

Distributional Robust Kelly Gambling

Qingyun Sun · Stephen Boyd

Abstract In classic Kelly gambling, bets are chosen to maximize the expected log growth of wealth, under a known probability distribution. In this note we consider the distributional robust version of the Kelly gambling problem, in which the probability distribution is not known, but lies in a given set of possible distributions. The bet is chosen to maximize the worst-case (smallest) expected log growth among the distributions in the given set. This distributional robust Kelly gambling problem is convex, but in general need not be tractable. We show that it can be tractably solved in a number of useful cases when there is a finite number of outcomes.

1 Introduction

Gambling. We consider a setting where a gambler repeatedly allocates a fraction of her wealth (assumed positive) across n different bets in multiple rounds. We assume there are n bets available to the gambler, who can bet any nonnegative amount on each of the bets. We let $b \in \mathbf{R}^n$ denote the bet allocation, so $b \geq 0$ and $\mathbf{1}^T b = 1$, where $\mathbf{1}$ is the vector with all entries one. Letting S_n denote the probability simplex in \mathbf{R}^n , we have $b \in S_n$. With bet allocation b , the gambler is betting Wb_i (in dollars) on outcome i , where $W > 0$ is the gambler's wealth (in dollars).

We let $r \in \mathbf{R}_+^n$ denote the random returns on the n bets, with $r_i \geq 0$ the amount won by the gambler for each dollar she puts on bet i . With allocation b , the total she wins is $r^T b W$, which means her wealth increases by the (random) factor $r^T b$. We assume that the returns r in different rounds are IID. We will assume that $r_n = 1$ almost surely, so b_n corresponds to the fraction of wealth the gambler holds in cash; the allocation $b = e_n$ corresponds to not betting at all. Since her wealth is multiplied in each round by the IID random factor $r^T b$, the log of the wealth over time is therefore a random walk, with increment distribution given by $\log(r^T b)$.

Finite outcome case. We consider here the case where one of K events occurs, *i.e.*, r is supported on only K points. We let r_1, \dots, r_K denote the return vectors, and $\pi = (\pi_1, \dots, \pi_K) \in S_K$ the corresponding probabilities. We collect the K payoff vectors into a matrix $R \in \mathbf{R}^{n \times K}$, with columns r_1, \dots, r_K . The vector $R^T b \in \mathbf{R}^K$ gives the wealth growth factor in the K possible outcomes. The mean log growth rate is

$$G_\pi(b) = \mathbf{E} \log(r^T b) = \pi^T \log(R^T b) = \sum_{k=1}^K \pi_k \log(r_k^T b),$$

where the log in the middle term is applied to the vector elementwise. This is the mean drift in the log wealth random walk.

Qingyun Sun
Stanford University
E-mail: qysun@stanford.edu

Stephen Boyd
Stanford University
E-mail: boyd@stanford.edu

Kelly gambling. In a 1956 classic paper [1], John Kelly proposed to choose the allocation vector b so as to maximize the mean log growth rate $G_\pi(b)$, subject to $b \geq 0$, $\mathbf{1}^T b = 1$. This method was called the Kelly criterion; since then, much work has been done on this topic [2, 3, 4, 5, 6, 7]. The mean log growth rate $G_\pi(b)$ is a concave function of b , so choosing b is a convex optimization problem [8, 9]. It can be solved analytically in simple cases, such as when there are $K = 2$ possible outcomes. It is easily solved in other cases using standard methods and algorithms, and readily expressed in various domain specific languages (DSLs) for convex optimization such as CVX [10], CVXPY [11, 12], Convex.jl [13], or CVXR [14]. We can add additional convex constraints on b , which we denote as $b \in B$, with $B \subseteq S_K$ a convex set. These additional constraints preserve convexity, and therefore tractability, of the optimization problem. Kelly did not consider additional constraints, or indeed the use of a numerical optimizer to find the optimal bet allocation vector, we still refer to the problem of maximizing $G_\pi(b)$ subject to $b \in B$ as the Kelly (gambling) problem (KP).

There have been many papers exploring and extending the Kelly framework; for example, a drawdown risk constraint, that preserves convexity (hence, tractability) is described in [15]. In [16], it is proposed to learn the constraint set from data. The Bayesian version of Kelly optimal betting is described in [17]. The Bayesian framework is extended to the multi-agent game theory setting in [18]. In [19], Kelly gambling is generalized to maximize the proportion of wealth relative to the total wealth in the population.

Distributional robust Kelly gambling. In this note we study a distributional robust version of Kelly gambling, in which the probability distribution π is not known. Rather, it is known that $\pi \in \Pi$, a set of possible distributions. We define the worst-case log growth rate (under Π) as

$$G_\Pi(b) = \inf_{\pi \in \Pi} G_\pi(b).$$

This is evidently a concave function of b , since it is an infimum of a family of concave functions of b , *i.e.*, $G_\pi(b)$ for $\pi \in \Pi$. The *distributional robust Kelly problem* (DRKP) is to choose $b \in B$ to maximize $G_\Pi(b)$. This is in principle a convex optimization problem, specifically a distributional robust problem; but such problems in general need not be tractable, as discussed in [20, 21, 22]. The purpose of this note is to show how the DRKP can be tractably solved for some useful probability sets Π .

