}{p} \to \infty$ for $p \to 0$ then the condition follows from the fact that $w^-(1-p) + w^+(p) \leq 2$ two by definition. Suppose ${w^+}'(0)$ exists then

$$\lim_{p\searrow 0} w^-(1-p) + w^+(p) \leq 1 + \lim_{p\searrow 0} w^+(p) = 1\,.$$

Consequently, the sufficient limit condition for gambling is like in case 1, equation (3.4).

Full Proof of Theorem 3.2. Suppose that a_n and b_n are is in the proof of Theorem 3.1. For the power-S-shaped value function it is easily seen that

$$\frac{U(x) - U(a_n)}{x - a_n} = \frac{0 + \lambda\left(\frac{p}{n}\right)}{\left(\frac{p}{n}\right)^\alpha} = \lambda n^{1-\alpha} p^{\alpha-1}$$

and

$$\frac{U(b_n) - U(a_n)}{b_n - a_n} = \frac{\left(\frac{1-p}{n}\right)^\alpha + \lambda\left(\frac{p}{n}\right)^\alpha}{\frac{1-p}{n} + \frac{p}{n}} = n^{1-\alpha}\left((1-p)^\alpha + \lambda p^\alpha\right)\,.$$

Hence,

$$\frac{\frac{U(x)-U(a)}{x-a}}{\frac{U(b)-U(a)}{b-a}} = \lambda\frac{p^{\alpha-1}}{\left((1-p)^\alpha + \lambda p^\alpha\right)}\,.$$

Therefore, according to equation (8.4), L is attractive if

$$\begin{aligned}
\iff 0 &< p(U(b) - U(a))\left(\frac{w^+(p)}{p} - \lambda\frac{p^{\alpha-1}}{\left((1-p)^\alpha + \lambda p^\alpha\right)}\left(w^-(1-p) + w^+(p)\right)\right) \\
&= \left(\frac{w^+(p)}{p^\alpha} - \lambda\frac{1}{\left((1-p)^\alpha + \lambda p^\alpha\right)}\left(w^-(1-p) + w^+(p)\right)\right)\,.
\end{aligned}$$

Since $\lambda > 1$, similarly to the proof of Theorem 3.1, case 3, it follows that

$$\lim_{p\to 0}\frac{w^+(p)}{p^\alpha} > \lambda$$

is a sufficient condition for gambling. □

Example for Never Stopping in Discrete and Finite Time. Consider the five-period binomial decision tree of Barberis (2012). Assume a casino that offers a fair version of French Roulette. We assume a fair casino to be close to the model of Barberis (2012). Then the basic gamble considered by Barberis is the fair analogue to a bet on Red or Black, which occur with equal probability. Now suppose the agent can also bet on a single number, which occurs with probability $\frac{1}{37}$. Consider an agent who only considers to bet 10 units of money on a single number. He is not even able to form a gambling strategy over several periods. This implies a rather coarse strategy space, a feature which is actually working against our never stopping result. However, the basic gamble is skewed whereas the basic gamble Barberis (2012) is symmetric. Let

