WARWICK
THE UNIVERSITY OF WARWICK

Original citation:
Alpern, Steve and Howard, J. V.. (2017) Winner-takes-all games : strategic optimisation of rank. Operations Research .

Permanent WRAP URL:
http://wrap.warwick.ac.uk/87473

Copyright and reuse:
The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:
Published version: https://doi.org/10.1287/opre.2017.1635

A note on versions:
The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

warwick.ac.uk/lib-publications

Winner-Takes-All Games: Strategic Optimisation of Rank

Steve Alpern* J. V. Howard†

23 January, 2017

Abstract

In many competitive situations (including nearly all sports) a player's aim is not simply to maximise his score but to maximise its rank among all scores. Examples include sales contests (where the salesman with highest monthly sales gets a bonus) and patent races (where lowest time is best). We assume the score X_i of player i is obtained costlessly, so that his utility is the probability of having the best score. This gives a constant-sum game. All that matters for player i is the distribution of his score, so we assume he chooses from a given convex set of distributions $\mathscr{F}_i$. We call such games Distribution Ranking Games, and characterise their solution for various classes of the distribution sets $\mathscr{F}_i$, such as distributions with given mean or moment, where we extend a result of Bell and Cover.

Keywords: Game theory, gambling, contests, league tables, sport, fund management.

*Steve Alpern is at the Warwick Business School, University of Warwick, Coventry CV4 7AL, UK `Steve.Alpern@wbs.ac.uk`

†J. V. Howard is at the London School of Economics, Houghton Street, London WC2A 2AE, UK `j.v.howard@lse.ac.uk`

1 Introduction

A common objective in many political, economic and sporting contexts is to maximise one's rank in a list. A mutual fund wants to obtain a higher annual return than its competitors, a weightlifter wishes to lift the heaviest weight, and a baseball team wants to score more runs. The subject of this article can be fairly summarised by saying that "relative performance is more important than absolute performance".

To model this type of problem, we consider games where each player i gets a *score* X_i and he wins the game if his score is the highest. One could be slightly more general and say that his utility is a function of his rank, but here we will stick to the simpler scenario of highest score wins. Ties will be broken fairly (equiprobably): we will ignore this possibility in our discussions (but not in our analyses).

In some games, it takes *effort* to generate a (high) score. This is the case in the definitive analysis of such ranking games given by Goldberg et al. (2013). Here we make the opposite assumption, that scores are generated costlessly by taking risky decisions. This makes our games constant-sum games, where each player's utility is simply his probability of having the highest score (i.e. of winning). If scores were costly, the game would not be constant-sum, and our methods would not apply. More accurately, we assume not that effort does not contribute to a player's score, but that maximum effort can be assumed, and that the variability in a player's score arises elsewhere (through choice and chance).

Our approach assumes independent choices by the players. This is not essential to our analysis, but makes it simpler. Various analyses of multistage sports events involve dependent choices, for example high jump competitions, as in Gerchak and Kilgour (1993), where the distributions have two atoms (success at the chosen height or failure). In a subsequent article we plan to allow for dependent scores. Thus if the game has rounds (innings, halves, primaries) and a player can take into account the current difference in scores, then the final scores are certainly not independent. So multiround games with intermediate observation of scores will not be considered here.

To introduce and motivate our model, consider a game between two darts players who simultaneously throw at a board, with the higher score winning. A player's strategy consists of many elements. He decides where to stand and where to aim, and can make corrections towards the end of his throw to compensate for what he has done earlier. Thus the strategy is dynamic. However, the outcome of all these complicated decisions can be summarised as a particular random variable X_i, his score, and all that matters in the game is its distribution. So in a reduction analogous to that from extensive to strategic form, we specify the set of random variables X_i that player i can choose from (or the set $\mathscr{F}_i$ of distributions that the X_i can have). In the darts game, the sets $\mathscr{F}_i$ might reflect the skill of each player i. We define a *Distribution Ranking (DR) game* $\Gamma = \Gamma(\mathscr{F}_1, \mathscr{F}_2, \ldots, \mathscr{F}_k)$ as one where each player i chooses a random variable X_i with distribution in $\mathscr{F}_i$, and his payoff is the probability that X_i is the maximum. Thus DR games have constant sum 1. We will often assume that the distribution sets $\mathscr{F}_i$ are all convex, so that mixtures are already in the set, and mixed strategies need not be considered. We will write $\Gamma(\mathscr{F})$ for $\Gamma(\mathscr{F}, \mathscr{F})$.

Our primary motivating example, the one which led us to study this subject, is the following infinite time horizon variation on an intriguing game of Anderson (2012).

Example 1 (Stock selling games and slot machines)**.** *A father gives independent unbiassed stocks (continuous stochastic processes) $Y_1(t)$ and $Y_2(t)$ to his two daughters, with initial prices μ_1 and μ_2. By unbiassed, we mean that they are martingales, and so $\mathbb{E}[Y_i(t)] = \mu_i$ for all t. We assume the processes will leave any bounded interval $[a, b]$, where $0 < a < \mu_i < b$, in finite expected time. Each daughter can secretly sell her stock at any time before it reaches price 0 (when she is forced to 'sell'); her strategy must have a finite expected selling time. The daughter who sells at the higher price wins the father's huge fortune, which is her only aim. A discrete version of this has the daughters sent to two casinos with μ_i chips each. As long as they still have chips left, they can put one into a fair slot machine which equiprobably returns no chips or two chips. Again, the one who leaves her casino with the most chips wins the game. When should they sell their stocks or leave the casino?*

In Anderson's example the players must sell by a fixed time; in our variation they are not limited by time, but they must stop if their capital balance falls to zero. In our examples, each daughter's distribution set $\mathscr{F}_i$ is the family of non-negative distributions which can be obtained by using some stopping rule (with finite expected stopping time). All these distributions will have mean μ_i. Such examples with fixed mean, or more generally, fixed moment, will be particularly significant in our analysis. Bell and Cover (1980) gave perhaps the simplest example of this type of game: two players move simultaneously, each choosing a distribution with mean 1 on the interval $[0, \infty)$. (All later references to Bell and Cover are to this 1980 paper unless a specific alternative is given.) As the set of distributions is closed under mixing, this game should either have no solution or else a solution in pure strategies. In fact Bell and Cover showed that the solution is the uniform distribution on $[0, 2]$ (perhaps a somewhat counter-intuitive result). In Section 2 we generalise their result to symmetric many-player games where the distributions may be restricted more generally than just by the mean, and in Section 4 to asymmetric two-player games. In the latter case, we find that if player I's score random variable X_1 is restricted by $\mathbb{E}\left[\phi(X_1)\right] \leq 1$ for a given continuous strictly increasing function ϕ, and if similarly player II has a constraint $\mathbb{E}\left[\psi(X_2)\right] \leq 1$, then the optimal solution has player I choosing a distribution function which is an affine transformation of an initial segment of ψ, and similarly for player II and ϕ. (Effectively, each player plays the other's constraint function.) The Bell-Cover result is recovered when $\phi(x) = \psi(x) = x$.

The paper is organised as follows. Section 2 presents our symmetric Nash solution to the games where k players have a common restriction on the mean (or generalised moment) of their distribution. Section 3 defines DR games. Section 4 analyses DR games where the sets $\mathscr{F}_i$ are given by restrictions on generalised moments of their distributions. Here we will give the general solution to the stock market/slot machine games of Example 1. The final part of Section 4 uses a graphical analysis of best responses to another player's strategy to give an intuitive understanding of many of the analytical results and to prove the uniqueness of our solution. This can be used to solve problems where the players use bounded support intervals. In Section 5 we give a more substantial example (the silent duel) to show an application of the theory: in this case it gives a new derivation of an old solution. The paper ends with suggestions for further work.

It may be useful to highlight the main results on the solutions we find for various classes of DR games. Theorem 2.1 (in Section 2) gives the explicit solution to the symmetric DR game played by k identical players with a generalised moment constraint. Theorems 4.1 and 4.2 solve the DR game for two asymmetric players (again with generalised moment constraints), the latter giving optimal response strategies and consequent uniqueness.