Related work on uncertainty aversion. In decision theory and economics, there are two important concepts, risk and uncertainty. Risk is about the situation when a probability can be assigned to each possible outcome of a situation. Uncertainty is about the situation when the probabilities of outcomes are unknown. Uncertainty aversion, also called ambiguity aversion, is a preference for known risks over unknown risks. Uncertainty aversion provides a behavioral foundation for maximizing the utility under the worst of a set of probability measures; see [23, 24, 25, 26] for more detailed discussion. The Kelly problem addresses risk; the distributional robust Kelly problem is a natural extension that considers uncertainty aversion.

Related work on distributional robust optimization. Distributional robust optimization is a well studied topic. Previous work on distribution robust optimization studied finite-dimensional parametrization for probability sets including moments, support or directional deviations constraints in [27, 28, 29, 30, 31, 32]. Beyond finite-dimensional parametrization of the probability set, researchers have also studied non-parametric distances for probability measure, like f -divergences (*e.g.*, Kullback-Leibler divergences) [33, 34, 35, 36, 37] and Wasserstein distances [38, 39, 40, 41].

2 Tractable distributional robust Kelly gambling

In this section we show how to formulate DRKP as a tractable convex optimization problem for a variety of distribution sets. The key is to derive a tractable description of the worst-case log growth $G_\Pi(b)$. We use duality to express $G_\Pi(b)$ as the value of a convex maximization problem, which allows us to solve DRKP as one convex problem. (This is a standard technique in robust optimization in general.)

2.1 Polyhedral distribution sets

Here we consider the case when Π is a polyhedron.

Convex hull of finite set. We start with the simplest example, when Π is a polyhedron defined as the convex hull of a finite set of points, $\Pi = \mathbf{conv}(\pi, \dots, \pi_s)$. The infimum of the log growth over $\pi \in \Pi$ is the same as the minimum over the vertices:

$$G_{\Pi}(b) = \min_i G_{\pi_i}(b) = \min_i \pi_i^T \log(R^T b).$$

Then the DRKP becomes

$$\begin{aligned} & \text{maximize } \min_i \pi_i^T \log(R^T b) \\ & \text{subject to } b \in B, \end{aligned}$$

with variable b . This problem is tractable, and indeed in modern domain specific languages for convex optimization, can be specified in just a few lines of simple code.

Polyhedron defined by linear inequalities and equalities. Here we consider the case when Π is given by a finite set of linear inequalities and equalities,

$$\Pi = \{\pi \in S_K \mid A_0 \pi = d_0, A_1 \pi \leq d_1\},$$

where $A_0 \in \mathbf{R}^{m_0 \times K}$, $b_0 \in \mathbf{R}^{m_0}$, $A_1 \in \mathbf{R}^{m_1 \times K}$, $b_1 \in \mathbf{R}^{m_1}$. The worst-case log growth rate $G_{\Pi}(b)$ is given by the optimal value of the linear program (LP)

$$\begin{aligned} & \text{minimize } \pi^T \log(R^T b) \\ & \text{subject to } \mathbf{1}^T \pi = 1, \quad \pi \geq 0, \\ & \quad \quad \quad A_0 \pi = d_0, \quad A_1 \pi \leq d_1, \end{aligned} \tag{1}$$

with variable π .

We form a dual of this problem, working with the constraints $A_0 \pi = d_0$, $A_1 \pi \leq d_1$; we keep the simplex constraints $\pi \geq 0$, $\mathbf{1}^T \pi = 1$ as an indicator function $I_S(\pi)$ in the objective. The Lagrangian is

$$L(\nu, \lambda, \pi) = \pi^T \log(R^T b) + \nu^T (A_0 \pi - d_0) + \lambda^T (A_1 \pi - d_1) + I_S(\pi),$$

where $\nu \in \mathbf{R}^{m_0}$ and $\lambda \in \mathbf{R}^{m_1}$ are the dual variables, with $\lambda \geq 0$. Minimizing over π we obtain the dual function,

$$g(\nu, \lambda) = \inf_{\pi \in S_K} L(\nu, \lambda, \pi) = \min(\log(R^T b) + A_0^T \nu + A_1^T \lambda) - d_0^T \nu - d_1^T \lambda,$$

where the min of a vector is the minimum of its entries. The dual problem associated with (1) is then

$$\begin{aligned} & \text{maximize } \min(\log(R^T b) + A_0^T \mu + A_1^T \lambda) - d_0^T \mu - d_1^T \lambda \\ & \text{subject to } \lambda \geq 0, \end{aligned}$$

with variables μ, λ . This problem has the same optimal value as (1), *i.e.*,

$$G_{\Pi}(b) = \sup_{\nu, \lambda \geq 0} \left(\min(\log(R^T b) + A_0^T \mu + A_1^T \lambda) - d_0^T \mu - d_1^T \lambda \right).$$

Using this expression for $G_{\Pi}(b)$, the DRKP becomes

$$\begin{aligned} & \text{maximize } \min(\log(R^T b) + A_0^T \mu + A_1^T \lambda) - d_0^T \mu - d_1^T \lambda \\ & \text{subject to } b \in B, \quad \lambda \geq 0, \end{aligned}$$

with variables b, μ, λ . This is a tractable convex optimization problem, readily expressed in domain specific languages for convex optimization.

