
Distributionally robust optimization with polynomial densities: theory, models and algorithms

Etienne de Klerk · Daniel Kuhn · Krzysztof Postek

Abstract In distributionally robust optimization the probability distribution of the uncertain problem parameters is itself uncertain, and a fictitious adversary, *e.g.*, nature, chooses the worst distribution from within a known ambiguity set. A common shortcoming of most existing distributionally robust optimization models is that their ambiguity sets contain pathological discrete distributions that give nature too much freedom to inflict damage. We thus introduce a new class of ambiguity sets that contain only distributions with sum-of-squares polynomial density functions of known degrees. We show that these ambiguity sets are highly expressive as they conveniently accommodate distributional information about higher-order moments, conditional probabilities, conditional moments or marginal distributions. Exploiting the theoretical properties of a measure-based hierarchy for polynomial optimization due to Lasserre [*SIAM J. Optim.* 21(3) (2011), pp. 864–885], we prove that certain worst-case expectation constraints are computationally tractable under these new ambiguity sets. We showcase the practical applicability of the proposed approach in the context of a stylized portfolio optimization problem and a risk aggregation problem of an insurance company.

Keywords distributionally robust optimization · semidefinite programming · sum-of-squares polynomials · generalized eigenvalue problem

Mathematics Subject Classification (2000) 90C22 · 90C26 · 90C15

1 Introduction

Since George Dantzig’s 1955 paper on linear programming under uncertainty [10], the field of stochastic programming has developed numerous methods for solving optimization problems that

Etienne de Klerk
Tilburg University & TU Delft, The Netherlands
E-mail: E.deKlerk@uvt.nl

Daniel Kuhn
EPFL, Lausanne, Switzerland
E-mail: daniel.kuhn@epfl.ch

Krzysztof Postek
Erasmus University Rotterdam, The Netherlands
E-mail: postek@ese.eur.nl

depend on uncertain parameters governed by a *known* probability distribution, see, *e.g.*, [5,36,42]. Stochastic programming usually aims to minimize a probability functional such as the expected value, a percentile or the conditional value-at-risk of a given cost function. In practice, however, the distribution that is needed to evaluate this probability functional is at best indirectly observable through independent training samples. Thus, the stochastic programming approach is primarily useful when there is abundant training data. If data is scarce or absent, on the other hand, it may be more adequate to use a robust optimization approach, which models the uncertainty through the set of all possible (or sufficiently likely) uncertainty realizations and minimizes the worst-case costs. Robust optimization is the appropriate modeling paradigm for safety-critical applications with little tolerance for failure and has been popularized in the late 1990's, when it was discovered that robust optimization models often display better tractability properties than stochastic programming models [1]. Distributionally robust optimization is a hybrid approach that attempts to salvage the tractability of robust optimization while maintaining the benefits of (limited) distributional information. In this context, uncertainty is modeled through an *ambiguity set*, that is, a family of typically infinitely many different distributions that are consistent with the available training data or any prior distributional information, and the objective is to minimize the worst-case expected costs across all distributions in the ambiguity set. A distributionally robust newsvendor model that admits an analytical solution has been investigated as early as in 1958 [39], and the theoretical properties of distributionally robust linear programs were first studied in 1966 [46]. Interest in distributionally robust optimization has also been fuelled by important applications in finance [34,33]. However, only recently it was recognized that many distributionally robust optimization problems of practical relevance can actually be solved in polynomial time. Tractability results are available both for *moment ambiguity sets*, which contain all distributions that satisfy a finite number of moment conditions [11,15,45], as well as for *metric-based ambiguity sets*, which contain all distributions within a prescribed distance from a nominal distribution with respect to some probability metric [2,31]. In all these cases, the extremal distributions that determine the worst-case expectation are discrete, and the number of their discretization points is often surprisingly small, *e.g.*, proportional to the number of moment constraints. As these unnatural discrete distributions are almost always inconsistent with the available training samples, distributionally robust optimization models with moment and metric-based ambiguity sets are often perceived as overly pessimistic.

In an attempt to mitigate the over-conservatism of traditional distributionally robust optimization, several authors have studied moment ambiguity sets that require their member distributions to satisfy additional structural properties such as symmetry, unimodality, monotonicity or smoothness etc. By leveraging ideas from Choquet theory and polynomial optimization, it has been shown that the resulting distributionally robust optimization problems admit hierarchies of increasingly accurate semidefinite programming (SDP) bounds [35]. An exact SDP reformulation for the worst-case probability of a polytope with respect to all unimodal distributions with known first and second moments is derived in [44], while second-order conic reformulations of distributionally robust individual chance constraints with moment and unimodality information are reported in [27]. For a survey of recent results on distributionally robust uncertainty quantification and chance constrained programming problems with moment and structural information we refer to [18]. Even though unimodality or monotonicity conditions eliminate all discrete distributions from a moment ambiguity set, the extremal distributions that critically determine all worst-case expectations remain pathological. For example, all extremal *unimodal* distributions are supported on line segments emanating from a single point in space (the mode) and thus fail to be absolutely continuous with respect to the Lebesgue measure. Thus, the existing distributionally robust optimization models with structural

information remain overly conservative. This observation motivates us to investigate a new class of ambiguity sets that contain only distributions with non-degenerate polynomial density functions.

This paper aims to study *worst-case expectation constraints* of the form

$$\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} f(x, z) \leq 0, \quad (1)$$

where $x \in \mathbb{R}^n$ is a decision vector, $z \in \mathbb{R}^m$ is an uncertain parameter governed by an ambiguous probability distribution $\mathbb{P} \in \mathcal{P}$, and $f(x, z)$ is an uncertainty-affected constraint function that can be interpreted as a cost. In words, the constraint (1) requires that the expected cost of the decision x be non-positive for every distribution in the ambiguity set \mathcal{P} . Throughout the paper we will assume that $f(x, z)$ depends polynomially on z and that each distribution $\mathbb{P} \in \mathcal{P}$ admits a sum-of-squares (hence non-negative) polynomial density function $h(z)$ with respect to some prescribed reference measure μ on \mathbb{R}^m (e.g., the Lebesgue measure). Imposing an upper bound on the polynomial degree of $h(z)$ thus yields a finite-dimensional parameterization of the ambiguity set \mathcal{P} . Moreover, many popular distributional properties can be expressed through linear constraints on the coefficients of $h(z)$ and are thus conveniently accounted for in the definition of \mathcal{P} . Examples include moment bounds, probability bounds for certain subsets of \mathbb{R}^m , bounds on conditional tail probabilities and marginal distribution conditions. Note that by fixing the marginal distributions of all components of z , the worst-case expectation problem on the left-hand side of (1) reduces to a Fréchet problem that seeks the worst-case copula of the uncertain parameters.

By leveraging a measure-based hierarchy for polynomial optimization due to Lasserre [24], we will demonstrate that the subordinate worst-case probability problem in (1) admits an exact SDP reformulation. Under mild additional conditions on $f(x, z)$, we will further prove that the feasible set of the constraint (1) admits a polynomial-time separation oracle. Moreover, we will analyze the convergence of the worst-case expectation in (1) as the polynomial degree of $h(z)$ tends to infinity, and we will illustrate the practical use of the proposed approach through numerical examples.

More succinctly, the main contributions of this paper can be summarized as follows:

- (i) *Modeling power*: We introduce a new class of ambiguity sets containing distributions that admit sum-of-squares polynomial density functions of degree at most $2r$, $r \in \mathbb{N}$, with respect to a given reference measure. Ambiguity sets of this type are highly expressive as they conveniently accommodate distributional information about higher-order moments, conditional probabilities or conditional moments. They also allow the modeler to prescribe (not necessarily discrete) marginal distributions that must be matched exactly by all distributions in the ambiguity set.
- (ii) *Computational tractability*: We identify general conditions under which the worst-case expectations over the new ambiguity sets can be reformulated *exactly* as tractable SDPs with $\mathcal{O}\binom{n+r}{r}$ variables. We also propose an efficient heuristic for computing the worst-case expectations *approximately* by solving a sequence of significantly smaller SDPs. Finally, we delineate conditions under which the feasible sets of the worst-case expectation constraints admit a polynomial-time separation oracle and thus lend themselves to efficient optimization via the ellipsoid method.
- (iii) *Convergence analysis*: We demonstrate that, as r tends to infinity, the worst-case expectations over the new ambiguity sets converge monotonically to classical worst-case expectations over larger ambiguity sets that relax the polynomial density requirement. At the same time, the extremal density functions converge to pathological discrete worst-case distributions characteristic for classical moment ambiguity sets without restrictions on the density functions.
- (iv) *Numerical results*: We showcase the practical applicability of the proposed approach in the context of a stylized portfolio optimization problem and a simple Fréchet problem inspired by [43] that models the risk aggregation problem of an insurance company.