2 Symmetric k-player games on $[0,\infty)$

The Bell-Cover game has two players who each choose a distribution with mean 1 on the interval $[0,\infty)$. Independent samples are drawn from these distributions and the higher value wins (ties are broken at random). We generalise this game in two ways. Firstly we allow for k players for finite $k \geq 2$. Secondly, instead of the players being restricted by the first moment of the distribution (the mean), we allow much broader restrictions (including, in particular, restrictions on any positive moment). So we allow what we will call *generalised moment* constraints, i.e. bounds on $\mathbb{E}[\phi(X)]$ for $X \sim F \in \mathscr{F}$, where

ϕ is a given continuous function on $[0,\infty)$, strictly increasing with an asymptotic value of at least k. We look for a symmetric equilibrium in which all the players choose the same distribution function, say F. Would any single player wish to change his choice, assuming the others did not? The single player would know that the maximum score, Z, of the other $k-1$ players has distribution function F^{k-1} (because $Z \le z$ if and only if each of the other players has a score $\le z$). He might think of it as a two-player game, except that he expects to win with probability $1/k$ at a symmetric equilibrium. This suggests the following Theorem.

Theorem 2.1. *Let $\phi(x)$ be a continuous strictly increasing function on $[0,\infty)$ with $\phi(0)=0$, $\lim_{x\to\infty}\phi(x)\ge k$. Let $\mathscr{F}$ be the set of non-negative distributions F with $\mathbb{E}[\phi(X)]\le 1$ when $X\sim F\in\mathscr{F}$. Suppose that in the k-player game $\Gamma=\Gamma(\mathscr{F},\mathscr{F},\dots,\mathscr{F})$ each player, i, must simultaneously choose a distribution, $F_i\in\mathscr{F}$. A score, $X_i\sim F_i$ is sampled from each distribution, and the highest score wins. Then Γ has an equilibrium when all players choose the distribution $F\in\mathscr{F}$ on $[0,\phi^{-1}(k)]$ with distribution function*

$$F(x)=\sqrt[k-1]{\frac{\phi(x)}{k}}\,.$$

Proof. Clearly F is a continuous distribution function. Also

$$\begin{aligned}\mathbb{E}_F\left[\phi(X)\right] &= \int_{x=0}^{\phi^{-1}(k)}\phi(x)\,dF(x)=\int_{F=0}^{1}kF^{k-1}\,dF\\ &=\left[F^k\right]_0^1=1.\end{aligned}$$

So $F\in\mathscr{F}$.

If the first $k-1$ players choose F, the distribution function of the maximum score on $[0,\phi^{-1}(k)]$ among players $1,\dots,(k-1)$ is $[F(x)]^{k-1}=\phi(x)/k$. If player k's score is y, his chance of winning will be $\min(\phi(y)/k,1)$; and so (law of iterated expectations) if he chooses $Y\sim G\in\mathscr{F}$, his chance of winning will be

$$=\mathbb{E}_G\left[\min\left(\frac{\phi(Y)}{k},1\right)\right]\le\mathbb{E}_G\left[\frac{\phi(Y)}{k}\right]\le\frac{1}{k}\,.$$

Player k can achieve this bound by choosing any distribution, G, on $[0,\phi^{-1}(k)]$ with $\mathbb{E}_G\left[\phi(Y)\right]=1$. F is such a distribution, so player k has no reason to change from it. Hence we have an equilibrium. □

Corollary 2.1.1. *Let $\mathscr{Z}_n$ be the set of non-negative distributions with n'th moment bounded by 1. The k-player game $\Gamma(\mathscr{Z}_n,\mathscr{Z}_n,\dots,\mathscr{Z}_n)$ has an equilibrium when all players choose the distribution $Q_{n,k}\in\mathscr{Z}_n$*

on $[0, \sqrt[n]{k}]$ *where*

$$Q_{n,k}(x) = \sqrt[k-1]{\frac{x^n}{k}}. \tag{2.1}$$

Proof. Take $\phi(x) = x^n$. □

Corollary 2.1.2. *In the two-player DR game* $\Gamma(\mathcal{I}_2)$ *with the players restricted to non-negative distributions with mean square of* 1*, each player can guarantee a payoff of* $1/2$ *by playing the distribution on* $[0, \sqrt{2}]$ *which has the density function* $f(x) = x$. □

Recall that the mean square equals the variance plus the mean squared, so if the distributions have mean μ and variance σ^2, the restriction is $\mu^2 + \sigma^2 \leq 1$. The solution distribution has $\mu = \sqrt{8/9}$ and $\sigma = 1/3$.

Corollary 2.1.3 (Bell and Cover)**.** *In the two-player DR game on* $[0,\infty)$ *with the players restricted to non-negative distributions with mean* 1*, each player can guarantee* $1/2$ *by playing the uniform distribution on* $[0,2]$.

Proof. Take $n = 1$ and $k = 2$ in Corollary 2.1.1. □

Bell and Cover also showed uniqueness, which we will do later in a more general context.

Note that when $k = 2$ the solution distribution function F is simply half the constraint function ϕ on $[0, \phi^{-1}(2)]$. We shall see later how this can be generalised to the asymmetric case.

2.1 Stock selling solution in the symmetric case

We are now in a position to give an elegant solution to the stock selling games of Example 1 in the symmetric case, where the value (if it exists) must be $1/k$ for k players. We begin with the two player case.

Proposition 2.2 (Two-player symmetric stock selling)**.** *For the two player symmetric stockselling game with initial stock price* 1*, the following strategy guarantees winning with probability at least* $1/2$*, which is therefore the value of the game. Sample a number* d *uniformly in the interval* $[0, 1]$ *and sell when the price reaches* $1 - d$ *or* $1 + d$.

Proof. Suppose d has been sampled. Under our assumptions, the first time, T, that the process moves a distance d from the starting position (i.e. the hitting time for $\{1 - d, 1 + d\}$) is a stopping time to which Doob's optional sampling theorem can be applied. So the expected stock value when we stop is equal to the starting value of 1. But we either sell at a loss at price $(1-d)$, or at a profit at price $(1+d)$, so these must be equally probable.

So half the time we will make a loss, and in these cases (as d has a uniform distribution) the selling price will be uniform on $[0,1]$, and similarly when we make a profit it will be uniform on $[1,2]$. Since these cases are equally likely, this strategy achieves a selling price, or score, which is uniform on $[0,2]$. So suppose player I adopts this strategy. Any mixed strategy that player II adopts will have a score (selling price) with mean 1, so Corollary 2.1.3 guarantees that player I wins with probability at least $1/2$, which is the value of the game. □

A similar approach solves the k player version, but we need a well known result (see, for example, Pinelis (2009)) which says that any distribution with mean 1 is a mixture of two-atom distributions with mean 1.

Proposition 2.3 (Many-player symmetric stock selling)**.** *In the k-player stock selling game, there is a symmetric equilibrium in which each player achieves the score distribution of equation* (2.1) *when* $n=1$.

Proof. By the result mentioned above, this score distribution can be represented as a mixture of two-point mean 1 distributions on points y and z with $y \le 1 \le z$. Sample $\{y,z\}$ from this distribution and sell when the price drops to y or rises to z (or immediately if $y=z=1$). This ensures winning with probability at least $1/k$ if the other players are doing the same, by Corollary 2.1.1. □

The following example shows that we can have a DR game whose solution has unbounded support.

Example 2 (Unbounded solution)**.** *In a symmetric two-player DR game each player must choose a distribution on the interval* $[0,\infty)$. *A referee will sample from each distribution, taking the sampled value as the intensity, λ, (expected number of events per unit time) of a Poisson process of favourable events. The player with the higher sampled intensity is the winner. However the probability (before the intensity is sampled) that the first event in the Poisson process occurs at or after time* 1 *must be at least* $1/2$. *What distribution should the players choose?*

The time to the first event in a Poisson process with intensity λ has an exponential distribution with mean $1/\lambda$. The probability that this time exceeds 1 is $\exp(-\lambda)$. So if the game is $\Gamma(\mathscr{F})$, the players are restricted to non-negative distributions $F \in \mathscr{F}$ such that $\mathbb{E}[\exp(-X)] \ge 1/2$ when $X \sim F$. This constraint can be rewritten as $\mathbb{E}[1-\exp(-X)] \leq 1/2$, so we can take $\phi(x) = 2(1-\exp(-x))$ and apply Theorem 2.1. Each player can guarantee $1/2$ by playing an exponential distribution with mean 1.

2.2 Avoiding elimination

So far we have assumed that the players care only about being first. In the case of two player games, this is the same as not being last. However, we can imagine a situation with k players where the player who is last will be eliminated, and all the other players go through to a later round of the competition. In this case their objective is to avoid having the worst score.