Box distribution set. We consider a special case of a polyhedral distribution set, with lower and upper bounds for each π_k :

$$\Pi = \{\pi \in S_K \mid |\pi - \pi^{\text{nom}}| \leq \rho\},$$

where $\pi^{\text{nom}} \in S_K$ is the nominal distribution, and $\rho \in \mathbf{R}_+^n$ is a vector of radii. (The inequality $|\pi - \pi^{\text{nom}}| \leq \rho$ is interpreted elementwise.)

Using the general method above, expressing the limits as $A_1\pi \leq d_1$ with

$$A_1 = \begin{bmatrix} I \\ -I \end{bmatrix}, \quad d_1 = \begin{bmatrix} \pi^{\text{nom}} + \rho \\ \rho - \pi^{\text{nom}} \end{bmatrix},$$

the DRKP problem becomes

$$\begin{aligned} & \text{maximize} \quad (\min(\log(R^T b) + \lambda_+ - \lambda_-) - (\pi^{\text{nom}})^T(\lambda_+ - \lambda_-) - \rho^T(\lambda_+ + \lambda_-)) \\ & \text{subject to} \quad b \in B, \quad \lambda_+ \geq 0, \quad \lambda_- \geq 0, \end{aligned}$$

with variables b, λ_+, λ_- . Defining $\lambda = \lambda_+ - \lambda_-$, we have $|\lambda| = \lambda_+ + \lambda_-$, so the DRKP becomes

$$\begin{aligned} & \text{maximize} \quad \min(\log(R^T b) + \lambda) - (\pi^{\text{nom}})^T \lambda - \rho^T |\lambda| \\ & \text{subject to} \quad b \in B, \end{aligned}$$

with variables b, λ .

2.2 Ellipsoidal distribution set

Here we consider the case when Π is the inverse image of a p -norm ball, with $p \geq 1$, under an affine mapping. This includes an ellipsoid (and indeed the box set described above) as a special case. We take

$$\Pi = \{\pi \in S_K \mid \|W^{-1}(\pi - \pi^{\text{nom}})\|_p \leq 1\},$$

where W is a nonsingular matrix. As usual we define q by $1/p + 1/q = 1$.

We define $x = -\log(R^T b)$, $z = W^{-1}(\pi - \pi^{\text{nom}})$, and $D_{p,W} = \{z \mid \|z\|_p \leq 1, \mathbf{1}^T Wz = 0, \pi^{\text{nom}} + Wz \geq 0\}$. Then we have

$$\begin{aligned} G_\Pi(b) &= -\sup_{\pi \in \Pi} ((\pi - \pi^{\text{nom}})^T x + (\pi^{\text{nom}})^T x) \\ &= -\sup_{z \in D_{p,W}} z^T W^T x + (\pi^{\text{nom}})^T x \\ &= \sup_{\mu, \lambda \geq 0} (-\sup_{\|z\|_p \leq 1} z^T W^T (x + \lambda - \mu \mathbf{1}) + (\pi^{\text{nom}})^T (\lambda + x)) \\ &= \sup_{\mu, \lambda \geq 0} (-\|W^T (x + \lambda - \mu \mathbf{1})\|_q + (\pi^{\text{nom}})^T (-\lambda - x)). \end{aligned}$$

Here the second last equation is the Lagrangian form where we keep the p -norm constraint as a convex indicator, and the last equation is based on the Hölder equality

$$\sup_{\|z\|_p \leq 1} z^T W^T (x + \lambda - \mu \mathbf{1}) = \|W^T (x + \lambda - \mu \mathbf{1})\|_q,$$

Using this expression for $G_\Pi(b)$, and let $u = -x - \lambda = \log(R^T b) - \lambda \leq \log(R^T b)$, then DRKP becomes

$$\begin{aligned} & \text{maximize} \quad (\pi^{\text{nom}})^T (u) - \|W^T (u - \mu \mathbf{1})\|_q \\ & \text{subject to} \quad u \leq \log(R^T b), \\ & \quad \quad \quad b \in B, \end{aligned}$$

with variables b, u, μ .

2.3 Divergence based distribution set

Let $\pi_1, \pi_2 \in S_K$ be two distributions. For a convex function $f : \mathbf{R}_+ \rightarrow \mathbf{R}$ with $f(1) = 0$, the f -divergence of π_1 from π_2 is defined as

$$D_f(\pi_1 \| \pi_2) = \pi_2^T f(\pi_1 / \pi_2),$$

where the ratio is meant elementwise. Recall that the Fenchel conjugate of f is $f^*(s) = \sup_{t \geq 0} (ts - f(t))$.

We can use the f -divergence from a nominal distribution to define a set of distributions. We take

$$\Pi = \{\pi \in S_K \mid D_f(\pi \| \pi^{\text{nom}}) \leq \epsilon\},$$

where $\epsilon > 0$ is a given value. (Such sets of distributions arise naturally when π^{nom} is the empirical distribution of a set of samples from π^{nom} , of a set of samples from a distribution.)