Figure 1. Gambling Utility for a Symmetric and a Skewed gamble

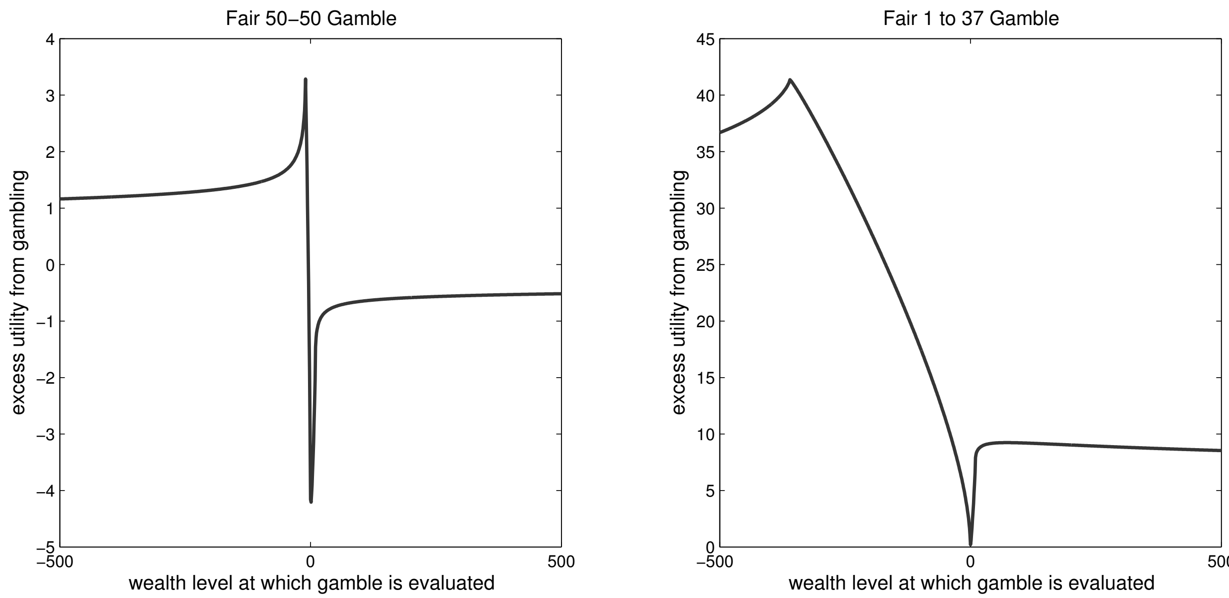

This figure shows the excess utility an agent gains from gambling (over not gambling) for different wealth levels. The left panel shows the utility from gambling a fair 50-50 bet, while the right panel shows the utility from gambling a fair 1 to 37 bet. The agent is a CPT maximizer with the parametrization of Tversky and Kahneman (1992) with parameters given by $\alpha = 0.88$, $\delta = 0.65$, and $\lambda = 2.25$. The agent's reference point is 0.

$(X_t)_{t \in \mathrm{R}_+}$ be the binomial random walk that represents his wealth. It increases by 360 with probability $\frac{1}{37}$ and decreases by 10 with probability $\frac{36}{37}$, starting at some level $X_0 \in \mathrm{R}$, i.e.,

$$\begin{aligned} \mathrm{P}(X_{t+1} = X_t + 360) &= \frac{1}{37} \text{ and} \\ \mathrm{P}(X_{t+1} = X_t - 10) &= \frac{36}{37}. \end{aligned}$$

The agent is forced to stop in the final period T, which is exogenous, or if the random walk reaches zero.

Suppose the agent has CPT preferences given by the original parametrization of Tversky and Kahneman (1992) with parameters as estimated by the authors. Figure 1 plots the excess utility from gambling for the two basic gambles described above. For the 50-50 gamble (left panel), gambling is attractive over the area of losses, and unattractive at the reference point and thereafter. This fits with the common intuition of risk

seeking over losses and risk aversion over gains, which is induced by the S-shaped value function.
Note that the probability weighting component has not much grip when evaluating 50-50 gambles. However, the right panel shows that gambling the skewed basic gamble is attractive *everywhere*. The lowest utility from gambling is at the reference point, but this utility is still positive (the exact value is $+0.56$). Therefore, at any node of the binomial tree, the agent will want to gamble. That is, the agent never stops even though we have finite time with an arbitrary number of gambling periods and a rather limited strategy space. Only one basic gamble is available, but this gamble is sufficiently skewed to be attractive to this very CPT agent. A stop-loss plan would grant even higher utility to the agent, but the one-shot gamble is attractive in itself already.

Bibliography

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. 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*, 11, 215–226.

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

Argote, L., S. Beckman, and D. Epple (1990): "The Persistence and Transfer of Learning in an Industrial Setting," *Management Science*, 36, 140–154.

Azevedo, E., and D. Gottlieb (forthcoming): "Risk-Neutral Firms Can Extract Unbounded Profits from Consumers with Prospect Theory Preferences," *Journal of Economic Theory*.

Barberis, N. (2012): "A Model of Casino Gambling," *Management Science*, 58(1), 35–51.

Barberis, N., and M. Huang (2008): "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review*, 98(5), 2066–2100.

Barberis, N., and W. Xiong (2009): "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation," *Journal of Finance*, 64(2), 751–784.

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

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

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

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

Benkard, C. L. (2000): "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 90, 1034–1054.

Bergemann, D., and J. Välimäki (2008): "Bandit Problems," *The New Palgrave Dictionary of Economics*, 2nd edition, ed. by S. Durlauf and L. Blume.

Bernoulli, D. (1738/1954): "Specimen Theoriae Novae de Mensura Sortis. Commentarii Academiae Scientiarum Imperialis Petropolitanae. Translated, Bernoulli, D., 1954. Exposition of a New Theory on the Measurement of Risk.," *Econometrica*, 22, 23–36.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010): "Learning-by-doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78, 453–508.

Blume, A., and P. Heidhues (2006): "Private monitoring in auctions," *Journal of Economic Theory*, 131(1), 179–211.

Bolton, P., and C. Harris (1999): "Strategic experimentation," *Econometrica*, 67(2), 349–374.

——— (2000): "Strategic experimentation: The undiscounted case," *Incentives, Organizations and Public Economics–Papers in Honour of Sir James Mirrlees*, pp. 53–68.

Bonatti, A., and J. Horner (2011): "Collaborating," *The American Economic Review*, 101(2), 632–663.

Boyer, B., T. Mitton, and K. Vorkink (2010): "Expected Idiosyncratic Skewness," *Review of Financial Studies*, 23(1), 169–202.