Theorem 2.4. *Let $\phi(x)$ be a continuous strictly increasing function on $[0,\infty)$ with $\phi(0) = 0$, $\lim_{x\to\infty}\phi(x) \geq k/(k-1)$. Let $\mathscr{F}$ be the set of non-negative distributions, F, with $\mathbb{E}\left[\phi(X)\right] \leq 1$ when $X \sim F \in \mathscr{F}$. The k-player Distribution Ranking game in which all players have the sole objective of avoiding being last has an equilibrium when all players choose the distribution $F \in \mathscr{F}$ on $[0,\phi^{-1}(k/(k-1))]$ with distribution function*

$$F(x) = 1 - \sqrt[k-1]{1-\left(\frac{k-1}{k}\right)\phi(x)}\,.$$

Proof. Clearly F is a continuous distribution function. Also (after some manipulation)

$$\mathbb{E}_F\left[\phi(X)\right] = \int_{x=0}^{\phi^{-1}(k/(k-1))} \phi(x)\,dF(x) = 1.$$

So $F \in \mathscr{F}$.

If the first $k-1$ players choose F the distribution function of the minimum score on $[0,\phi^{-1}(k/(k-1))]$ among players $1,\ldots,(k-1)$ is

$$1-[1-F(x)]^{k-1} = \left(\frac{k-1}{k}\right)\phi(x).$$

If player k chooses $Y \sim G \in \mathscr{F}$, then player k's chance of winning will be

$$\mathbb{E}_G\left[\min\left(\left(\frac{k-1}{k}\right)\phi(Y),1\right)\right] \leq \mathbb{E}_G\left[\left(\frac{k-1}{k}\right)\phi(Y)\right] \leq \frac{k-1}{k}\,.$$

Player k can achieve this bound by choosing any distribution, G, on $[0,\phi^{-1}(k/(k-1))]$ with $\mathbb{E}_G\left[\phi(Y)\right] = 1$. F is such a distribution, so player k has no reason to change from it. Hence we have an equilibrium.

□

Corollary 2.4.1. *Let $\mathcal{Z}_n$ be the set of non-negative distributions with n'th moment bounded by* 1*. The k-player distribution ranking game in which the players must choose a distribution in $\mathcal{Z}_n$ with the sole objective of avoiding being ranked last has an equilibrium when all players choose the distribution $R_{n,k} \in \mathcal{Z}_n$ on $[0, \sqrt[n]{k/(k-1)}]$ where*

$$R_{n,k}(x) = 1 - \sqrt[k-1]{1 - \left(\frac{k-1}{k}\right) x^n}\ .$$

Proof. Take $\phi(x) = x^n$. □

3 Distribution Ranking Games

In this section we provide a formal definition of Distribution Ranking (DR) Games and then discuss in more detail the two-player case. The Appendix shows by example that DR games may have no equilibrium solution, a unique solution, or many solutions.

3.1 Definition of a DR game

We begin by briefly describing the general form of a DR game, even though in this paper we will be primarily interested in the two player version. Suppose each player i, $i = 1, \ldots, k$, has a strategy set $\mathscr{F}_i$ of distributions on a closed interval $[a_i, b_i]$. We need lower bounds, for example when $\mathscr{F}_i$ consists of unit mean distributions, to prevent a player from putting a small amount of probability near $-\infty$ and having a high mean for the remainder of his distribution, thus essentially negating the mean restriction. The game is played by each player choosing a distribution function F_i from his set $\mathscr{F}_i$. The distributions F_i are then independently sampled and player i's score is the random variable $X_i \sim F_i$. The payoff to player i is a given function of the rank of his score (higher ranks are assumed to have higher payoffs). We consider henceforth only the 'winner-takes-all' payoff, where each player's payoff is the probability that he has the highest score (after any ties are fairly adjudicated). The Distribution Ranking Game $\Gamma(\mathscr{F}_1, \mathscr{F}_2, \ldots, \mathscr{F}_k)$ is then the k-player game in strategic form in which the strategy set for player i is $\mathscr{F}_i$ and his payoff is his probability of winning. It is a constant-sum game as these probabilities must sum to 1. More formally, let $E_{i,j}$ denote the event that there is a j-way tie for the highest score, and that player i's score is among them. Then the payoff P_i to player i is given by

$$P_i(F_1, F_2, \ldots, F_i, \ldots, F_n) = \sum_{j=1}^{k} \left(\frac{1}{j}\right) \mathbb{P}\left(E_{i,j}\right).$$

In the two person case we will simply denote the payoff to player 1, the probability that he wins, by

$$P = P_1(F_1, F_2) = \mathbb{P}(X_1 > X_2) + \frac{1}{2}\mathbb{P}(X_1 = X_2), \text{ where as usual } X_i \sim F_i. \tag{3.1}$$

3.2 The two person DR game

For the remainder of the paper we will be specialising to two person DR games, so we use the payoff function given above in (3.1). The existence of an equilibrium thus corresponds to the existence of a value $v = v(\mathscr{F}_1, \mathscr{F}_2) \in [0,1]$, which can be interpreted as the probability that player I wins, given optimal play on both sides. Each player has an optimal strategy $\bar{F}_i \in \mathscr{F}_i$, $i = 1,2$. By playing $\bar{F}_i$, player I guarantees P is at least v, and by playing $\bar{F}_2$, player II guarantees that P is no more than v. We also note that this is an anonymous game in the sense that $P(F_1, F_2) = 1 - P(F_2, F_1)$. If F_1 or F_2 is non-atomic, $P(F_1, F_2)$ is simply the probability that X_1 is larger than X_2.

We assume that both players have a lower bound on their distributions, and so (by rescaling if necessary) we may assume that both are restricted to a sub-interval of $[0, \infty)$. The probability that player I either wins or ties is then

$$\mathbb{P}(X_1 \geq X_2) = \int_0^\infty F_2(x)\, dF_1(x) = \mathbb{E}_{F_1}\left[F_2(X_1)\right].$$

Similarly the probability that player II either wins or ties is

$$\mathbb{P}(X_2 \geq X_1) = \int_0^\infty F_1(x)\, dF_2(x) = \mathbb{E}_{F_2}\left[F_1(X_2)\right].$$

If either distribution has no atoms, the probability of winning does not depend on the tie breaking rule. Suppose that F_1 is non-atomic (so that it is continuous and monotonic increasing from 0 to 1 on the real line). The probability of a tie will then be zero. A parametric plot of $F_2(x)$ against $F_1(x)$ might be as shown in Figure 1. (Here F_2 has an atom of probability at a value x in an interval of x-values where F_1 places no probability.) The probability that player I wins is then the area below the curve, and that player II wins is the area above. They clearly sum to 1 (integration by parts). If both distributions

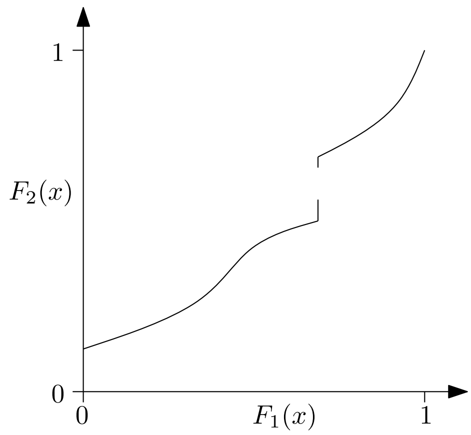

Figure 1: Parametric plot of $(F_1(x), F_2(x))$

have atoms, we can still draw a diagram, but it may then have horizontal gaps where F_1 has atoms, and diagonal gaps where both players place an atom of probability in the same location. In the latter case, it is not clear what is the area below the curve, and the probability of winning will depend on the tie breaking rule. Our assumption that ties are broken at random implies that we should fill all the gaps with straight lines.

A solution to this two-player game is then a pair of optimal pure strategies $\bar{F}_i \in \mathscr{F}_i$ satisfying

$$P\left(F_1, \bar{F}_2\right) \leq P\left(\bar{F}_1, \bar{F}_2\right) = v \leq P\left(\bar{F}_1, F_2\right),$$

for all $F_i \in \mathscr{F}_i$, where $v = v\left(\mathscr{F}_1, \mathscr{F}_2\right)$ is the value of the game.

4 Two-player games on $[0, \infty)$ with generalised moment constraints

In this section we look at DR Games $\Gamma(\mathscr{F}, \mathscr{G})$ where $\mathscr{F}$ and $\mathscr{G}$ are sets of non-negative distributions with constraints on their 'generalised moments' as defined below

$$\mathbb{E}\left[\phi(X)\right] \leq k_1, \text{ i.e. } \int_0^\infty \phi(x)\, dF(x) \leq k_1 \tag{4.1}$$

$$\mathbb{E}\left[\psi(Y)\right] \leq k_2, \text{ i.e. } \int_0^\infty \psi(y)\, dG(y) \leq k_2, \tag{4.2}$$

where ϕ and ψ are given continuous strictly increasing functions on $[0, \infty)$ with $\phi(0) = \psi(0) = 0$ and $X \sim F \in \mathscr{F}$ and $Y \sim G \in \mathscr{G}$. We also need additional conditions (equations (4.3) and (4.4)) which ensure that both constraints are binding in the solution: these are automatically satisfied if ϕ and ψ are unbounded.