We define $x = -\log(R^T b)$ again. Our goal is to minimize $-G_\Pi(b) = \sup_{\pi \in \Pi} \pi^T x$. We form a dual of this problem, working with the constraints $D_f(\pi \| \pi_0) \leq \epsilon$ and $\mathbf{1}^T \pi = 1$; we keep the constraint $\pi \geq 0$ implicit. With dual variables $\lambda \in \mathbf{R}_+$, $\gamma \in \mathbf{R}$, then for $\pi \geq 0$, the Lagrangian is

$$L(\gamma, \lambda, \pi) = \pi^T x + \lambda(-(\pi^{\text{nom}})^T f(\frac{\pi}{\pi^{\text{nom}}}) + \epsilon) - \gamma(e^T \pi - 1) + I_+(\pi),$$

where I_+ is the indicator function of \mathbf{R}_+^K . The dual objective function is

$$\begin{aligned} \sup_{\pi \geq 0} L(\gamma, \lambda, \pi) &= \sup_{\pi \geq 0} (\sum_{i=1}^K \pi_i^{\text{nom}} (\frac{\pi_i}{\pi_i^{\text{nom}}} x_i - \frac{\pi_i}{\pi_i^{\text{nom}}} \gamma - \lambda f(\frac{\pi_i}{\pi_i^{\text{nom}}})) + \lambda \epsilon + \gamma) \\ &= \sum_{i=1}^K \pi_{0,i} \sup_{t_i \geq 0} (t_i (x_i - \gamma) - \lambda f(t_i)) + \lambda \epsilon + \gamma \\ &= \sum_{i=1}^K \pi_{0,i} \lambda f^*(\frac{x_i - \gamma}{\lambda}) + \lambda \epsilon + \gamma. \end{aligned}$$

We can write the problem as

$$\begin{aligned} &\text{maximize} \quad -(\pi^{\text{nom}})^T \lambda f^*(\frac{-\log(R^T b) - \gamma}{\lambda}) - \lambda \epsilon - \gamma \\ &\text{subject to} \quad \lambda \geq 0, \quad b \in B, \end{aligned}$$

with variables b, γ, λ . We transform the problem to follow the disciplined convex programming (DCP) rules by convex relaxation of the equality constraint. Now DRKP becomes

$$\begin{aligned} &\text{maximize} \quad -(\pi^{\text{nom}})^T w - \epsilon \lambda - \gamma \\ &\text{subject to} \quad w \geq \lambda f^*(\frac{z}{\lambda}) \\ &\quad \quad \quad z \geq -\log(R^T b) - \gamma \\ &\quad \quad \quad \lambda \geq 0, \quad b \in B, \end{aligned}$$

with variables b, γ, λ, w, z .

Here

$$\lambda f^*(\frac{z}{\lambda}) = (\lambda f)^*(z) = \sup_{t \geq 0} (tz - \lambda f(t)),$$

is the perspective function of the non-decreasing convex function $f^*(z)$, so it is also a convex function that is non-decreasing in z . Additionally, $-\log(R^T b) - \gamma$ is a convex function of b and γ ; then from the DCP composition rule, we know this form of DRKP is convex.

We remark that there is a one-parameter family of f -divergences generated by the α -function with $\alpha \in \mathbf{R}$, where we can define the generalization of natural logarithm by

$$\log_\alpha(t) = \frac{t^{\alpha-1} - 1}{\alpha - 1}.$$

For $\alpha \neq 1$, it is a power function, for $\alpha \rightarrow 1$, it is converging to the natural logarithm. Now if we assume $f_\alpha(1) = 0$ and $f'_\alpha(t) = \log_\alpha(t)$, then we have

$$f_\alpha(t) = \frac{t^\alpha - 1 - \alpha(t - 1)}{\alpha(\alpha - 1)}.$$

The Fenchel conjugate is

$$f_\alpha^*(s) = \frac{1}{\alpha} ((1 + (\alpha - 1)s)^{\frac{\alpha}{\alpha-1}} - 1).$$

We now show some more specific examples of f -divergences; for a more detailed discussion see [36].

- *KL-divergence*. With $f(t) = t \log(t) - t + 1$, we obtain the KL-divergence. We have $f^*(s) = \exp(s) - 1$. This corresponds to $\alpha = 1$.
- *Reverse KL-divergence*. With $f(t) = -\log(t) + t - 1$, the f -divergence is the reverse KL-divergence. We have $f^*(s) = -\log(1 - s)$ for $s < 1$. This corresponds to $\alpha = 0$.
- *Pearson χ^2 -divergence*. With $f(t) = \frac{1}{2}(t - 1)^2$, we obtain the Pearson χ^2 -divergence. We have $f^*(s) = \frac{1}{2}(s + 1)^2 - \frac{1}{2}$, $s > -11$. This corresponds to $\alpha = 2$.
- *Neyman χ^2 -divergence*. With $f(t) = \frac{1}{2t}(t - 1)^2$, we obtain the Neyman χ^2 -divergence. We have $f^*(s) = 1 - \sqrt{1 - 2s}$, $s < \frac{1}{2}$. This corresponds to $\alpha = -1$.
- *Hellinger-divergence*. With $f(t) = 2(\sqrt{t} - 1)^2$, we obtain the Hellinger-divergence. We have $f^*(s) = \frac{2s}{2-s}$, $s < 2$. This corresponds to $\alpha = -1$.
- *Total variation distance*. With $f(t) = |t - 1|$, the f -divergence is the total variation distance. We have $f^*(s) = -1$ for $s \leq -1$ and $f^*(s) = s$ for $-1 \leq s \leq 1$.