The intimate relation between polynomial optimization and the problem of moments has already been exploited in several papers on distributionally robust optimization. For example, ideas from polynomial optimization give rise to SDP bounds on the probability of a semi-algebraic set [4] or the expected value of a piecewise polynomial [47] across all probability distributions satisfying a given set of moment constraints. These SDP bounds are tight in the univariate case or if only marginal moments are specified. Otherwise, one may obtain hierarchies of asymptotically tight SDP bounds. As an application, these techniques can be used to derive bounds on the prices of options with piecewise polynomial payoff functions, based solely on the knowledge of a few moments of the underlying asset prices [3]. Moreover, asymptotically tight SDP bounds that account for both moment and structural information are proposed in [35]. However, all these approaches differ from our work in that the ambiguity sets have discrete or otherwise degenerate extremal distributions.

Distributionally robust polynomial optimization problems over non-degenerate polynomial density functions that are close to a nominal density estimate (obtained, *e.g.*, via a Legendre series density estimator) in terms of the L_2 -distance are considered in [30]. In this work the non-negativity of the candidate density functions is not enforced explicitly, which considerably simplifies the problem and may be justified if the distance to the nominal density is sufficiently small. It is shown that the emerging distributionally robust optimization problems are equivalent to deterministic polynomial optimization problems that are not significantly harder than the underlying nominal problem and can be addressed by solving a sequence of tractable SDP relaxations.

Distributionally robust chance constraints with ambiguity sets containing all possible mixtures of a given parametric distribution family are studied in [24]. The mixtures are encoded through a probability density function on a compact parameter space. The authors propose an asymptotically tight SDP hierarchy of inner approximations for the feasible set of the distributionally robust chance constraint. In contrast, we explicitly represent all probability distributions in the ambiguity set through polynomial density functions that can capture a wide range of distributional features.

The remainder of this paper develops as follows. Section 2 reviews Lasserre’s measure-based approach to polynomial optimization, which is central to this paper. A major drawback of the resulting SDP hierarchies is their limited scalability. This prompts us to devise an efficient heuristic solution algorithm in Section 3. Sections 4 and 5 develop SDP hierarchies for worst-case expectation constraints of the form (1) with and without moment information, respectively, and investigate the convergence of the underlying worst-case expectations as the degree of the polynomial density functions tends to infinity. Section 6 highlights the modeling power of the proposed approach, while Section 7 reports on numerical results for a portfolio design problem as well as a risk aggregation problem of an insurance company. Conclusions are drawn in Section 8.

2 Lasserre’s measure-based hierarchy for polynomial optimization

In what follows, we denote by $x^\alpha := \prod_{i=1}^n x_i^{\alpha_i}$ the monomial of the variables $x = (x_1, \dots, x_n)$ with respective exponents $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{N}_0^n$, and we define $N(n, r) := \{\alpha \in \mathbb{N}_0^n : \sum_{i=1}^n \alpha_i \leq r\}$ as the set of all exponents that give rise to monomials with degree at most r . We let $\Sigma[x]$ denote the set of all sum-of-squares (SOS) polynomials in the variables x , and we define $\Sigma[x]_r$ as the subset of all SOS polynomials with degree at most $2r$.

Now consider the polynomial global optimization problem

$$p_{\min, \mathbf{K}} := \min_{x \in \mathbf{K}} p(x) = \min_{x \in \mathbf{K}} \sum_{\alpha \in N(n, d)} p_\alpha x^\alpha, \quad (2)$$

where $p(x) = \sum_{\alpha \in N(n,d)} p_\alpha x^\alpha$ is an n -variate polynomial of degree d , and $\mathbf{K} \subset \mathbb{R}^n$ a closed set with nonempty interior. (We assume existence of a global minimizer.)

We also assume that the moments of a finite Borel measure μ supported on \mathbf{K} are known in the sense that they are either available in closed-form or efficiently computable. To be clear, we view a finite Borel measure μ on \mathbb{R}^n as a nonnegative set function defined on the Borel σ -algebra of \mathbb{R}^n . (Recall that the Borel σ -algebra is generated by all open sets in \mathbb{R}^n .) By definition, μ must satisfy $\mu(\emptyset) = 0$ and $\mu(\cup_{i=1}^\infty S_i) = \sum_{i=1}^\infty \mu(S_i)$ for any collection of disjoint, measurable sets $S_i \subset \mathbb{R}^n$, $i \in \mathbb{N}$, and $\mu(\mathbb{R}^n) < \infty$. The support of μ , denoted by $\text{supp}(\mu)$, is defined as the smallest closed set \mathbf{K} so that $\mu(\mathbb{R}^n \setminus \mathbf{K}) = 0$.

We denote the (known) moments of μ by

$$m_\alpha(\mathbf{K}) := \int_{\mathbf{K}} x^\alpha d\mu(x) \quad \text{for } \alpha \in \mathbb{N}_0^n. \quad (3)$$

Lasserre [24] introduced the following upper bound on $p_{\min, \mathbf{K}}$,

$$\begin{aligned} \underline{p}_{\mathbf{K}}^{(r)} &:= \min_{h \in \Sigma_r} \left\{ \int_{\mathbf{K}} p(x)h(x)d\mu(x) : \int_{\mathbf{K}} h(x)d\mu(x) = 1 \right\} \\ &= \min_{h \in \Sigma[x]_r} \mathbb{E}_{x \sim (\mathbf{K}, h)}[p(x)], \end{aligned} \quad (4)$$

where r is a fixed integer, and $x \sim (\mathbf{K}, h)$ indicates that x is a random vector supported on \mathbf{K} that is governed by the probability measure $h \cdot d\mu$. It is known that if μ is the Lebesgue measure, then $\underline{p}_{\mathbf{K}}^{(r)}$ is equal to the the smallest generalized eigenvalue of the system

$$Av = \lambda Bv, \quad (5)$$

with $v \neq 0$, where the symmetric matrices A and B are of size $\binom{n+r}{r}$ with rows and columns indexed by $N(n, r)$, and

$$A_{\alpha, \beta} = \sum_{\delta \in N(n,d)} p_\delta m_{\alpha+\beta+\delta}(\mathbf{K}), \quad B_{\alpha, \beta} = m_{\alpha+\beta}(\mathbf{K}) \quad \text{for } \alpha, \beta \in N(n, r). \quad (6)$$

Lasserre [24] establishes conditions on μ and \mathbf{K} so that $\lim_{r \rightarrow \infty} \underline{p}_{\mathbf{K}}^{(r)} = p_{\min, \mathbf{K}}$, and the rate of convergence was subsequently studied in [9, 7, 8] for special choices of μ and \mathbf{K} . The most general condition under which convergence holds, as shown in [25, Theorem 2.2], is when \mathbf{K} is closed with nonempty interior, and the moments of μ on \mathbf{K} satisfy the following conditions:

$$\int_{\mathbf{K}} x_i^{2k} d\mu(x) \leq (2k)!M \quad \forall i \in \{1, \dots, n\}, k \in \mathbb{N}, \quad (7)$$

for some $M > 0$. For example, if one defines μ in terms of a finite Borel measure φ with $\text{supp}(\varphi) = \mathbf{K}$ via

$$d\mu(x) = \exp(-|x_1| - \dots - |x_n|) d\varphi(x), \quad (8)$$

then this choice satisfies the conditions (7); see [24, §3.2].

We summarize the known convergence results in Table 1.