Equations (4.1) and (4.2) define half-spaces. Half-spaces are the simplest convex sets: every convex set of distributions in the vector space $C(S)$ is the intersection of half spaces. So these constraints give a very natural DR game to analyse. Theorem 2.1 solved the symmetric case, and we noted that (for two-player games) the players used a rescaling of an initial segment of ϕ $(=\psi)$ as the solution distribution.

To find the solution when $\phi \neq \psi$, we first rescale so that $k_1 = k_2 = 1$. (We could also make a change of variable so that ϕ, say, is the identity, but we have chosen to retain symmetry.) Then we consider what happens if each player tries to use an initial segment of the other's constraint function, suitably normalised, as his distribution function. So player I would plan to use the distribution function $\psi(x)/\psi(b_1)$ on $[0, b_1]$ for some b_1.

We want constraint (4.1) to be binding, and we will show that we can do this if we add the requirement that

$$\lim_{z\to\infty} \frac{1}{\psi(z)} \int_0^z \phi(x)\, d\psi(x) \geq 1\ ,$$

with a similar requirement for player II. Then we can find b_1 to make constraint (4.1) bind, and similarly find b_2 to make constraint (4.2) bind. By relabelling the players if necessary we may assume that $0 < b_2 \leq b_1 \leq \infty$. (It will become clear that if $\phi, \psi \to \infty$ as $x \to \infty$, our requirements will be satisfied automatically, and b_1 will be finite.) The theorem will show that both players use the larger interval $[0, b_1]$, and that player II places an atom of probability at 0.

Theorem 4.1. *Let ϕ and ψ be continuous strictly increasing functions on $[0,\infty)$ with $\phi(0) = \psi(0) = 0$ satisfying*

$$\lim_{z\to\infty} \frac{1}{\psi(z)} \int_0^z \phi(x)\, d\psi(x) \geq 1\ , \textit{ and} \tag{4.3}$$

$$\lim_{z\to\infty} \frac{1}{\phi(z)} \int_0^z \psi(x)\, d\phi(x) \geq 1. \tag{4.4}$$

Let $\mathscr{F}$ be the set of non-negative distributions with $\mathbb{E}\left[\phi(X)\right] \leq 1$ when $X \sim F \in \mathscr{F}$, and $\mathscr{G}$ be the set of non-negative distributions with $\mathbb{E}\left[\psi(Y)\right] \leq 1$ when $Y \sim G \in \mathscr{G}$. Let b be the unique solution to the equation

$$\frac{1}{\psi(b)} \int_0^b \phi(x)\, d\psi(x) = 1, \tag{4.5}$$

and suppose (wlog) that

$$\frac{1}{\phi(b)} \int_0^b \psi(x)\, d\phi(x) \geq 1. \tag{4.6}$$

Then a solution to the two-player game $\Gamma(\mathscr{F},\mathscr{G})$ *is for both players to play on* $[0,b]$, *player* I *using the distribution* $F(x)$ *and player* II *the distribution* $G(x)$ *where*

$$F(x) = \frac{\psi(x)}{\psi(b)}, \text{ and}$$
$$G(x) = 1 - \frac{\phi(b)}{\phi(b)\psi(b) - \psi(b)} + \frac{\phi(x)}{\phi(b)\psi(b) - \psi(b)}.$$

The game has value $1 - 1/\psi(b)$.

Proof. We first show that b exists and is unique. Let

$$\eta(z) = \frac{1}{\psi(z)} \int_0^z \phi(x)\, d\psi(x).$$

This is a weighted average of $\phi(x)$-values for $x < z$, so is less than $\phi(z)$, and tends to zero as z tends to zero from above. Also, if we increase z to z^*, the new value, $\eta(z^*)$ will be a weighted average of the old value, $\eta(z)$, and an average of $\phi(x)$-values for $z < x < z^*$. So we have an average of the old value (less than $\phi(z)$) and a new quantity (greater than $\phi(z)$). Hence the function $\eta(z)$ is strictly increasing, and (by continuity) equation (4.5) has a unique solution (which may be ∞). So F is a distribution function.

Next note that equation (4.5) states that a weighted average of ϕ-values on $[0,b]$ equals 1, so $\phi(b) > 1$. (Inequality (4.6) shows similarly that $\psi(b) > 1$.) Hence $G(x)$ is a strictly increasing function of x. To show $G(0) \geq 0$ we must show that $\phi(b) \leq \phi(b)\psi(b) - \psi(b)$. However we know that (integration by parts)

$$\int_0^b \phi(x)\, d\psi(x) + \int_0^b \psi(x)\, d\phi(x) = \phi(b)\psi(b).$$

So using equation (4.5) we have

$$\begin{aligned} \phi(b)\psi(b) - \psi(b) &= \int_0^b \psi(x)\, d\phi(x) \qquad (4.7)\\ &\geq \phi(b), \end{aligned}$$

by assumption (4.6). (If the assumption was not satisfied, we would relabel the players.) So G is also a distribution function. F is continuous everywhere and G has one atom at 0 if inequality (4.6) is not binding.

If $X \sim F$,

$$\mathbb{E}\left[\phi(X)\right] = \frac{1}{\psi(b)} \int_0^b \phi(x)\, d\psi(x) = 1,$$

using equation (4.5). If $Y \sim G$,

$$\mathbb{E}\left[\psi(Y)\right] = \frac{1}{\phi(b)\psi(b) - \psi(b)} \int_0^b \psi(x)\, d\phi(x) = 1,$$

using equation (4.7). So F and G are valid strategies for the players.

Also if $X \sim F$ and $Y^* \sim G^* \in \mathscr{G}$:

$$\begin{aligned} \mathbb{P}\left(Y^* \geq X\right) &= \mathbb{E}\left[F\left(Y^*\right)\right] = \frac{1}{\psi(b)} \mathbb{E}\left[\psi\left(Y^*\right)\right] \leq \frac{1}{\psi(b)} \\ \mathbb{P}\left(X > Y^*\right) &\geq 1 - \frac{1}{\psi(b)}; \end{aligned}$$

so F gives player I the required guarantee.

Similarly if $Y \sim G$ and $X^* \sim F^* \in \mathscr{F}$:

$$\begin{aligned} \mathbb{P}\left(X^* \geq Y\right) &= \mathbb{E}\left[G\left(X^*\right)\right] \\ &= \mathbb{E}\left[1 - \frac{\phi(b)}{\phi(b)\psi(b) - \psi(b)} + \frac{\phi(X^*)}{\phi(b)\psi(b) - \psi(b)}\right] \\ &\leq 1 - \frac{\phi(b)}{\phi(b)\psi(b) - \psi(b)} + \frac{1}{\phi(b)\psi(b) - \psi(b)} \\ &= 1 - \frac{1}{\psi(b)} \\ \mathbb{P}\left(Y > X^*\right) &\geq \frac{1}{\psi(b)}; \end{aligned}$$

so G gives player II the required guarantee. □

In applications of this theorem, we may refer to the player who uses a continuous (atomless) distribution function as *dominant*. The next example shows that the dominant player may have a winning probability which is less than $1/2$.

Example 3. *In a DR Game on* $[0,\infty)$*, player* I*'s third moment and player* II*'s second moment are both constrained to be no more than* 1*. What distributions should they choose?*

Straightforward calculations show that $b_1 = \sqrt[3]{5/2} \approx 1.36$, and $b_2 = \sqrt{5/3} \approx 1.29$. So player I is dominant and $b = b_1$. Also

$$F(x) \approx 0.54x^2, \text{ and } G(x) \approx 0.095 + 0.362x^3.$$

The value of the game is ≈ 0.457. Note that although player I is dominant, he wins less than half the time.

Another interesting special case occurs when both players are constrained in the same way, but with different bounds. For example, one player might have his mean square constrained to a maximum of 1, the other to a maximum of 2.