2.4 Wasserstein distance distribution set

The Wasserstein distance $D_c(\pi, \pi^{\text{nom}})$ with cost $c \in \mathbf{R}_+^{K \times K}$ is defined as the optimal value of the problem

$$\begin{aligned} & \text{minimize } \sum_{i,j} Q_{ij} c_{ij} \\ & \text{subject to } Q\mathbf{1} = \pi, \quad Q^T\mathbf{1} = \pi^{\text{nom}}, \quad Q \geq 0, \end{aligned}$$

with variable Q . The Wasserstein distance distribution set is

$$\Pi = \{\pi \in S_K \mid D_c(\pi, \pi^{\text{nom}}) \leq s\},$$

with $s > 0$. The Wasserstein distance has several other names, including Monge-Kantorovich, earth-mover, or optimal transport distance [38, 39, 40, 41].

The worst-case log growth $G_\Pi(b)$ is given by the value of the following LP,

$$\begin{aligned} & \text{minimize } \pi^T \log(R^T b) \\ & \text{subject to } Q\mathbf{1} = \pi, \quad Q^T\mathbf{1} = \pi^{\text{nom}}, \quad Q \geq 0, \\ & \quad \quad \quad \sum_{i,j} Q_{ij} c_{ij} \leq s, \end{aligned}$$

with variable Q . Using strong duality for LP, the DRKP becomes

$$\begin{aligned} & \text{maximize } \left(\sum_j \pi_j^{\text{nom}} \min_i (\log(R^T b)_i + \lambda c_{ij}) - s\lambda \right) \\ & \text{subject to } b \in B, \quad \lambda \geq 0. \end{aligned}$$

where $\lambda \in \mathbf{R}_+$ is the dual variable.

3 Numerical example

In this section we illustrate distributional robust Kelly gambling with an example. Our example is a simple horse race with n horses, with bets placed on each horse placing, *i.e.*, coming in first or second. There are thus $K = n(n - 1)/2$ outcomes (indexed as j, k with $j < k \leq n$), and n bets (one for each horse to place).

We first describe the nominal distribution of outcomes π^{nom} . We model the speed of the horses as independent random variables, with the fastest and second fastest horses placing. With this model, π^{nom} is entirely described by the probability that horse i comes in first, we which denote β_i . For $j < k$, we have

$$\begin{aligned} \pi_{jk}^{\text{nom}} &= P(\text{horse } j \text{ and horse } k \text{ are the first and second fastest}) \\ &= P(j \text{ is 1st, } k \text{ is 2nd}) + P(k \text{ is 1st, } j \text{ is 2nd}) \\ &= P(j \text{ is 1st})P(k \text{ is 2nd} \mid j \text{ is 1st}) + P(k \text{ is 1st})P(j \text{ is 2nd} \mid k \text{ is 1st}) \\ &= \beta_j(\beta_k/(1 - \beta_j)) + \beta_k(\beta_j/(1 - \beta_k)) \\ &= \beta_j\beta_k\left(\frac{1}{1-\beta_j} + \frac{1}{1-\beta_k}\right). \end{aligned}$$

The fourth line uses $P(k \text{ is 2nd} \mid j \text{ is 1st}) = \beta_k/(1 - \beta_j)$.

Growth rate	b^K	b^{RK}
π^{nominal}	4.3%	2.2%
π^{worst}	-2.2%	0.7%

Table 1: For box uncertainty set with $\eta = 0.26$, we compare the growth rate and worst-case growth rate for the Kelly optimal and the distributional robust Kelly optimal bets.

Growth rate	b^K	b^{RK}
π^{nominal}	4.3%	2.2%
π^{worst}	-2.2%	0.4%

Table 2: For ball uncertainty set with $c = 0.016$, we compare the growth rate and worst-case growth rate for the Kelly optimal and the distributional robust Kelly optimal bets.

For the return matrix, we use parimutuel betting, with the fraction of bets on each horse equal to β_i , the probability that it will win (under the nominal probability distribution). The return matrix $R \in \mathbf{R}^{n \times K}$ then has the form

$$R_{i,jk} = \begin{cases} \frac{n}{1+\beta_j/\beta_k} & i = j \\ \frac{n}{1+\beta_k/\beta_j} & i = k \\ 0 & i \notin \{j, k\}, \end{cases}$$

where we index the columns (outcomes) by the pair jk , with $j < k$.

Our set of possible distributions is the box

$$\Pi_\eta = \{\pi \mid |\pi - \pi^{\text{nom}}| \leq \eta \pi^{\text{nom}}, \mathbf{1}^T \pi = 1, \pi \geq 0\},$$

where $\eta \in (0, 1)$, *i.e.*, each probability can vary by η from its nominal value.