Table 1: Known rates of convergence for the Lasserre hierarchy

$\mathbf{K} \subset \mathbb{R}^n$	$p_{\mathbf{K}}^{(r)} - p_{\min, \mathbf{K}}$	measure μ , $\text{supp}(\mu) = \mathbf{K}$	reference
closed, nonempty interior	$o(1)$	satisfies (7)	[24]
compact, nonempty interior	$o(1)$	finite Borel measure	[24]
compact, satisfies interior cone condition	$O\left(\frac{1}{\sqrt{r}}\right)$	Lebesgue measure	[9]
convex body	$O\left(\frac{1}{r}\right)$	Lebesgue measure	[7]
$[-1, 1]^n$	$\Theta\left(\frac{1}{r^2}\right)$	$d\mu(x) = \prod_{i=1}^n (1 - x_i^2)^{-1/2} dx$	[8]

2.1 Examples of known moments

The moments (3) are available in closed-form, for example, if μ is the Lebesgue measure and \mathbf{K} is an ellipsoid or triangulated polytope; see, *e.g.*, [24] and [9]. For the canonical simplex, $\Delta_n = \{x \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq 1\}$, we have

$$m_\alpha(\Delta_n) = \frac{\prod_{i=1}^n \alpha_i!}{(\sum_{i=1}^n \alpha_i + n)!}, \quad (9)$$

see, *e.g.*, [22, Equation (2.4)] or [17, Equation (2.2)]. One may trivially verify that the moments for the hypercube $\mathbf{Q}_n = [0, 1]^n$ are given by

$$m_\alpha(\mathbf{Q}_n) = \int_{\mathbf{Q}_n} x^\alpha dx = \prod_{i=1}^n \int_0^1 x_i^{\alpha_i} dx_i = \prod_{i=1}^n \frac{1}{\alpha_i + 1}.$$

The moments for the unit Euclidean ball are given by

$$m_\alpha(B_1(0)) = \begin{cases} \frac{\pi^{(n-1)/2} 2^{(n+1)/2} \prod_{i=1}^n (\alpha_i - 1)!!}{(n + \sum_{i=1}^n \alpha_i)!!} & \text{if } \alpha_i \text{ is even for all } i, \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

where the double factorial of any integer k is defined through

$$k!! = \begin{cases} k \cdot (k-2) \cdots 3 \cdot 1 & \text{if } k > 0 \text{ is odd,} \\ k \cdot (k-2) \cdots 4 \cdot 2 & \text{if } k > 0 \text{ is even,} \\ 1 & \text{if } k = 0 \text{ or } k = -1. \end{cases}$$

When \mathbf{K} is an ellipsoid, one may obtain the moments from (10) by applying an affine transformation of variables. Another tractable support set that will become relevant in Section 7.1 of this paper is the knapsack polytope, that is, the intersection of a hypercube and a half-space; the moments for this and related polytopes are derived in [29]. Finally, in Section 7.2 we will work with the nonnegative orthant $\mathbf{K} = \mathbb{R}_+^n$. Since \mathbf{K} is unbounded in this case, we need to introduce a measure of the form (8). A suitable choice that corresponds to $d\varphi(x) = 2^n \exp(-\sum_{i=1}^n x_i) dx$ in (8), is

$$d\mu(x) = \exp\left(-\sum_{i=1}^n x_i\right) dx.$$

This is the exponential measure associated with the orthogonal Laguerre polynomials. For more information, the reader is referred to [24, §3.2]. We will also use another choice of measure for

$\mathbf{K} = \mathbb{R}_+^n$ in Section 7.2, namely the lognormal measure,

$$d\mu(x) = \prod_{i=1}^n \frac{1}{x_i v_i \sqrt{2\pi}} \exp\left(-\frac{(\ln(x_i) - \bar{z}_i)^2}{2v_i^2}\right) dx_i, \quad (11)$$

where \bar{z}_i and v_i represent prescribed location and scale parameters, $i = 1, \dots, n$. The moments of μ are given by

$$m_\alpha(\mathbf{K}) = \prod_{i=1}^n \exp(\alpha_i \bar{z}_i + (\alpha_i v_i)^2 / 2). \quad (12)$$

One may readily verify that these moments do *not* satisfy the bounds on the moments in (7). When using this measure we are therefore not guaranteed convergence of the Lasserre hierarchy.

We stress that, even though these examples of known moments are limited, they include typical sets that are routinely used in (distributionally) robust optimization to represent uncertainty sets or supports, most notably budget uncertainty sets and ellipsoids.

3 An efficient, heuristic implementation of a Lasserre-type hierarchy

The drawback of solving problem (4) is that it involves operations with matrices of order $\binom{n+r}{r}$ for increasing values of r . Thus one is limited to relatively small values of n and r .

In this section we describe a weaker hierarchy of bounds that is similar in spirit to the hierarchy in (4), but where the sizes of the corresponding generalized eigenvalue problems remain the same at each level of the hierarchy. Conceptually, the idea is to use the optimal density function, say $h \in \Sigma_r$ at level r , to approximate the optimal density function at a higher level in the hierarchy.

To explain the idea, consider again the global optimization problem (2), which minimizes an n -variate polynomial p over a convex body \mathbf{K} . We assume that the moments of a prescribed reference measure μ supported on \mathbf{K} are known, and we denote these moments by

$$m_\alpha^\mu(\mathbf{K}) := \int_{\mathbf{K}} x^\alpha d\mu(x) \quad \text{for } \alpha \in \mathbb{N}_0^n,$$

where we add the superscript μ to make the dependence on the reference measure explicit. Next we compute the upper bound $\underline{p}_{\mathbf{K}}^{(r)}$ as in (4), where r is a fixed (small) integer. Denoting the resulting optimal density by $h(x) = \sum_{\beta \in N(n, 2r)} h_\beta x^\beta \in \Sigma[x]_r$, we can then define a new probability measure μ' on \mathbf{K} through

$$d\mu'(x) = h(x) \cdot d\mu(x). \quad (13)$$

Note that we may obtain the moments of μ' from the moments of μ via

$$\begin{aligned} m_\alpha^{\mu'}(\mathbf{K}) &= \int_{\mathbf{K}} x^\alpha d\mu'(x) \\ &= \int_{\mathbf{K}} x^\alpha h(x) d\mu(x) \\ &= \sum_{\beta \in N(n, 2r)} h_\beta \int_{\mathbf{K}} x^{\alpha+\beta} d\mu(x) \\ &= \sum_{\beta \in N(n, 2r)} h_\beta m_{\alpha+\beta}^\mu(\mathbf{K}) \quad \text{for } \alpha \in \mathbb{N}_0^n. \end{aligned}$$

Finally, one may now replace μ by μ' and repeat the same process R times for some fixed $R \in \mathbb{N}$. The complete procedure is summarized in Algorithm 1.

Data: Polynomial p of degree d , allowed degree $r \geq d$, integer order R ; moments of some measure μ up to order $2r + d + 2(R - 1)r = 2rR + d$, *i.e.*, the values $m_\alpha^\mu(\mathbf{K})$ for all $\alpha \in N(n, 2rR + d)$

Result: Upper bound of order R on $p_{\min, \mathbf{K}}$

for $k \leftarrow 1$ **to** R **do**

 Form the matrices A and B defined in (6) ;

 Solve the generalized eigenvalue problem for A and B in (5) to obtain the optimal density $h \in \Sigma_r$;

 Define the measure μ' via (13) ;

 Obtain the moments of μ' via $m_\alpha^{\mu'}(\mathbf{K}) = \sum_{\beta \in N(n, 2r)} h_\beta m_{\alpha+\beta}^\mu(\mathbf{K})$ for all $\alpha \in N(n, 2r(R - k) + d)$;

 Replace $\mu \leftarrow \mu'$ and $m_\alpha^\mu(\mathbf{K}) \leftarrow m_\alpha^{\mu'}(\mathbf{K})$ for all $\alpha \in N(n, 2r(R - k) + d)$;

end

Algorithm 1: Algorithm to compute the upper bound $\underline{p}_{\mathbf{K}}^{(r, R)}$ of order R on $\underline{p}_{\mathbf{K}}^{(r)}$.

The following remarks on this heuristic procedure are in order:

1. The bound computed by the algorithm is no better than the bound $\underline{p}_{\mathbf{K}}^{(r, R)}$, but is much cheaper to compute because, in each iteration, it only involves generalized eigenvalue problems of order $\binom{n+r}{r}$ for a small fixed integer r , *e.g.*, $r = 4$.
2. The bounds generated by the algorithm, as indexed by R , are not guaranteed to converge to $p_{\min, \mathbf{K}}$ as $R \rightarrow \infty$. However, one may easily obtain a convergent variant (at a computational cost) by increasing the value r inside a given iteration, if no improvement in the upper bound is obtained in that iteration.
3. One has to store and update a moment table indexed by $\alpha \in N(n, 2r(R - k) + d)$ in iteration k , and the updating process involves simple linear algebra.

We now give some numerical examples to indicate how this algorithm performs. We will consider the test functions in Table 2.

Table 2: Test functions, all with $n = 2$, domain $\mathbf{K} = [-1, 1]^2$, and minimum $p_{\min, \mathbf{K}} = 0$.

Name	$p(x)$
Matyas function	$26(x_1^2 + x_2^2) - 48x_1x_2$
Motzkin polynomial	$64(x_1^4x_2^2 + x_1^2x_2^4) - 48x_1^2x_2^2 + 1$

For ease of reference, we will denote the bound generated by Algorithm 1 by $\underline{p}_{\mathbf{K}}^{(r, R)}$. As this bound corresponds to a density function of degree rR with respect to the initial reference measure, it is natural to compare it to the stronger, but more expensive, bound $\underline{p}_{\mathbf{K}}^{(r \cdot R)}$. The following table lists different bounds for the case $r \cdot R = 20$, each corresponding to a density function of degree 40.