Corollary 4.1.1. *Let χ be a continuous strictly increasing unbounded function on $[0,\infty)$ with $\chi(0) = 0$. Suppose both players must choose distributions on $[0,\infty)$, player I from the set of distributions $\mathscr{F}$ satisfying $\mathbb{E}\left[\chi(X)\right] \le k_1$ when $X \sim F \in \mathscr{F}$, and player II from $\mathscr{G}$ satisfying $\mathbb{E}\left[\chi(Y)\right] \le k_2$ when $Y \sim G \in \mathscr{G}$. Then the value of the DR game $\Gamma(\mathscr{F},\mathscr{G})$ to player I equals $1 - k_2/(2k_1)$ if $k_1 \ge k_2$ and equals $k_1/(2k_2)$ if $k_1 < k_2$.*

Proof. If $k_1 \ge k_2$, take $\phi(x) = \chi(x)/k_1$ and $\psi(x) = \chi(x)/k_2$. We then find the players play on $[0,b]$ where $b = \chi^{-1}(2k_1)$, $F(x) = \chi(x)/2k_1$, and $G(x) = (1 - k_2/k_1) + (k_2/k_1)F(x)$, and the value of the game is $1 - k_2/(2k_1)$ to player I. If $k_1 < k_2$, interchange the players. $\square$

We can simplify further to the case $\chi(x) = x$ in the following result noted in Bell and Cover (1988).

Corollary 4.1.2. *Suppose both players must choose distributions on $[0,\infty)$, player I with mean 1 and player II with mean μ, $0 < \mu < \infty$. Then the value of the game to player II is $v(\mu)$ where*

$$\begin{cases} v(\mu) = \frac{\mu}{2} & \text{if } \mu < 1, \\ v(\mu) = 1 - \frac{1}{2\mu} & \text{if } 1 \le \mu. \end{cases}$$

Note that v is continuous with a continuous derivative, but the second derivative is discontinuous at 1.

Proof. If $\mu < 1$, take $\chi(x) = x$, $k_1 = 1$, and $k_2 = \mu$ in Corollary 4.1.1. We then find $b = 2$, $F(x) = x/2$, $G(x) = (1-\mu) + \mu(x/2)$, and the value of the game is $1 - \mu/2$ to player I. If $\mu > 1$, interchange the players and rescale so the larger mean is 1. $\square$

When $\mu < 1$, player II is 'stealing' player I's strategy a fraction μ of the time (scoring $1/2$ on average when she does so), and playing 0 (and losing) the remainder of the time. In fact, Theorem 4.1 shows that player II is actually stealing her own strategy – i.e. the strategy she would use if she were less constrained.

Corollary 4.1.3. *Let $\mathscr{F}_i$ be the set of distributions whose support is $[a,b]$ and whose mean is μ_i where $a < \mu_2 \le \mu_1 \le \frac{a+b}{2}$. Then the two-player distribution ranking game $\Gamma(\mathscr{F}_1, \mathscr{F}_2)$ has a pure strategy solution with value $1-\mu_2/2\mu_1$. Player I chooses the uniform distribution on $[a, 2\mu_1 - a]$; player II uses the same distribution with probability μ_2/μ_1 and otherwise plays at a.*

Proof. If the players are restricted to non-negative distributions with means 1 (for player I) and $\mu \le 1$ (for player II) the theorem gives the solution immediately, with both players having support $[0,2]$. Hence the same solution applies if the players are restricted to $[0,b]$ where $b \ge 2$, because they are restricted to subsets of distributions which still contain the maximal elements. A linear transformation now gives the result. □

Example 1 revisited again We saw that any distributions of the final values $F \in \mathscr{F}_1$ and $G \in \mathscr{F}_2$ have means μ_1 and μ_2, and that both distributions have support on $\mathbb{R}_{\ge 0}$. We assume that $\mu_1 < \mu_2$, and rescale stock prices so that $\mu_2 = 1$. Now suppose player II chooses a number d uniformly in $[0,1]$ and then adopts the rule of selling when $Y_2(t)$ reaches $1-d$ or $1+d$. This gives her the stopping distribution $\bar{G}$, the uniform distribution on $[0,2]$, which guarantees she wins with probability at least $1-\mu_1/2$ against any distribution in $\mathscr{F}_1$.

Similarly, we can write the distribution $\bar{F}$, which places probability $1-\mu_1$ at 0, and otherwise plays uniform on $[0,2]$, as a distribution over two point distributions p at $a < \mu_1$ and $(1-p)$ at $b > \mu_1$. Specifically,

$$\bar{F}(x) = \int_0^2 T(z) g(z)\, dz.$$

Here $T(z)$ is the equiprobable measure on $\{2\mu_1 - z, z\}$ for $\mu_1 \le z \le 2\mu_1$, and the measure on $\{0, z\}$ with weight μ_1/z on z for $2\mu_1 \le z \le 2$. The density g is given by

$$g(z) = \begin{cases} \mu_1 & \text{if } \mu_1 \le z \le 2\mu_1, \\ z/2 & \text{if } 2\mu_1 \le z \le 2. \end{cases}$$

The distribution defined by the above equations has an atom at 0 of weight

$$= \int_{2\mu_1}^2 \left(1 - \frac{\mu_1}{z}\right) \frac{z}{2}\, dz = 1 - \mu_1.$$

It also has a density $f(x)$ for $0 < x \le 2$, and we easily see this density is constant, $\mu/2$. Hence the equations define the desired distribution.

Thus player I wins this game with probability $\mu_1/2\mu_2$. We again note that this analysis assumes only that the processes will leave any bounded interval in finite expected time.

4.1 Best response distributions

In a standard two-player game, if one player's (mixed) strategy is known, the other player will always have a best response which is a pure strategy, he will not need to use a mixed strategy. (We can think of a pure strategy as a mixed strategy concentrated on an atom of probability 1.) At an equilibrium each player plays a best response to the other's (mixed) strategy, and typically each player will have available many possible best (pure) strategy responses, and can guarantee the value of the game by an appropriate mixture of best response pure strategies.

How does this change for DR Games of the type occurring in Theorem 4.1? The next theorem answers this question, showing that although one-atom distributions may not be sufficient, we need extend only so far as to allow two-atom distributions as well. The theorem is also helpful in suggesting and verifying solution strategies, and in addition gives some geometric intuition about the problem.

Theorem 4.2. *Let $\Gamma(\mathscr{F},\mathscr{G})$ be a DR Game in which $\mathscr{G}$ is the set of distributions, G, on $[a,\infty)$ such that $\mathbb{E}[\chi(Y)] \le \mu$ when $Y \sim G \in \mathscr{G}$. Here $\chi(y)$ is a given strictly increasing continuous function on $[a,\infty)$, where $\chi(a) < \mu = \chi(k) < \chi(\infty)$. Suppose player I has already fixed his distribution function $F' \in \mathscr{F}$ with associated score random variable $X \sim F'$. F' can be written as a function $F(\chi)$ of χ (i.e. F is defined by $F(x) = \mathbb{P}(\chi(X) \le x)$ or by $F = F' \circ \chi^{-1}$) for the interval of χ-values $[\chi(a),\chi(\infty))$.*

1. *Let $\bar{F}$ be the concavification of F on $[\chi(a),\chi(\infty))$. Then player II cannot achieve more than $\bar{F}(\mu)$, but can either achieve this or else has strategies which will give her an expected payoff arbitrarily close to it.*

2. *Suppose now that player II is restricted to $[a,b]$ instead of $[a,\infty)$. Let F^* be F restricted to the domain $[\chi(a),\chi(b)]$, except that if F has a jump of m at $\chi(b)$ (i.e. an atom of probability m at $\chi(b)$), F^* has a jump of $m/2$ there. Let $\bar{F}$ be the concavification of F^* on $[\chi(a),\chi(b)]$. Then player II cannot achieve more than $\bar{F}(\mu)$, but can either achieve this or else has strategies which will give her an expected payoff arbitrarily close to it.*

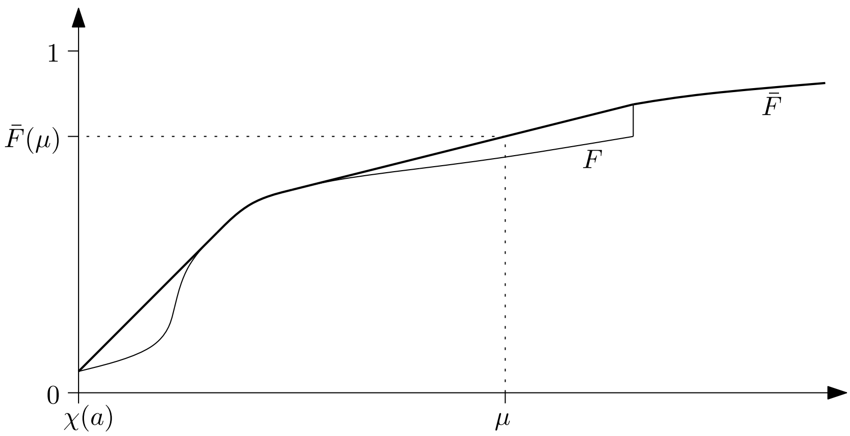

Figure 2: F and its concavification, $\bar{F}$

Proof. 1. Consider two-atom strategies for player II. One atom must be placed at, say, c so that $u = \chi(c) < \mu$, and the other at, say, d so that $\mu < \chi(d) = v$. The weighted average of u and v must be no more than μ, so the maximum probability we can assign to d is $(\mu - u)/(v - u)$. The probability that player II wins or draws is then

$$\frac{v-\mu}{v-u}F(u) + \frac{\mu-u}{v-u}F(v).$$

This has a simple geometric interpretation. Draw the graph of F against χ, and draw the line joining the points on the graph at u and v – i.e. the line joining $(u, F(u))$ to $(v, F(v))$. The height of this line at μ is the required probability. Since this includes the possibility of ties, it gives an upper bound for player II's payoff. Hence $\bar{F}(\mu)$ is an upper bound to what player II can obtain by two atom pure strategies.