Another uncertainty set is the ball

$$\Pi_c = \{\pi \mid \|\pi - \pi^{\text{nom}}\|_2 \leq c, \mathbf{1}^T \pi = 1, \pi \geq 0\}$$

For our specific example instance, we take $n = 20$ horses, so there are $K = 190$ outcomes. We choose β_i , the probability distribution of the winning horse, proportional to $\exp z_i$, where we sample independently $z_i \sim \mathcal{N}(0, 1/4)$. This results in β_i ranging from around 20% (the fastest horse) to around 1% (the slowest horse).

First, we show growth rate and worst-case growth rate for the Kelly optimal and the distributional robust Kelly optimal bets under two uncertainty sets. In table 1, we show the comparison for box uncertainty set with $\eta = 0.26$; in table 2, and for ball uncertainty set with $c = 0.016$. The two parameters are chosen so that the worst case growth of Kelly bets for both uncertainty sets are -2.2% . In particular, using standard Kelly betting, we lose money (when the distribution is chosen as the worst one for the Kelly bets).

We can see that, as expected, the Kelly optimal bet has higher log growth under the nominal distribution, and the distributional robust Kelly bet has better worst-case log growth. We see that the worst-case growth of the distributional robust Kelly bet is significantly better than the worst-case growth of the nominal Kelly optimal bet. In particular, with robust Kelly betting, we make money, even when the worst distribution is chosen.

The nominal Kelly optimal bet b^K and the distributional robust Kelly bet b^{RK} for both uncertainty sets in figure 1. For each of our bets b^K and b^{RK} shown above, we find a corresponding worst case distribution, denoted $\pi^{\text{wc},K}$ and $\pi^{\text{wc},RK}$, which minimize $G_\pi(b)$ over $\pi \in \Pi$. These distributions, shown for box uncertainty set in figure 2 and for ball uncertainty set in figure 3, achieve the corresponding worst-case log growth for the two bet vectors.

Finally, we compare the expected wealth logarithmic growth rate as we increase the size of the uncertainty sets. For the box uncertainty set we choose $\eta \in [0, 0.3]$, and for the ball uncertainty set we choose $c \in [0, 0.02]$, we look at the expected growth for both the kelly bet b^K and the distributional robust kelly bet b^{RK} under both the nominal probability π^{nom} and the worst case probability π^{worst} . Specifically, in figure 4 we plot the four expected growth: $\pi^{\text{nom},T} \log(R^T b^K)$, $\pi_\eta^{\text{worst},T} \log(R^T b^K)$, $\pi^{\text{nom},T} \log(R^T b_\eta^{RK})$, $\pi_\eta^{\text{wc},T} \log(R^T b_\eta^{RK})$.




Fig. 1: The Kelly optimal bets b^K for the nominal distribution, and the distributional robust optimal bets for the box and ball uncertainty sets, ordered by the descending order of b^K .




Fig. 2: For box uncertainty set with $\eta = 0.26$, we show the nominal distribution π^{nom} (sorted) and the two worst-case distributions $\pi^{\text{wc,K}}$ and $\pi^{\text{wc,RK}}$.




Fig. 3: For ball uncertainty set with $c = 0.016$, we show the nominal distribution π^{nom} (sorted) and the two worst-case distributions $\pi^{\text{wc,K}}$ and $\pi^{\text{wc,RK}}$.