Note that the $\underline{p}_{\mathbf{K}}^{(r, R)}$ bounds with largest r are the strongest in the examples, as one may expect.

Function	$p_{\mathbf{K}}^{(20)}$	$p_{\mathbf{K}}^{(10,2)}$	$p_{\mathbf{K}}^{(5,4)}$	$p_{\mathbf{K}}^{(4,5)}$	$p_{\mathbf{K}}^{(2,10)}$
Matyas	0.4811	0.4989	0.7285	0.9604	1.1070
Motzkin	0.1817	0.1969	0.2907	0.3603	0.4588

4 Distributionally robust constraints involving polynomial uncertainty

We now consider a worst-case feasibility constraint of the form (1), where $z \in \mathbb{R}^m$ represents a random vector with a support $\mathbf{K} \subset \mathbb{R}^m$, assumed to be closed and with nonempty interior. Assume that the constraint function $f(x, z)$ displays a polynomial dependence on z . In particular, assume that $f(x, z) = \sum_{\beta \in N(m, d)} f_{\beta}(x) z^{\beta}$ has degree d in z , where the $f_{\beta} : \mathbb{R}^n \rightarrow \mathbb{R}$ are functions of x only.

If the ambiguity set \mathcal{P} contains all distributions that have an SOS polynomial density of degree at most $2r$, $r > 1$, with respect to a fixed, finite Borel measure μ supported on \mathbf{K} , then the worst-case expected feasibility constraint (1) reduces to

$$f_{\mathbf{K}}^{(r)}(x) := \sup_{h \in \Sigma[z]_r} \left\{ \int_{\mathbf{K}} f(x, z) h(z) d\mu(z) : \int_{\mathbf{K}} h(z) d\mu(z) = 1 \right\} \leq 0. \quad (14)$$

Formally speaking, we consider an ambiguity set of the form

$$\mathcal{P} = \left\{ h \cdot d\mu : h \in \Sigma[z]_r, \int_{\mathbf{K}} h(z) d\mu(z) = 1 \right\}. \quad (15)$$

We assume that the moments of the measure μ on \mathbf{K} are available, and we again use the notation

$$m_{\alpha}(\mathbf{K}) := \int_{\mathbf{K}} z^{\alpha} d\mu(z) \quad \text{for } \alpha \in \mathbb{N}_0^m.$$

Expressing $h \in \Sigma[z]_r$ as $h(z) = \sum_{\alpha \in N(m, 2r)} h_{\alpha} z^{\alpha}$, the left-hand-side of (14) may be re-written as

$$\begin{aligned} & \sup_{h_{\alpha} : \alpha \in N(m, 2r)} \sum_{\beta \in N(m, d)} f_{\beta}(x) \sum_{\alpha \in N(m, 2r)} h_{\alpha} m_{\alpha+\beta}(\mathbf{K}) \\ & \text{s.t.} \quad \sum_{\alpha \in N(m, 2r)} h_{\alpha} m_{\alpha}(\mathbf{K}) = 1, \\ & \quad \sum_{\alpha \in N(m, 2r)} h_{\alpha} z^{\alpha} \in \Sigma[z]_r. \end{aligned} \quad (16)$$

Since the sum-of-squares condition on h is equivalent to a linear matrix inequality in its coefficients h_{α} , problem (16) constitutes a tractable semidefinite program (SDP) in h_{α} , $\alpha \in N(m, 2r)$, if x is fixed. The next theorem establishes that we can also efficiently optimize over the feasible set of the constraint (14) whenever the coefficient functions f_{β} are convex and $\mathbf{K} \subset \mathbb{R}_+^m$.

Theorem 1 *Consider the constraint (14) and assume that all f_{β} are convex functions of x whose subgradients are efficiently computable. Moreover, assume that $\mathbf{K} \subset \mathbb{R}_+^m$. Then, the set of $x \in \mathbb{R}^n$ that satisfy (14) is convex and admits a polynomial-time separating hyperplane oracle.*

Proof We have to show that the function $f_{\mathbf{K}}^{(r)}(x)$ from (14) is convex in x . We may rewrite the function as

$$f_{\mathbf{K}}^{(r)}(x) = \sup_{h \in \Sigma[z]_r} \sum_{\beta \in N(m,d)} \mathbb{E}_{z \sim (\mathbf{K}, h)} [z^\beta] f_\beta(x)$$

For each $h \in \Sigma[z]_r$, the function $\mathbb{E}_{z \sim (\mathbf{K}, h)} [z^\beta] f_\beta(x)$ is convex in x , since $\mathbf{K} \subset \mathbb{R}_+^n$ implies $\mathbb{E}_{z \sim (\mathbf{K}, h)} [z^\beta] \geq 0$. Thus $f_{\mathbf{K}}^{(r)}(x)$ is the point-wise supremum of an infinite collection of convex functions, and therefore convex itself (see, *e.g.*, [37, Theorem 5.5]). Thus the set $\mathcal{C} := \{x \in \mathbb{R}^n \mid f_{\mathbf{K}}^{(r)}(x) \leq 0\}$ is convex.

If $\bar{x} \notin \mathcal{C}$, *i.e.*, $f_{\mathbf{K}}^{(r)}(\bar{x}) > 0$, then we may construct a hyperplane that separates \bar{x} from \mathcal{C} as follows. Let $\bar{h} \in \Sigma[z]_r$ be such that

$$f_{\bar{h}}(\bar{x}) := \sum_{\beta \in N(m,d)} \mathbb{E}_{z \sim (\mathbf{K}, \bar{h})} [z^\beta] f_\beta(\bar{x}) > 0.$$

One may obtain such an \bar{h} in polynomial time by solving the SDP (16) with fixed $x = \bar{x}$. Now let $\partial f_{\bar{h}}(\bar{x})$ denote a subgradient of $f_{\bar{h}}$ at \bar{x} . (By assumption such a subgradient is available in polynomial time.) By the definition of a subgradient, we now have

$$\partial f_{\bar{h}}(\bar{x})^T (x - \bar{x}) \leq f_{\bar{h}}(x) - f_{\bar{h}}(\bar{x}) \leq -f_{\bar{h}}(\bar{x}) \quad \forall x \in \mathcal{C}.$$

The outer linear inequality now separates \bar{x} from \mathcal{C} . \square

Theorem 1 implies that if all coefficient functions f_β are convex, one may optimize a convex function of x over a set given by constraints of the type (14) in polynomial time, *e.g.*, by using the ellipsoid method, provided that an initial ellipsoid is known that contains an optimal solution [16].

Finally, we point out that, due to the convergence properties of the Lasserre hierarchy, one recovers the usual robust counterpart (robust against the single worst-case realization of z as in [1]) in the limit as r tends to infinity.

Theorem 2 *Assume that $\mathbf{K} \subset \mathbb{R}^n$ is closed with nonempty interior. Then, in the limit as $r \rightarrow \infty$, the constraint (14) reduces to the usual robust counterpart constraint*

$$\max_{z \in \mathbf{K}} f(x, z) \leq 0.$$

More precisely, if $x \in \mathbf{K}$ is fixed, and (\mathbf{K}, μ) satisfies one of the assumptions in Table 1, one has

$$\lim_{r \rightarrow \infty} f_{\mathbf{K}}^{(r)}(x) = \max_{z \in \mathbf{K}} f(x, z).$$

Moreover the rate of convergence is as given in Table 1, depending on the choice of (\mathbf{K}, μ) .

Proof For fixed x , the computation of $f_{\mathbf{K}}^{(r)}(x)$ is an SDP problem of the form (4), and the required convergence result therefore follows from the convergence of the Lasserre hierarchy (4), as summarized in Table 1. \square

5 Approximate solution of the general problem of moments

In applications it is often possible to inject moment information into the ambiguity set \mathcal{P} . For example, if there is prior information about the location or the dispersion of the random vector z , one can include constraints on its mean vector or its covariance matrix into the definition of the ambiguity set. Specifically, if it is known that $\mathbb{E}_{z \sim (\mathbf{K}, h)}[z^{\beta_i}] = \gamma_i$ for some $\beta_i \in \mathbb{N}_0^n$ and $\gamma_i \in \mathbb{R}$ for $i = 1, \dots, p$, one can restrict the ambiguity set (15) by including the moment constraints

$$\int_{\mathbf{K}} z^{\beta_i} h(z) d\mu(z) = \sum_{\alpha \in N(m, 2r)} h_{\alpha} m_{\alpha + \beta_i}(\mathbf{K}) = \gamma_i \quad \forall i = 1, \dots, p,$$

which reduce to simple linear equations for the coefficients h_{α} , $\alpha \in N(m, 2r)$, of the density function h . In this setup, the maximization over the ambiguity set corresponds to a general problem of moments, see, e.g., [41]. We show how our approach may be used to solve this problem, and we will illustrate this through concrete examples in Section 7.