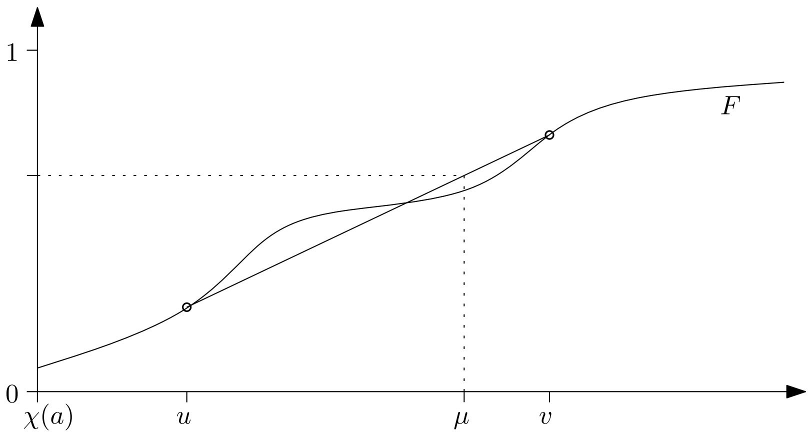

Figure 3: Line joining $(u, F(u))$ to $(v, F(v))$

Now any non-negative distribution with mean μ can be expressed as a mixture of two-point distributions with mean μ (see, for example, Pinelis (2009)), and so any mixture of pure strategies is also a mixture of two-point distributions. The expectation of a set of values cannot exceed the least upper bound of the values, and so $\bar{F}(\mu)$ is an upper bound to what player II can obtain by any pure or mixed strategies. Can she achieve this bound?

There will be a line L above $\bar{F}$ which touches it at μ. If this line also touches F at μ (i.e. if $\bar{F}(\mu) = F(\mu)$), player II can place an atom of probability 1 at $\chi = \mu$ (i.e at $y = k$) and either win or draw with probability $F(\mu) = \bar{F}(\mu)$. She will draw with positive probability only if player I has also placed an atom at μ, but in this case she can place an atom of probability $\frac{\mu-\chi(a)}{\mu-\chi(a)+\varepsilon}$ at $\chi = \mu + \varepsilon$, with the remaining probability at $\chi(a)$. Since ε can be arbitrarily small, she can then obtain an expected payoff arbitrarily close to $\bar{F}(\mu)$.

Alternatively, if L does not touch F at μ (but passes above it), it must touch F at one or more points below μ and also at one or more points above. Let u and v be two such touching points, $0 \le u < \mu < v$. Player II can consider placing two atoms of probability at u and v with mean μ.

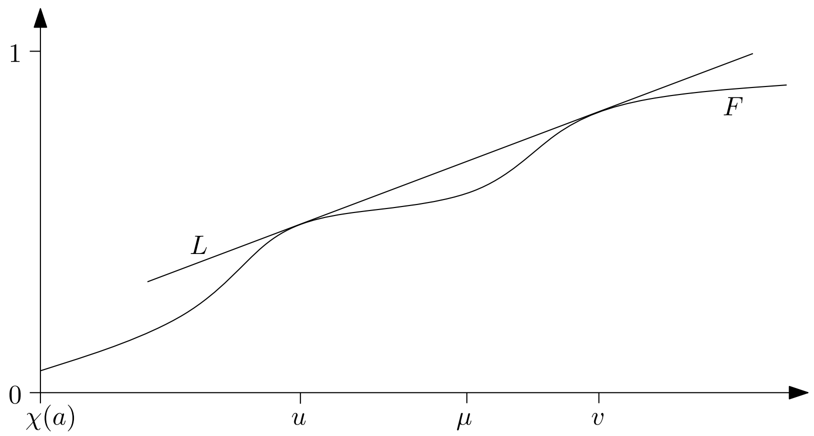

Figure 4: Line L tangent to $\bar{F}$ at μ.

The probability that player II either wins or draws is then

$$
\begin{aligned}
&= \frac{v-\mu}{v-u} F(u) + \frac{\mu-u}{v-u} F(v), \\
&= \bar{F}(\mu).
\end{aligned}
$$

She will draw with positive probability only if player I has placed an atom at u or v or both. But she can then place her atoms at $u + \varepsilon$ and $v + \delta$ for arbitrarily small ε and δ, and so obtain an expected payoff arbitrarily close to $\bar{F}(\mu)$.

2. The proof for the finite interval $[a, b]$ is similar, but must allow for the fact that if player I places

an atom of probability of mass m at $\chi(b)$, player II achieves only $F(b) - m/2$ when she plays there.

□

Corollary 4.2.1. *The solutions found to the DR Game with generalised constraints (Theorem 4.1) are unique.*

Proof. Returning to the notation of Theorem 4.1, Figure 5 shows a parametric plot of F against ψ.

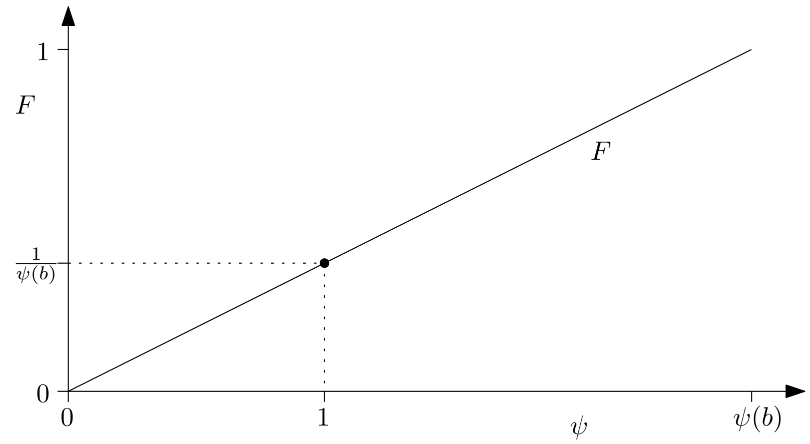

Figure 5: Parametric plot of $\big(\psi(x), F(x)\big)$

As expected, this graph (which is already concavified) is a line passing through $(1, 1/\psi(b))$ (marked with a circle in the figure) i.e. through $(\mu, \bar{F}(\mu))$ in the notation of Theorem 4.2.

Suppose F^* is another distribution which gives the same guarantee to player I. Suppose $Y \sim G$. If player II is playing on $[0, b]$ and F^* placed any probability to the right of b, player I could improve by moving all this probability to b, and scaling up F^* within the interval. So F^* must also have support $[0, b]$, and we can write F^* as a function of ψ.

Clearly, if $F^*(x) > F(x)$ for some $x \in [0, b]$, the concavification of F^* will pass above the point $(1, 1/\psi(b))$, so that player II will be able to obtain more than $1/\psi(b)$. (This holds even if $b = \infty$.) Hence $F^*(x) \leq F(x)$ for all $x \in [0, b]$. But if $X \sim F$, $\mathbb{E}\big[\phi(X)\big] = 1$. So if $X \sim F^*$ and for some x, $F^*(x) < F(x)$, we would have $\mathbb{E}\big[\phi(X)\big] > 1$ (ϕ is strictly increasing). So $F^* = F$.