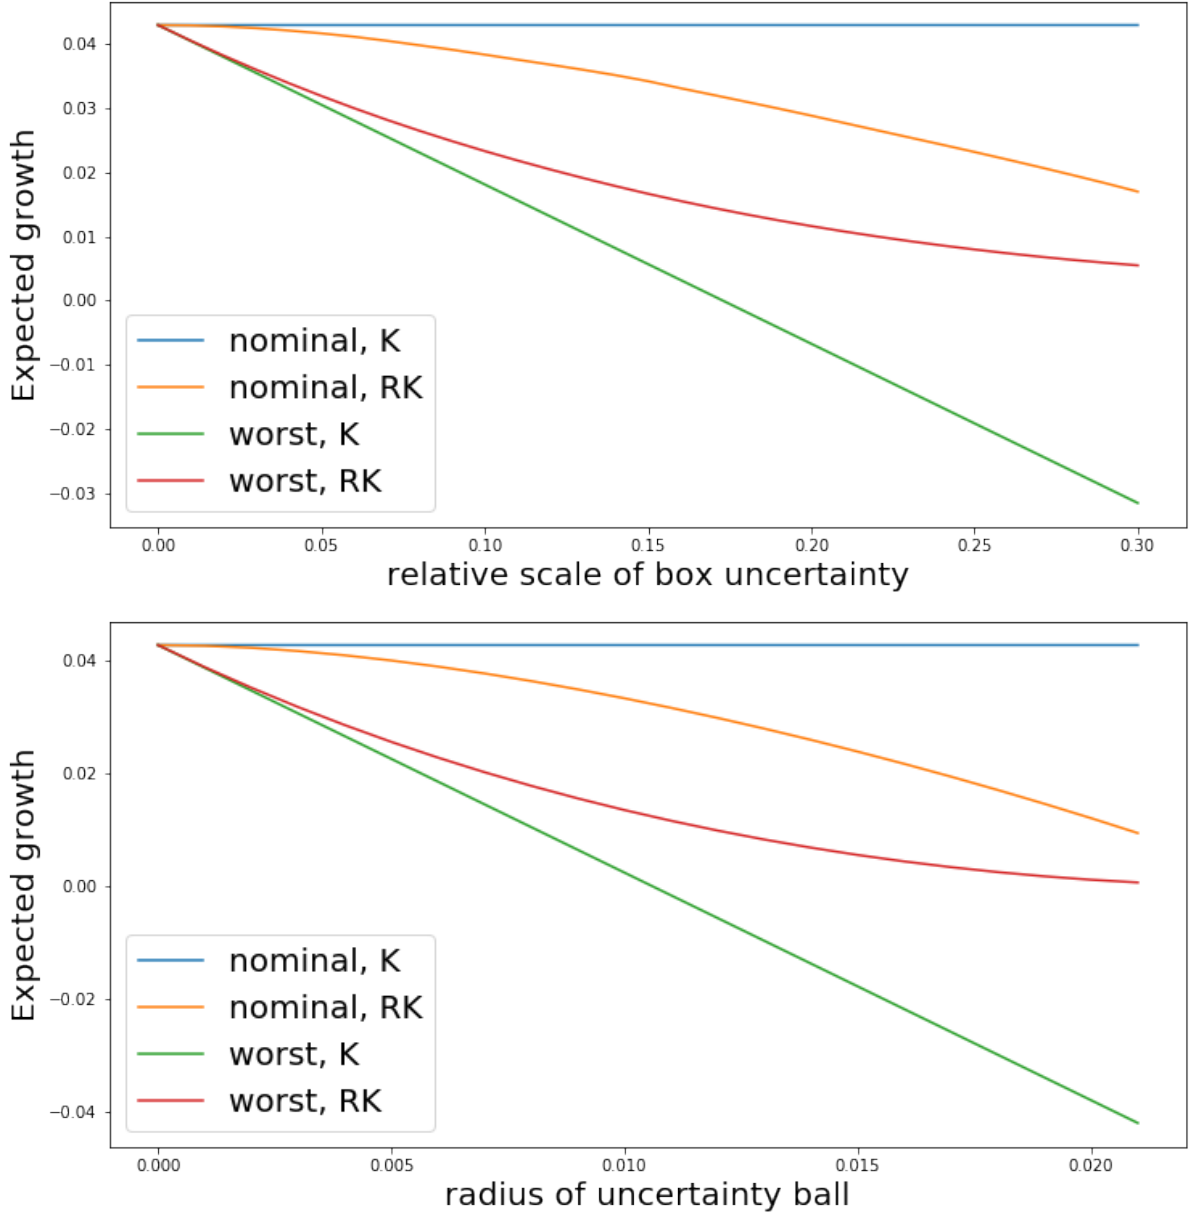


Fig. 4: The box and ball uncertainty set family. The blue, green, orange, red line are $\pi^{\text{nom},T} \log(R^T b^K)$, $\pi_{\eta}^{\text{worst},T} \log(R^T b^K)$, $\pi^{\text{nom},T} \log(R^T b_{\eta}^{RK})$, $\pi_{\eta}^{\text{wc},T} \log(R^T b_{\eta}^{RK})$.

Acknowledgments

We thank David Donoho for suggesting to look at the expected growth for uncertainty sets of varying size. We thank Darrell Duffie for pointing out references related to uncertainty aversion, and we thank Enzo Busseti for useful discussions on Kelly gambling.

Appendix

All of the formulations of DRKP are not only tractable, but easily expressed in DSLs for convex optimization. The CVXPY code to specify and solve the DRKP for ball and box constraints, for example, is given below.

For box uncertainty set, $\Pi_{\rho} = \{\pi \mid |\pi - \pi^{\text{nom}}| \leq \rho, \mathbf{1}^T \pi = 1, \pi \geq 0\}$, the CVXPY code is

```

pi_nom = Parameter(K, nonneg=True)
rho = Parameter(K, nonneg=True)
b = Variable(n)
mu = Variable(K)
wc_growth_rate = min(log(R.T*b) + mu) - pi_nom.T*abs(mu) - rho.T*mu
constraints = [sum(b) == 1, b >= 0]
DRKP = Problem(Maximize(wc_growth_rate), constraints)
DRKP.solve()

```

For ball uncertainty set, $\Pi_c = \{\pi \mid \|\pi - \pi^{\text{nom}}\|_2 \leq c, \mathbf{1}^T \pi = 1, \pi \geq 0\}$, the CVXPY code is

```

pi_nom = Parameter(K, nonneg=True)
c = Parameter((1,1), nonneg=True)
b = Variable(n)
U = Variable(K)
mu = Variable(K)
log_growth = log(R.T*b)
wc_growth_rate = pi_nom.T*F-c*norm(U- mu,2)
constraints = [sum(b) == 1, b >= 0, U <= log_growth]
DRKP = Problem(Maximize(wc_growth_rate), constraints)
DRKP.solve()

```

Here R is the matrix whose columns are the return vectors, $\mathbf{pi_nom}$ is the vector of nominal probabilities. \mathbf{rho} is K dimensional box constraint and c is radius of the ball. For each problem, the second to last line forms the problem, and in the last line the problem is solved. The robust optimal bet is written into $\mathbf{b.value}$. Here

The code for this example can be found at

<https://github.com/QingyunSun/Distributional-Robust-Kelly-Gambling>.