Throughout this section we assume that $\mathbf{K} \subset \mathbb{R}^k$ is a nonempty closed set, while f_0, f_1, \dots, f_p are real-valued Borel-measurable functions on \mathbf{K} . Moreover, we assume that μ is a finite Borel measure on \mathbf{K} such that f_0, \dots, f_p are μ -integrable.

Theorem 3 *Let \mathbf{K}_i , $i = 0, \dots, p$, be Borel-measurable subsets of \mathbf{K} . Then there exists an atomic Borel measure μ' on \mathbf{K} with a finite support of at most $p + 2$ points so that*

$$\int_{\mathbf{K}_i} f_i(z) d\mu(z) = \int_{\mathbf{K}_i} f_i(z) d\mu'(z) \quad \forall i = 0, \dots, p.$$

Proof This result is due to Rogosinsky [38], but an elementary proof is given by Shapiro [41, Lemma 3.1]; see also Lasserre [23]. \square

As a consequence one has the following result for the problem of moments.

Corollary 1 *Consider the problem of moments*

$$val := \inf_{\mu \in \mathcal{P}} \left\{ \int_{\mathbf{K}_0} f_0(z) d\mu(z) : \int_{\mathbf{K}_i} f_i(z) d\mu(z) = b_i \quad \forall i = 1, \dots, p \right\}, \quad (17)$$

where \mathcal{P} is the set of all Borel probability measures supported on \mathbf{K} , and \mathbf{K}_i is a Borel-measurable subset of \mathbf{K} for each $i = 0, \dots, p$. If the problem has a solution, it has a solution that is an atomic measure supported on at most $p + 2$ points in \mathbf{K} , i.e., a convex combination of at most $p + 2$ Dirac delta measures supported in \mathbf{K} .

In what follows we show how the atomic measure solution, whose existence is guaranteed by Corollary 1, may be approximated arbitrarily well by SOS polynomial density functions.

Theorem 4 *Consider problem (17) with the additional assumptions that $\mathbf{K} \subset \mathbb{R}^n$ has nonempty interior and that the functions f_0, f_1, \dots, f_p are polynomials. Also assume that \mathbf{K} and μ satisfy one of the assumptions in Table 1. Then, for any $\epsilon > 0$ there exists a $d \in \mathbb{N}$ and a probability density $h \in \Sigma_d[z]$ such that one has*

$$\begin{aligned} \int_{\mathbf{K}_0} f_0(z) h(z) d\mu(z) &\in (val - \epsilon, val + \epsilon) \\ \int_{\mathbf{K}_i} f_i(z) h(z) d\mu(z) &= (b_i - \epsilon, b_i + \epsilon) \quad \forall i = 1, \dots, p. \end{aligned}$$

Moreover, for the choices of \mathbf{K} and μ in Table 1 where a rate of convergence is known, one may bound d in terms of ϵ . For example, if \mathbf{K} is a convex body and μ the Lebesgue measure, then one may assume that $d = O(1/\epsilon^2)$.

Proof Fix $a \in \mathbf{K}$, and consider the polynomials $z \mapsto (f_i(z) - f_i(a))^2$, $i = 0, \dots, p$. Moreover, let p be a polynomial with global minimizer a such that $p(a) = 0$ and p upper bounds all these polynomials on \mathbf{K} , i.e.,

$$p(z) \geq (f_i(z) - f_i(a))^2 \text{ for all } z \in \mathbf{K} \text{ and all } i \in \{0, \dots, p\}. \quad (18)$$

For a given probability density $h \in \Sigma[z]_r$ with $\int_{\mathbf{K}} h(z) d\mu(z) = 1$, we denote as before

$$\mathbb{E}_{z \sim (\mathbf{K}, h)}[p(z)] = \int_{\mathbf{K}} p(z) h(z) d\mu(z).$$

Note that by (18) we have $\mathbb{E}_{z \sim (\mathbf{K}, h)}[p(z)] \geq \mathbb{E}_{z \sim (\mathbf{K}, h)}[(f_i(z) - f_i(a))^2]$. Combining with Jensen's inequality, we therefore conclude

$$\begin{aligned} (\mathbb{E}_{z \sim (\mathbf{K}, h)}[(f_i(z) - f_i(a))])^2 &\leq \mathbb{E}_{z \sim (\mathbf{K}, h)}[(f_i(z) - f_i(a))^2] \\ &\leq \mathbb{E}_{z \sim (\mathbf{K}, h)}[p(z)]. \end{aligned}$$

Recalling the notation of the Lasserre hierarchy from (4), we denote $\underline{p}_{\mathbf{K}}^{(r)} = \min_{h \in \Sigma[z]_r} \mathbb{E}_{z \sim (\mathbf{K}, h)}[p(z)]$. If μ and \mathbf{K} satisfy one of the conditions from Table 1, one has $\lim_{r \rightarrow \infty} \underline{p}_{\mathbf{K}}^{(r)} = 0$, with the rate of convergence as indicated in the table. Thus, for any $\epsilon > 0$ there is a sufficiently large $d \in \mathbb{N}$ such that

$$\min_{h \in \Sigma[z]_r} (\mathbb{E}_{z \sim (\mathbf{K}, h)}[(f_i(z) - f_i(a))])^2 \leq \epsilon \quad \forall r \geq d, i \in \{0, \dots, p\}.$$

Letting h^* denote the minimizer, one has

$$|\mathbb{E}_{z \sim (\mathbf{K}, h^*)}[(f_i(z) - f_i(a))]| \leq \sqrt{\epsilon} \quad \forall r \geq d, i \in \{0, \dots, p\}.$$

To complete the proof, we simply have to associate a with an atom of the optimal atomic distribution from Corollary 1. \square

As a consequence of Theorem 4, we may obtain approximate solutions to the generalized problem of moments (17) by solving SDPs of the form:

$$\min_{h \in \Sigma[z]_r} \left\{ \int_{\mathbf{K}_0} f_0(z) h(z) d\mu(z) : \int_{\mathbf{K}_i} f_i(z) h(z) d\mu(z) = [b_i - \epsilon, b_i + \epsilon] \forall i = 1, \dots, p \right\},$$

for given $r \in \mathbb{N}$ and $\epsilon \geq 0$, and we will do precisely that in the example of Section 7.1. We remark that these SDPs are different from the ones studied by Lasserre [23], where an *outer* approximation of the cone of finite Borel measures supported on \mathbf{K} is used, whereas we use an inner approximation.

6 Modeling power

The ambiguity set \mathcal{P} defined in (15) contains all distributions supported on a convex body \mathbf{K} that have an SOS polynomial density $h \in \Sigma[z]_r$ with respect to a prescribed reference measure μ . For any fixed x , the worst-case expectation $f_{\mathbf{K}}^{(r)}(x)$ on the left-hand-side of the worst-case feasibility constraint (14) can be computed efficiently by solving the SDP (16). The ambiguity set \mathcal{P} admits several generalizations that preserve the SDP-representability of the worst-case expectation.

Moment information As already discussed in Section 5, conditions on (mixed) moment values of different random variables give rise to simple linear conditions on the polynomial coefficients of h .

Confidence information If the random vector z is known to materialize inside a given Borel set $\mathbf{C} \subset \mathbb{R}^m$ with probability $\gamma \in [0, 1]$, we can add the condition $\mathbb{P}_{z \sim (\mathbf{K}, h)}[z \in \mathbf{C}] = \gamma$ to the definition of the ambiguity set \mathcal{P} . Moreover, if the moments $m_\alpha(\mathbf{K} \cap \mathbf{C})$ of the reference measure μ over $\mathbf{K} \cap \mathbf{C}$ are either available analytically or efficiently computable for all $\alpha \in N(m, 2r)$, then this condition can be re-expressed as the following simple linear equation in the polynomial coefficients of h .

$$\int_{\mathbf{K}} \mathbf{1}_{z \in \mathbf{C}} h(z) d\mu(z) = \sum_{\alpha \in N(m, 2r)} h_\alpha m_\alpha(\mathbf{K} \cap \mathbf{C}) = \gamma$$

Upper and lower bounds on $\mathbb{P}_{z \sim (\mathbf{K}, h)}[z \in \mathbf{C}]$ can be handled similarly in the obvious manner. In the context of purely moment-based ambiguity sets, such probability bounds have been studied in [45].