Finally, suppose G^* is another distribution which gives the same guarantee to player II as G. As above, G^* must have the same support as G, and can be written as function of ϕ. Also, the concavification of G^* (expressed as a function of ϕ) must pass through or below the point $(1, 1-1/\psi(b))$. (Figure 6 shows how G passes through this point, marked with a square.). So there must be a line through the point which lies above (or touches) G^* for all $\phi \geq 0$. The dashed line in Figure 6 shows a

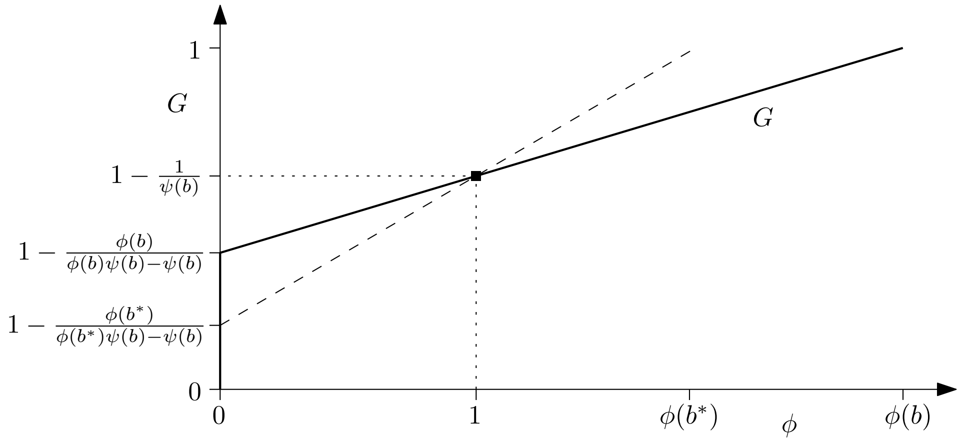

Figure 6: Parametric plot of $\big(\phi(x), G(x)\big)$

possible supporting line. If $Y \sim G^*$, we must have $\mathbb{E}\left[\psi(Y)\right] \leq 1$. Taking the dashed line as a distribution function gives a lower bound for $\mathbb{E}\left[\psi(Y)\right]$, so we take H to be the distribution function corresponding to the dashed line and find

$$\begin{aligned}
\mathbb{E}\left[\psi(Y)\right] &= \int_0^{b^*} \psi(x)\, dH(x) \\
&= \frac{1}{\phi(b^*)\psi(b) - \psi(b)} \int_0^{b^*} \psi(x)\, d\phi(x)\,.
\end{aligned}$$

Differentiating with respect to $\phi(b^*)$ gives

$$\frac{\psi(b)\psi(b^*)(\phi(b^*)-1) - \psi(b)\int_0^{b^*} \psi(x)\, d\phi(x)}{\left[\psi(b)(\phi(b^*)-1)\right]^2}.$$

This is negative if

$$\begin{aligned}
\psi(b^*)(\phi(b^*)-1) &< \int_0^{b^*} \psi(x)\, d\phi(x) \\
\phi(b^*)\psi(b^*) - \psi(b^*) &< \phi(b^*)\psi(b^*) - \int_0^{b^*} \phi(x)\, d\psi(x) \\
\int_0^{b^*} \phi(x)\, d\psi(x) &< \psi(b^*).
\end{aligned}$$

But this inequality holds for $b^* < b$, so if $Y \sim H$, $\mathbb{E}\left[\psi(Y)\right]$ decreases as b^* increases until $b^* = b$, at which value $H = G$ with expectation 1. So the supporting line must coincide with the distribution function G (taken as function of ϕ for $\phi > 0$). Then $G^*(x) \leq G(x)$ for all x, and as with F^* and F, we must have $G^* = G$. □

Corollary 4.2.2. *Suppose both players must choose distributions on the same finite interval* $[a, b]$ *but having (possibly) different means in* (a, b). *Then the game has the same value as the game with the same means played on* $[a, \infty)$. *The upper bound does not alter the balance of advantage.*

Proof. We can rescale so that $a = 0$ and the larger mean (player I) is 1. Let the smaller mean then be $\mu < 1$. If $b \geq 2$ (or in fact if $b = \infty$), Corollary 4.1.3 gives the solution. If $1 < b < 2$, Figures 7 and 8 give the solution (the diagrams have different values for b and μ). Player I places probability $b - 1$ in a

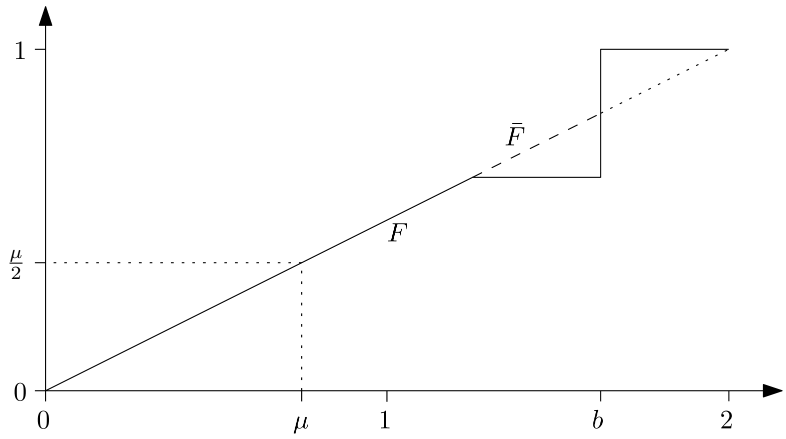

Figure 7: Strategy for player I on $[0, b]$ with larger mean

uniform distribution on $[0, 2(b - 1)]$, and probability $2 - b$ at b. This holds player II to $\frac{\mu}{2}$. Player II

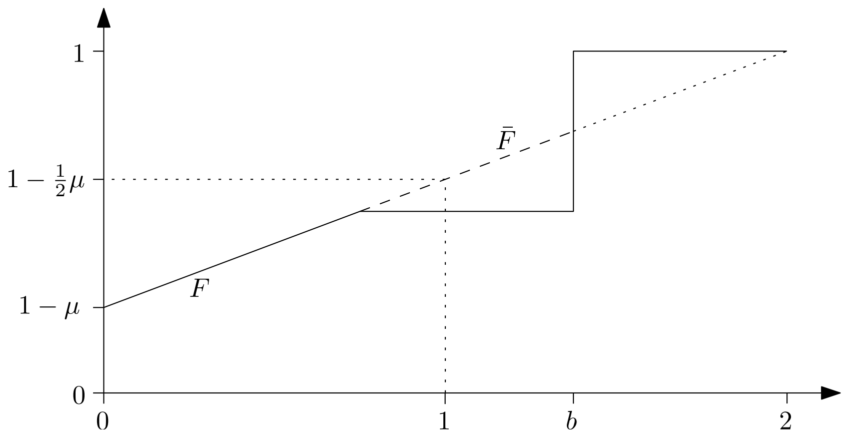

Figure 8: Strategy for player II on $[0, b]$ with smaller mean

chooses a distribution which places probability $1 - \mu$ at 0, $(2 - b)\mu$ at b, and spreads the remaining probability $(b - 1)\mu$ uniformly over $[0, 2(b - 1)]$. This holds player I to $1 - \frac{\mu}{2}$. □

5 The silent duel

As an application of our solution procedure for DR games, we simplify and shorten the original solution, e.g. Karlin (1959), to the classical problem of the silent duel. A modern treatment is given in Owen (2013). Unlike our previous examples, the constraints on the players' distributions $\mathscr{F}$ are not on their means or higher moments but are bounds on the integral

$$\int_0^\infty \left(\frac{x}{1-x}\right) dF(x) \, .$$

So this example illustrates the importance of the general approach taken in Section 4.

Suppose that two equally skilled competitors attempt a task whose level of difficulty can be varied – for example lifting a heavy barbell (vary the weight), jumping over a bar (vary the height), or proving a mathematical theorem (vary the theorem). Each competitor chooses his own level of difficulty and is allowed only one attempt, not knowing what level the other is attempting nor whether he succeeds or fails. The player who is successful at the highest level wins, or shares the prize if there is a tie. Of course, this protocol does not follow normal weight-lifting or high-jumping competition rules. However, it does give a simple model of a research tournament where two entrants compete to find the best solution to a problem posed by a firm – better solutions being more difficult to find. This classical zero-sum 'silent duel' game was formulated and solved by researchers at RAND around 1948-1952. (The solution is given in Karlin (1959).) The story involved two antagonistic duellists walking towards each other. A more friendly formulation has two equally skilled duellists approaching targets at which they may silently fire at distances of their own choice. The probability of hitting the target decreases with its distance. The winner, who gets a unit prize, is the duellist who hits his target at the greatest distance; if both miss, they share the prize (each gets a 'consolation prize' of one half). If they hit simultaneously it is a tie, with the (posthumous) winner determined by a fair coin. When should they fire?

If the probability of success increases continuously from 0 to 1 as they approach each other, we can measure distance by the probability of missing. Then a player who fires at distance $y \in [0, 1]$ hits with probability $1 - y$ and misses with probability y. Suppose that we introduce a score variable, x: if a player hits at distance x, he is given a score of x, and if he misses he is not given any score.