References

1. John Kelly. A new interpretation of information rate. *IRE Transactions on Information Theory*, 2(3):185–189, 1956.
2. Leonard Maclean, Edward Thorp, and William Ziemba. *The Kelly Capital Growth Investment Criterion: Theory and Practice*, volume 3. World Scientific Publishing, 2011.
3. Edward Thorp. Portfolio choice and the Kelly criterion. In *Stochastic Optimization Models in Finance*, pages 599–619. Elsevier, 1975.
4. Mark Davis and Sebastien Lleo. Fractional Kelly strategies in continuous time: Recent developments. In *Handbook of the Fundamentals of Financial Decision Making*, pages 753–788. World Scientific Publishing, 2012.
5. Leonard Maclean, Edward Thorp, and William Ziemba. Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria. *Quantitative Finance*, 10(7):681–687, 2010.
6. Edward Thorp. Understanding the Kelly criterion. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 509–523. World Scientific, 2011.
7. Joseph Kadane. Partial-Kelly strategies and expected utility: Small-edge asymptotics. *Decision Analysis*, 8(1):4–9, 2011.
8. Martin Andersen, Joachim Dahl, and Lieven Vandenbergh. CVXOPT: A Python package for convex optimization, version 1.1.5., 2013.
9. Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
10. Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
11. Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
12. Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
13. Madeleine Udell, Karanveer Mohan, David Zeng, Jenny Hong, Steven Diamond, and Stephen Boyd. Convex optimization in Julia. *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*, 2014.
14. Anqi Fu, Balasubramanian Narasimhan, and Stephen Boyd. CVXR: An R package for disciplined convex optimization. *arXiv preprint arXiv:1711.07582*, 2017.
15. Enzo Busseti, Ernest K Ryu, and Stephen Boyd. Risk-constrained Kelly gambling. *arXiv preprint arXiv:1603.06183*, 2016.
16. Morteza Mardani, Qingyun Sun, Shreyas Vasawanala, Vardan Papyan, Hatef Monajemi, John Pauly, and David Donoho. Neural proximal gradient descent for compressive imaging. *NIPS2018*, 2018.
17. Sid Browne and Ward Whitt. Portfolio choice and the Bayesian Kelly criterion. *Advances in Applied Probability*, 28(4):1145–1176, 1996.
18. Sam Ganzfried and Qingyun Sun. Bayesian opponent exploitation in imperfect-information games. *arXiv preprint arXiv:1603.03491*, 2016.
19. Andrew Lo, Allen Orr, and Ruixun Zhang. The growth of relative wealth and the Kelly criterion. *Journal of Bioeconomics*, 20(1):49–67, 2018.
20. Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
21. Arkadi Nemirovski and Dmitri Yudin. On Cesari’s convergence of the steepest descent method for approximating saddle points of convex-concave functions. *Doklady Akademii Nauk SSSR*, 239:1056–1059, 1978.
22. Arkadi Nemirovski and Dmitri Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley, 1983.
23. Craig Fox and Amos Tversky. Ambiguity aversion and comparative ignorance. *The quarterly journal of economics*, 110(3):585–603, 1995.
24. James Dow and Sergio Ribeiro da Costa Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica: Journal of the Econometric Society*, pages 197–204, 1992.
25. Paolo Ghirardato and Massimo Marinacci. Risk, ambiguity, and the separation of utility and beliefs. *Mathematics of operations research*, 26(4):864–890, 2001.
26. Larry Epstein. A definition of uncertainty aversion. In *Uncertainty in Economic Theory*, pages 187–224. Routledge, 2004.
27. Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010.
28. Kai Yang, Yihong Wu, Jianwei Huang, Xiaodong Wang, and Sergio Verdú. Distributed robust optimization for communication networks. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1157–1165. IEEE, 2008.
29. Mathias Bürger, Giuseppe Notarstefano, and Frank Allgöwer. Distributed robust optimization via cutting-plane consensus. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 7457–7463. IEEE, 2012.
30. Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010.
31. Xin Chen, Melvyn Sim, and Peng Sun. A robust optimization perspective on stochastic programming. *Operations Research*, 55(6):1058–1071, 2007.
32. Almir Mutapcic and Stephen Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods & Software*, 24(3):381–406, 2009.
33. Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015.
34. John Duchi, Peter Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425*, 2016.
35. Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292, 2018.
36. Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

37. Hongseok Namkoong and John Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in Neural Information Processing Systems*, pages 2208–2216, 2016.
38. Jose Blanchet, Lin Chen, and Xunyu Zhou. Distributionally robust mean-variance portfolio selection with Wasserstein distances. *arXiv preprint arXiv:1802.04885*, 2018.
39. Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust Wasserstein profile inference and applications to machine learning. *arXiv preprint arXiv:1610.05627*, 2016.
40. Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, pages 1–52, 2017.
41. Soroosh Shafieezadeh-Abadeh, Peyman Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 1576–1584, 2015.