Conditional probabilities Given any two Borel sets $\mathbf{C}_1, \mathbf{C}_2 \subset \mathbb{R}^m$ and a probability $\gamma \in [0, 1]$, we can also enforce the condition $\mathbb{P}_{z \sim (\mathbf{K}, h)}[z \in \mathbf{C}_2 | z \in \mathbf{C}_1] = \gamma$ in the definition of \mathcal{P} . If the moments $m_\alpha(\mathbf{K} \cap \mathbf{C}_1)$ and $m_\alpha(\mathbf{K} \cap \mathbf{C}_1 \cap \mathbf{C}_2)$ of the reference measure μ are either available analytically or efficiently computable for all $\alpha \in N(m, 2r)$, then this condition can be re-expressed as

$$\begin{aligned} \int_{\mathbf{K}} \mathbf{1}_{z \in \mathbf{C}_1 \cap \mathbf{C}_2} h(z) d\mu(z) &= \gamma \int_{\mathbf{K}} \mathbf{1}_{z \in \mathbf{C}_1} h(z) d\mu(z) \\ \Leftrightarrow \sum_{\alpha \in N(m, 2r)} h_\alpha (m_\alpha(\mathbf{K} \cap \mathbf{C}_1 \cap \mathbf{C}_2) - \gamma m_\alpha(\mathbf{K} \cap \mathbf{C}_1)) &= 0, \end{aligned}$$

which is again linear in the coefficients of h . Upper and lower bounds on conditional probabilities can be handled similarly.

Conditional moment information If it is known that $\mathbb{E}_{z \sim (\mathbf{K}, h)}[z^\beta | \mathbf{C}] = \gamma$ for some $\beta \in \mathbb{N}_0^n$, Borel set $\mathbf{C} \subset \mathbb{R}^m$ and $\gamma \in \mathbb{R}$, while the moments $m_{\alpha+\beta}(\mathbf{K} \cap \mathbf{C})$ of the reference measure μ over set $\mathbf{K} \cap \mathbf{C}$ are either available analytically or efficiently computable for all $\alpha \in N(m, 2r)$, then one can add the following condition to the ambiguity set \mathcal{P} , which is linear in the coefficients of h .

$$\begin{aligned} \int_{\mathbf{K}} z^\beta \mathbf{1}_{z \in \mathbf{C}} h(z) d\mu(z) &= \gamma \int_{\mathbf{K}} \mathbf{1}_{z \in \mathbf{C}} h(z) d\mu(z) \\ \Leftrightarrow \sum_{\alpha \in N(m, 2r)} h_\alpha (m_{\alpha+\beta}(\mathbf{K} \cap \mathbf{C}) - \gamma m_\alpha(\mathbf{K} \cap \mathbf{C})) &= 0, \end{aligned}$$

Multiple reference measures The distributions in the ambiguity set \mathcal{P} defined in (15) depend both on the reference measure μ as well as the density function h . A richer ambiguity set can be constructed by specifying multiple reference measures μ_i with corresponding density functions $h^i \in \Sigma[z]_r$, $i = 1, \dots, p$. The distributions in the resulting ambiguity set are of the form $\sum_{i=1}^p h^i \cdot d\mu_i$. If the moments $m_\alpha^i(\mathbf{K})$ of the reference measure μ_i over \mathbf{K} are either available analytically or efficiently computable for all $\alpha \in N(m, 2r)$ and $i = 1, \dots, p$, then the normalization constraint can be recast as

$$\sum_{\alpha \in N(m, 2r)} h_\alpha^i m_\alpha^i(\mathbf{K}) = \gamma_i \quad \forall i = 1, \dots, p \quad \text{and} \quad \sum_{i=1}^p \gamma_i = 1,$$

where $\gamma = (\gamma_1, \dots, \gamma_p) \geq 0$ constitutes an auxiliary decision vector. The resulting ambiguity set can be interpreted as a convex combination of p ambiguity sets of the form (15) and thus lends itself for modeling multimodality information; see, *e.g.*, [20]. In this case, γ_i captures the probability of the i -th mode, which may itself be uncertain. Thus, γ should range over a subset of the probability simplex, *e.g.*, a ϕ -divergence uncertainty set of the type studied in [2].

Marginal distributions It is often easier to estimate the marginal distributions of all m components of a random vector z instead of the full joint distribution. Marginal distribution information can also be conveniently encoded in ambiguity sets of the type (15). To see this, assume that the marginal distribution of z_i is given by μ_i and is supported on a compact interval $\mathbf{K}_i \subset \mathbb{R}$, $i = 1, \dots, m$. In this case it makes sense to set $\mathbf{K} = \times_{i=1}^m \mathbf{K}_i$ and to define the reference measure μ through $d\mu = \prod_{i=1}^m d\mu_i$. Thus, μ coincides with the product of the known marginals. The requirement

$$\int_{\times_{j \neq i} \mathbf{K}_j} h(z) \prod_{j \neq i} d\mu_j(z_j) = 1 \quad \forall z_i \in \mathbf{K}_i, \quad \forall i = 1, \dots, m$$

then ensures that the marginal distribution of z_i under $h \cdot d\mu$ exactly matches μ_i . If the moments $m_{\alpha_i}(\mathbf{K}_i)$ of the marginal distribution μ_i over \mathbf{K}_i are either available analytically or efficiently computable for all $\alpha_i = 1, \dots, 2r$, then the above condition simplifies to the linear equations

$$\sum_{\substack{\alpha \in N(m, 2r) \\ \alpha_i = 0}} h_\alpha \prod_{j \neq i} m_{\alpha_j}(\mathbf{K}_j) = 1 \quad \text{and} \quad \sum_{\substack{\alpha \in N(m, 2r) \\ \alpha_i = \ell}} h_\alpha \prod_{j \neq i} m_{\alpha_j}(\mathbf{K}_j) = 0 \quad \forall \ell = 1, \dots, 2r, \quad \forall i = 1, \dots, m. \quad (19)$$

Situations where the marginals of groups of random variables are known can be handled analogously. Note that when all marginals are known, there is only ambiguity about the dependence structure or *copula* of the components of z [40]. Quantifying the worst-case copula amounts to solving a so-called Fréchet problem. In distributionally robust optimization, Fréchet problems with discrete marginals or approximate marginal matching conditions have been studied in [13, 12, 43].

Besides the ambiguity set \mathcal{P} , the constraint function f also admits some generalizations that preserve the SDP-representability of the worst-case expectation in (14).

Uncertainty quantification problems If the constraint function f in (14) is given by $f(x, z) = \mathbf{1}_{z \in \mathbf{C}}$ for some Borel set $\mathbf{C} \subset \mathbb{R}^m$, then the worst-case expectation reduces to the worst-case probability of the set \mathbf{C} . Moreover, if the moments $m_\alpha(\mathbf{K} \cap \mathbf{C})$ of the reference measure μ over $\mathbf{K} \cap \mathbf{C}$ are either available analytically or efficiently computable for all $\alpha \in N(m, 2r)$, then the worst-case probability can be computed by solving a variant of the SDP (16) with the alternative objective function

$$\sum_{\alpha \in N(m, 2r)} h_\alpha m_\alpha(\mathbf{K} \cap \mathbf{C}).$$

7 Numerical experiments

In the following we will exemplify the proposed approach to distributionally robust optimization in the context of financial portfolio analysis (Section 7.1) and risk aggregation (Section 7.2).

7.1 Portfolio analysis

Consider a portfolio optimization problem, where the decision vector $x \in \mathbb{R}^n$ captures the percentage weights of the initial capital allocated to n different assets. By definition, one thus has $x_i \in [0, 1]$ for all $i = 1, \dots, n$ and $\sum_i x_i = 1$. We assume that the asset returns $r_i = (u_i + l_i)/2 + z_i(u_i - l_i)/2$ depend linearly on some uncertain risk factors $z_i \in [-1, 1]$ for all $i = 1, \dots, n$, where l_i and u_i represent known upper and lower bounds on the i -th return, respectively. In this framework, we denote by $z \in \mathbb{R}^n$ the vector of all risk factors and by $\mathbf{K} = [-1, 1]^n$ its support. Moreover, the portfolio return can be expressed as

$$f(x, z) = \sum_{i=1}^n x_i \cdot ((u_i + l_i)/2 + z_i(u_i - l_i)/2).$$

Unless otherwise stated, we set μ to the Lebesgue measure on \mathbb{R}^n . Modeling the probability density functions as SOS polynomials allows to account for various statistical properties and stylized facts of real asset returns as described in [6]. For example, the proposed approach can conveniently capture *gain loss asymmetry*, *i.e.*, the observation that large drawdowns in stock prices and stock index values are more common than equally large upward movements. This feature can be modeled by assigning a higher probability to an individual asset's large upward returns than to its low downward returns. Specifically, the ambiguity set may include the conditions $\mathbb{P}_{z \sim (\mathbf{K}, h)}(z_i \leq a_i) = \gamma_1$ and $\mathbb{P}_{z \sim (\mathbf{K}, h)}(z_i \geq b_i) = \gamma_2$ for some thresholds $a_i < b_i$ and confidence levels $\gamma_1 > \gamma_2$.