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), 1–16.

Compte, O. (1998): "Communication in repeated games with imperfect private monitoring," *Econometrica*, pp. 597–626.

de Georgi, E., and T. Hens (2006): "Making Prospect Theory Fit for Finance," *Financial Markets and Portfolio Management*, 20(3), 339–360.

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*.

Dechenaux, E., D. Kovenock, and V. Lugovskyy (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.

Dixit, A. K., and R. S. Pindyck (1994): *Investment under Uncertainty*, vol. 15. Princeton University Press.

Ebert, S. (2011): "A Note on Skewness and Binary Risks," Working Paper, University of Bonn.

Ebert, S., and D. Wiesen (2011): "Testing for Prudence and Skewness Seeking," *Management Science*, 57(7), 1334–1349.

Goldstein, W., and H. Einhorn (1987): "Expression Theory and the Preference Reversal Phenomena," *Psychological Review*, 94(2), 236–254.

Golec, J., and M. Tamarkin (1998): "Bettors Love Skewness, Not Risk, at the Horse Track," *Journal of Political Economy*, 106(1), 205–225.

Green, T., and B. Hwang (forthcoming): "Initial Public Offerings as Lotteries: Skewness Preference and First-Day Returns," *Management Science*.

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.

Henderson, V. (forthcoming): "Prospect Theory, Liquidation and the Disposition Effect," *Management Science*.

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.

Hopenhayn, H., and F. Squintani (2011): "Preemption Games with Private Information," *Review of Economic Studies*, 78, 667–692.

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

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–292.

Kandori, M., and H. Matsushima (1998): "Private observation, communication and collusion," *Econometrica*, pp. 627–652.

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.

Kaustia, M. (2010): "Prospect Theory and the Disposition Effect," *Journal of Financial and Quantitative Analysis*, 45(3), 781–812.

Keller, G., and S. Rady (2010): "Strategic experimentation with Poisson bandits," *Theoretical Economics*, 5(2), 275–311.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic experimentation with exponential bandits," *Econometrica*, 73(1), 39–68.

Köbberling, V., and P. Wakker (2005): "An Index of Loss Aversion," *Journal of Economic Theory*, 122(1), 119–131.

Kraus, A., and R. Litzenberger (1976): "Skewness Preference and the Valuation of Risky Assets," *Journal of Finance*, 31(4), 1085–1100.

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.

Kyle, A., H. Ou-Yang, and W. Xiong (2006): "Prospect Theory and Liquidation Decisions," *Journal of Economic Theory*, 129(1), 273–288.

Lazear, E. P., and S. Rosen (1981): "Rank-Order Tournaments as Optimal 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*.

Nash, J. (1951): "Non-cooperative games," *The Annals of Mathematics*, 54(2), 286–295.

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

Odean, T. (1998): "Are Investors Reluctant to Realize Their Losses?," *Journal of Finance*, 53(5), 1775–1799.

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

Prelec, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497–527.

Revuz, D., and M. Yor (1999): *Continuous Martingales and Brownian Motion*. Springer.

Rieger, M. O., and M. Wang (2006): “Prospect Theory and the St. Petersburg Paradox.,” *Economic Theory*, 28, 665–679.

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

Rosenberg, D., E. S., and N. Vieille (2007): “Social Learning in One-Armed Bandit Problems,” *Econometrica*, 75, 1591–1611.

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

Rothschild, M. (1974): “A two-armed bandit theory of market pricing,” *Journal of Economic Theory*, 9(2), 185–202.

Schmidt, U., and H. Zank (2008): “Risk Aversion in Cumulative Prospect Theory,” *Management Science*, 54, 208–216.

Seel, C., and P. Strack (2009): “Gambling in Dynamic Contests,” *Working Paper*.

Shefrin, H., and M. Statman (1985): “The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence,” *Journal of Finance*, 40(3), 777–790.

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*.

Skorokhod, A. V. (1965): *Studies in the theory of random processes*, Translated from the Russian by Scripta Technica, Inc. Addison-Wesley Publishing Co., Inc., Reading, Mass.

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

Thompson, P. (2001): “How Much Did the Liberty Shipbuilders Forget? New Evidence for an Old Case Study,” *Journal of Political Economy*, 109, 103–137.

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

Tversky, A., and D. Kahneman (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 5, 297–323.

von Neumann, J., and O. Morgenstern (1944): *Theory of Games and Economic Behavior*. Princeton University Press.

Wakker, P. (2010): *Prospect Theory for Risk and Ambiguity*. Cambridge University Press.

Wakker, P., and A. Tversky (1995): “Risk Attitudes and Decision Weights,” *Econometrica*, 63(6), 1255–1280.

Xu, Z., and X. Zhou (forthcoming): “Optimal Stopping under Probability Distortion,” *Annals of Applied Probability*.