Suppose that a duellist has a continuous distribution for the range y at which he fires, having a density $r(y)$ on the interval $[0,b]$. His overall probability of missing is

$$p = \int_0^b y r(y)\, dy.$$

The distribution of his score x, conditional on his not missing, then has the density function

$$f(x) = \frac{(1-x)r(x)}{1-p}.$$

But now

$$\int_0^b \left(\frac{x}{1-x}\right) f(x)\, dx = \frac{1}{1-p}\int_0^b x r(x)\, dx = \frac{p}{1-p}.$$

So the (cumulative) score distribution F (given that he does not miss) must belong to the set

$$\mathscr{F}_p = \left\{F : \int_0^1 \chi(x)\, dF(x) = \chi(p)\right\}, \quad \text{where } \chi(x) = \frac{x}{1-x}.$$

We now show how the silent duel can be solved as a DR game. We imagine the game is played as follows. The first player chooses a probability of missing, p, and a compatible score distribution $F \in \mathscr{F}_p$. Similarly the second player chooses q and $F^* \in \mathscr{F}_q$. Then a referee samples independent events with probabilities p and q to see who has missed. If both players miss they each get a prize z which we take as $1/2$. If one misses and the other hits, the one who hits gets 1. Finally, if neither misses, a DR game is played using the distributions F and F^*. Clearly at an equilibrium, both players will play optimally in the DR game.

Corollary 4.1.1 gives the solution to the DR game, taking χ in the corollary as the function $\chi(x) = x/(1-x)$, and letting $k_1 = \chi(p)$ and $k_2 = \chi(q)$. (Although $\chi(x)$ goes to infinity on $[0,1)$ rather than $[0,\infty)$, the proof of Theorem 4.1 and the corollary easily generalise to this case.) If $p > q$ player I will be dominant and we find

$$\begin{aligned} b &= \frac{2p}{1+p}, \quad \text{and} \\ F(x) &= \left(\frac{1-p}{2p}\right)\left(\frac{x}{1-x}\right) \quad \text{for } 0 \le x \le b. \end{aligned}$$

Also, the value of the DR game (for $p > q$) is

$$\frac{2p - q(1+p)}{2p(1-q)}.$$

Now we can calculate the value of the silent duel as a function of p and q if the players play optimally in the induced DR game. If $p > q$, and z is the utility to each player when they both miss, then the value to player I is given by

$$\begin{aligned} v(p,q) &= (z)pq + (1-p)q + (1-p)(1-q)\frac{2p - q(1+p)}{2p(1-q)}, \\ &= \left(z - \frac{1}{2}\right)pq + 1 - p + q - \frac{q}{2p}, \\ &= 1 - p + q - \frac{q}{2p} \quad \text{for the constant-sum case } z = 1/2. \end{aligned}$$

Also (by interchanging the players) if $p < q$ it will be $(z - 1/2)pq + q - p + p/(2q)$ in general and $q - p + p/(2q)$ in the constant-sum case.

Assuming $z = 1/2$, we see player I's best response to $q < 1/2$ is $p = \sqrt{q/2}$, while if $q > 1/2$, his best response is 0 (shoot at point-blank range). The only equilibrium occurs at $p = q = 1/2$ (if $q = 1/2$, any $p \leq 1/2$ is a best response for player I, and similarly for player II). Hence $b = 2/3$ and for $0 \leq x \leq b$ we have the symmetric solution to the silent duel with

$$\begin{aligned} F(x) &= \frac{x}{2(1-x)} \quad \text{as the optimal score distribution,} \\ f(x) &= \frac{1}{2(1-x)^2} \quad \text{as the optimal score density, and} \\ r(x) &= \left(\frac{1-p}{1-x}\right) f(x) = \frac{1}{4(1-x)^3} \quad \text{as the density of the firing distance.} \end{aligned}$$

So a player should fire at distance y from the target for $y \in [0, 2/3]$ with probability density $r(y) = 1/\left[4(1-y)^3\right]$. He should never fire while $y > 2/3$.

6 Future work

This work could be extended in various ways. To give just three examples, the paper considers situations where there is only one constraint on the distributions, but there could be several. Secondly, we have not looked at sequential DR games, where the final outcome may depend on the sum of intermediate scores (as in the two halves of a football match), or the maximum of a number of attempts (as in weightlifting). Thirdly, the objective may not be to obtain first place in a league table, but rather to retain a position in (say) the top three. Also, even within the topics of this paper, it should be possible to extend Theorem 4.1 to the many-player case.

We feel that this is an exciting new area of research in game theory.

Acknowledgment

Steve Alpern wishes to acknowledge support from the Air Force Office of Scientific Research under grant FA9550-14-1-0049.

Appendix: Solution possibilities

We note that DR games are not guaranteed to have any pure strategy equilibria. We illustrate this for symmetric two-player games with a finite support set (e.g. biassing a die).

non-existence of equilibrium For this, we consider the (non convex) set of three unbiassed dice $\{D, E, F\}$ where the faces of D have two 1's and four 4's, E has all 3's and F has four 2's and two 5's. If one player chooses D, the other can win more than half the time by choosing F: we write $F \succ D$. Then $D \succ E \succ F \succ D$, and the non-transitivity precludes a maximal element. The symmetric game with this strategy set has no equilibrium and no value. However if mixed strategies are allowed, the strategy set becomes the convex hull of D, E, and F, and then the solution is to play the pure strategies with probabilities $3/7$, $1/7$, and $3/7$, which is equivalent to a seven sided die G in which the numbers 1, 3, and 5 appear once and 2 and 4 appear twice. Note that all four dice have mean 3 and that G wins half the time (after adjudicating ties) against D, E, and F. These non-transitive dice give an example of the Steinhaus-Trybula paradox, see Steinhaus and Trybula (1959).

existence and uniqueness of equilibrium Consider the set of all dice with distributions on $\{1, 2, 3, 4, 5, 6\}$ having the usual mean of 3.5. It can be shown that the unique optimal solution to this symmetric game is the usual unbiassed die which takes on each of the six numbers equiprobably.

existence and non-uniqueness of equilibrium To obtain non-uniqueness we simply modify the above example to 7 sided dice on $\{0, 1, 2, 3, 4, 5, 6\}$ with a mean of 3. Possible solutions are then uniform $\left(\frac{1}{7}, \frac{1}{7}, \ldots, \frac{1}{7}\right)$, or uniform on the even numbers $\left(\frac{1}{4}, 0, \frac{1}{4}, 0, \frac{1}{4}, 0, \frac{1}{4}\right)$, or uniform on the odd numbers $\left(0, \frac{1}{3}, 0, \frac{1}{3}, 0, \frac{1}{3}, 0\right)$. Computational work on discrete DR games has been carried out by Apicella (2013).

A fourth possibility can occur with continuous support sets and asymmetric players.

existence of value with an ε-equilbrium only Suppose both players must select a distribution with mean 1, but player I is confined to the support set $[0, 3/2]$ while player II is restricted to $[0, 5/4]$. Player I has ε-optimal strategies in which he places an atom at $5/4 + \varepsilon$, avoiding the possibility of a tied score of $5/4$. But there is no equilibrium.

References

ANDERSON, E. (2012): "Ranking games and gambling: when to quit when you're ahead," *Oper. Res.*, 60, 1229–1244.

APICELLA, T. (2013): "Strategic maximisation of discrete distributional fixed mean ranking games," Master's thesis, University of Warwick, supervised by Steve Alpern.

BELL, R. AND T. M. COVER (1988): "Game-theoretic optimal portfolios," *Management Science*, 34, 724–733.

BELL, R. M. AND T. M. COVER (1980): "Competitive optimality of logarithmic investment," *Mathematics of Operations Research*, 5, 161–166.

GERCHAK, Y. AND D. M. KILGOUR (1993): "Sequential Competitions with Nondecreasing Levels of Difficulty," *Oper. Res. Lett.*, 13, 49–58.

GOLDBERG, L. A., P. W. GOLDBERG, P. KRYSTA, AND C. VENTRE (2013): "Ranking games that have competitiveness-based strategies," *Theoret. Comput. Sci.*, 476, 24–37.

KARLIN, S. (1959): *The Theory of Infinite Games*, Addison-Wesley series in statistics, Addison-Wesley.

OWEN, G. (2013): *Game theory*, Emerald Group Publishing Limited, Bingley, fourth ed.

PINELIS, I. (2009): "Optimal two-value zero-mean disintegration of zero-mean random variables," *Electron. J. Probab.*, 14, no. 26, 663–727.

STEINHAUS, H. AND S. TRYBULA (1959): "On a paradox in applied probabilities," *Bull. Acad. Polon. Sci*, 7, 67–69.