Similarly, our approach can handle *correlations of extreme returns*. As pointed in [6], in spite of the widespread use of the covariance matrix, 'in circumstances when stock prices undergo large fluctuations [...], a more relevant quantity is the conditional probability of a large (negative) return in one stock given a large negative movement in another stock.' An example constraint on the conditional probability of one asset's low performance given another assets' lower performance is $\mathbb{P}_{z \sim (\mathbf{K}, h)}(z_i \leq \underline{r}_i | z_j \leq \underline{r}_j) \leq \gamma$, where \underline{r}_i and \underline{r}_j are given thresholds, while γ is a confidence level.

In this numerical experiment we evaluate the probability that the return of a fixed portfolio x materializes below a prescribed threshold \underline{r} , that is, we evaluate the worst case of the probability

$$\mathbb{P}_{z \sim (\mathbf{K}, h)}(r(x, z) \leq \underline{r})$$

over an ambiguity set \mathcal{P} of the form (15) with the additional moment constraints $\mathbb{E}_{z \sim (\mathbf{K}, h)}[z^{\beta_i}] = \gamma_i$ for some given exponents $\beta_i \in \mathbb{N}_0^n$ and targets $\gamma_i \in \mathbb{R}$ for $i = 1, \dots, p$. This corresponds to computing the integral of the density function over the knapsack polytope $\mathbf{K} \cap \mathbf{A}(x, u, l, \underline{r})$, where

$$\mathbf{A}(x, u, l, \underline{r}) = \left\{ z \in \mathbb{R}^n : \sum_{i=1}^n x_i(u_i - l_i)z_i/2 \leq \underline{r} - \sum_{i=1}^n x_i(u_i + l_i)/2 \right\}$$

represents a halfspace in \mathbb{R}^n that depends on the fixed portfolio x , the return bounds $l = (l_1, \dots, l_n)$ and $u = (u_1, \dots, u_n)$, and the threshold \underline{r} . To formulate this problem as an SDP, we first need to compute the moments of the monomials with respect to the Lebesgue measure over the given knapsack polytope by using the results of [29]. The worst-case probability problem can then be

Table 3: Worst-case probabilities for the portfolio return falling below r computed by directly solving the SDP (20) ($R = 1$) and by using Algorithm 1 ($R > 1$). Missing values indicate the occurrence of numerical instability.

r	R													
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
0	0.17													
1	0.39	0.47	0.49	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.57
2	0.48	0.52	0.54	0.55	0.56	0.57	0.60	-	-	-	-	-	-	-
3	0.50	0.56	0.59	0.62	-	-	-	-	-	-	-	-	-	-
4	0.53	0.58	0.61	-	-	-	-	-	-	-	-	-	-	-
5	0.55	0.59	-	-	-	-	-	-	-	-	-	-	-	-
6	0.56	0.60	-	-	-	-	-	-	-	-	-	-	-	-
7	0.58	-	-	-	-	-	-	-	-	-	-	-	-	-
8	0.59	-	-	-	-	-	-	-	-	-	-	-	-	-
9	0.59	-	-	-	-	-	-	-	-	-	-	-	-	-
10	0.60	-	-	-	-	-	-	-	-	-	-	-	-	-
11	0.61	-	-	-	-	-	-	-	-	-	-	-	-	-
12	0.61	-	-	-	-	-	-	-	-	-	-	-	-	-
∞	0.69	-	-	-	-	-	-	-	-	-	-	-	-	-

reformulated as the SDP

$$\begin{aligned}
& \sup_{h(z)} \sum_{\alpha \in N(n, 2r)} h_{\alpha} m_{\alpha}(\mathbf{K} \cap \mathbf{A}(x, u, l, \underline{r})) \\
& \text{s.t.} \quad \sum_{\alpha \in N(n, 2r)} h_{\alpha} m_{\alpha}(\mathbf{K}) = 1, \\
& \quad \sum_{\alpha \in N(n, 2r)} h_{\alpha} m_{\alpha + \beta_i}(\mathbf{K}) = \gamma_i \quad \forall i = 1, \dots, p, \\
& \quad \sum_{\alpha \in N(n, 2r)} h_{\alpha} z^{\alpha} \in \Sigma[z]_r.
\end{aligned} \tag{20}$$

In the numerical experiment we assume that there are $n = 2$ assets with lower and upper return bounds $l = (0.8, 0.7)^{\top}$ and $u = (1.2, 1.3)^{\top}$, respectively. We evaluate the probability that the return of the fixed portfolio $x = (0.75, 0.25)^{\top}$ falls below the threshold $\underline{r} = 0.9$ (the minimum possible return of the portfolio is 0.775). We assume that the only known moment information about the asset returns is that their means both vanish, that is, we set $p = 2$, $\beta_1 = (1, 0)$, $\beta_2 = (0.1)$ and $\gamma_1 = \gamma_2 = 0$. Table 3 reports the *exact* optimal values of the SDP (20) for $r = 1, \dots, 12$ ($R = 1$). The value in the last row of the table (labeled $r = \infty$) provides the worst-case probability across *all* distributions satisfying the prescribed moment conditions (not only those with a polynomial density) and was computed using the methods described in [19]. In this case, one can also show that there exists a worst-case distribution with only two atoms. It assigns probability 0.31 to the scenario $z = (1, 1)^{\top}$ and probability 0.69 to the scenario $z = (0.28, 0.28)^{\top}$. We further computed the worst-case probabilities *approximately* by using Algorithm 1 from Section 3 for SOS parameters $r = 1, \dots, 6$ and for up to $R = 14$ iterations. The results are given in Table 3 ($R > 1$).

7.2 Risk aggregation

In the second experiment we study the risk aggregation problem of an insurer holding a portfolio of different random losses z_i , $i = 1, \dots, n$, corresponding to different types of insurance claims, *e.g.*,

life, vehicle, health or home insurance policies, etc. Inspired by [43, § 6], we aim to estimate the worst-case probability that the sum of the n losses exceeds a critical threshold $b = 10$ beyond which the insurance company would be driven into illiquidity. Formally, we aim to maximize

$$\mathbb{P}_{z \sim (\mathbf{K}, h)}(z_1 + \dots + z_n \geq b) \quad (21)$$

across all distributions in an ambiguity set \mathcal{P} , which reflects the prior distributional information available to the insurer. We will consider different models for the domain \mathbf{K} of $z = (z_1, \dots, z_n)$, the reference measure μ on \mathbf{K} and the ambiguity set \mathcal{P} . Throughout the experiments we will always assume that the reference measure is separable with respect to the losses, that is, we assume that

$$d\mu(z) = \varrho_1(z_1) \cdots \varrho_n(z_n) dz,$$

where ϱ_i denotes a given density function (with respect to the Lebesgue measure) of the random variables z_i for each $i = 1, \dots, n$. We will consider the following complementary settings:

1. *Lognormal densities:* We set $\mathbf{K} = \mathbb{R}_+^n$ and let ϱ_i be a lognormal density function defined earlier in (11), but repeated here for convenience:

$$\varrho_i(z_i) = \frac{1}{z_i v_i \sqrt{2\pi}} \exp\left(-\frac{(\log(z_i) - \bar{z}_i)^2}{2v_i^2}\right), \quad (22)$$

where \bar{z}_i and v_i represent prescribed location and scale parameters, $i = 1, \dots, n$.

2. *Exponential densities:* We set $\mathbf{K} = \mathbb{R}_+^n$ and let ϱ_i be the exponential density function with unit rate parameter defined through $\varrho_i(z_i) = \exp(-z_i)$, $i = 1, \dots, n$. The resulting reference measure is intimately related to the orthogonal Laguerre polynomials.
3. *Uniform densities:* We set $\mathbf{K} = [0, M]^n$ for some constant $M > 0$ and let ϱ_i be the uniform density function defined through $\varrho_i(z_i) = 1/M$, $i = 1, \dots, n$. Note that under this choice the reference measure is proportional to the Lebesgue measure.

In order to reformulate the risk aggregation problem as a tractable SDP, we need the moments of the reference measure μ over the hypercube \mathbf{K} and over the knapsack polytope $\mathbf{K} \cap \mathbf{C}$, where

$$\mathbf{C} = \{z \in \mathbb{R}^n : z_1 + \dots + z_n \geq b\}.$$

For all classes of density functions described above, the moments of μ are indeed accessible. Specifically, under the lognormal densities, the moments of μ over \mathbf{K} are given by (12), and repeated here for convenience:

$$m_\alpha(\mathbf{K}) = \int_{\mathbf{K}} \left(\prod_{i=1}^n z_i^{\alpha_i} \right) \prod_{i=1}^n \frac{1}{z_i v_i \sqrt{2\pi}} \exp\left(-\frac{(\log(z_i) - \bar{z}_i)^2}{2v_i^2}\right) dz = \prod_{i=1}^n \exp(\alpha_i \bar{z}_i + (\alpha_i v_i)^2 / 2).$$

Moreover, the moments of μ over $\mathbf{K} \cap \mathbf{C}$ can be expressed as

$$\begin{aligned} m_\alpha(\mathbf{K} \cap \mathbf{C}) &= m_\alpha(\mathbf{K}) - m_\alpha(\mathbf{K} \setminus \mathbf{C}) \\ &= m_\alpha(\mathbf{K}) - \int_{\mathbf{K} \setminus \mathbf{C}} \left(\prod_{i=1}^n z_i^{\alpha_i} \right) \prod_{i=1}^n \frac{1}{z_i v_i \sqrt{2\pi}} \exp\left(-\frac{(\log(z_i) - \bar{z}_i)^2}{2v_i^2}\right) dz. \end{aligned}$$

To evaluate the integral in the last expression, we use the MATLAB routine `adsimp(·)` from [14]. Furthermore, under the exponential and the uniform densities, the moments of the reference measure μ over \mathbf{K} and $\mathbf{K} \cap \mathbf{C}$ are all available in closed form.

We assume that the insurance company is able to estimate the marginal distributions of the individual losses either exactly or approximately by using a combination of statistical analysis and probabilistic modeling. However, the insurer has no information about the underlying copula. This type of distributional information is often justified in practice because obtaining reliable marginal information requires significantly less data than obtaining exact dependence structures; see, *e.g.*, [28]. Throughout the experiment we assume that there are $n = 2$ random losses governed by lognormal probability density functions of the form (22) with parameters $\bar{z}_1 = -0.3$, $\bar{z}_2 = 0.4$, $v_1 = 0.8$ and $v_2 = 0.5$. The ambiguity set \mathcal{P} then contains all distributions of the form $h \cdot d\mu$, $h \in \Sigma[z]_r$, under which the marginals of the losses follow the prescribed lognormal distributions either exactly or approximately. More precisely, we model the marginal distributional information as follows:

1. *Marginal distribution matching:* The lognormal distributions of the individual losses are matched *exactly* by any distribution $h \cdot d\mu$ in the ambiguity set. This can be achieved by defining the reference measure μ as the product of the marginal lognormal distributions and by requiring that h satisfies (19). Note that under the alternative reference measures corresponding to the exponential or uniform density functions, lognormal marginals cannot be matched exactly with polynomial densities of any degrees. Note also that an exact matching of (non-discrete) marginal distributions cannot be enforced with the existing numerical techniques for solving Fréchet problems proposed in [13, 12, 43].
2. *Marginal moment matching:* The marginals of the individual losses have the same moments of order 0, 1 or 2 as the prescribed lognormal distributions. Note that this kind of moment matching can be enforced under any of the reference measures corresponding to lognormal, exponential or uniform density functions. Moreover, moment matching is also catered for in [43] bar the extra requirement that the joint distribution of the losses must have an SOS polynomial density.
3. *Marginal histogram matching:* We may associate a histogram with each marginal lognormal distribution as illustrated in Figure 1 and require that the marginals of the losses under the joint distribution $h \cdot d\mu$ have the same histograms. This condition can be enforced under any of the reference measures corresponding to lognormal, exponential or uniform density functions. In the numerical experiments, we use histograms with 20 bins of width 0.25 starting at the origin. Histogram matching is also envisaged in [43].




Fig. 1: Histograms of the lognormal marginal distributions of z_1 (left) and z_2 (right).

For $\mathbf{K} = \mathbb{R}_+^n$ and the reference measure corresponding to the lognormal density functions, the worst-case values of the probability (21) are reported in Table 4. Results are shown for $r \leq$

5, which corresponds to polynomial densities of degrees at most 10. The last row of the table ($r = \infty$) provides the worst-case probabilities across *all* distributions satisfying the prescribed moment or histogram conditions (not only those with a polynomial density) and was computed using the methods described in [43]. Note that under moment matching up to order 2, the worst-case probability for $r = 5$ amounts to 0.0021, as opposed to the much higher probability of 0.0615 obtained with the approach from [43]. A similar observation holds for histogram matching. The requirement that the distributions in the ambiguity set be sufficiently regular in the sense that they admit a polynomial density function with respect to the reference measure is therefore restrictive and effectively rules out pathological discrete worst-case distributions. Moreover, the worst-case probabilities under exact distribution matching and under histogram matching are of the same order of magnitude for all $r \leq 5$ but significantly smaller than the worst-case probability under histogram matching for $r = \infty$. A key question to be asked in practice is thus whether one deems the class of distributions $h \cdot d\mu$ with $h \in \Sigma[z]_r$ to be rich enough to contain all ‘reasonable’ distributions.

Table 4: Worst-case probabilities for the lognormal reference measure.

r	Moment matching up to order			Histogram matching	Distribution matching
	0	1	2		
0	0.0017	0.0017	0.0017	0.0017	0.0017
1	0.1432	0.0042	0.0017	0.0017	0.0017
2	0.8255	0.0106	0.0020	0.0019	0.0018
3	0.9982	0.0114	0.0021	0.0022	0.0019
4	1.0000	0.0117	0.0021	0.0026	0.0023
5	1.0000	0.0118	0.0021	0.0026	0.0023
∞	1.0000	1.0000	0.0615	0.0198	n/a

Table 5: Worst-case probabilities for the exponential reference measure.

r	Moment matching up to order			Histogram matching		
	0	1	2	ℓ_1 -dist. ≤ 0.1	ℓ_1 -dist. ≤ 0.05	ℓ_1 -dist. ≤ 0.02
0	0.0005	-	-	-	-	-
1	0.0214	0.0147	-	-	-	-
2	0.2058	0.0823	-	-	-	-
3	0.6481	0.1484	-	-	-	-
4	0.9393	0.1497	0.0086	-	-	-
5	0.9953	0.1699	0.0104	-	-	0.0104
6	0.9998	0.1709	0.0139	-	-	-
7	1.0000	0.1800	0.0158	-	-	-
8	1.0000	0.1860	0.0182	0.0802	-	-
9	1.0000	0.1862	0.0207	0.1076	-	-
10	1.0000	0.1928	0.0224	0.1144	0.0515	-
11	1.0000	0.1968	0.0244	0.1156	0.0633	0.0204
12	1.0000	0.1971	0.0262	0.1160	0.0652	0.0320
∞	1.0000	1.0000	0.0615	n/a	n/a	n/a

Table 5 reports the worst-case probabilities corresponding to the reference measure on $\mathbf{K} = \mathbb{R}_+^r$ induced by the exponential density functions. For low values of r , the polynomial densities lack the necessary degrees of freedom to match all imposed moment constraints. In these situations, the

worst-case probability problem becomes infeasible (indicated by '-'). When feasible, however, we managed to solve the problem for r up to 12. The density functions corresponding to large values of r are highly flexible and thus result in worst-case probabilities that are closer to those obtained by the benchmark method from [43], which relaxes the restriction to a subspace of polynomial densities. Similar phenomena are also observed in the context of histogram matching. It was impossible to match the prescribed histogram probabilities exactly for all $r \leq 12$. We thus relaxed the histogram matching conditions in the definition of the ambiguity set to allow for densities whose implied marginal histograms are within a prescribed ℓ_1 -distance from the target histograms. This approximate histogram matching condition is easily captured in our framework and gives rise to a few extra linear constraints on the coefficients of the polynomial density function. Table 5 reports the worst-case probabilities for three different tolerances on the histogram mismatch in terms of the ℓ_1 -distance. We observe that the resulting worst-case probabilities are significantly larger than those obtained under the lognormal reference measure and increase with the ℓ_1 -tolerance.

Finally, Table 6 reports the worst-case probabilities corresponding to the uniform reference measure on $\mathbf{K} = [0, 10]^2$. The results are qualitatively similar to those of Table 5, but they also show that the choice of the reference measure plays an important role when r is small.

Table 6: Worst-case probabilities for the uniform reference measure.

r	Moment matching up to order			Histogram matching		
	0	1	2	ℓ_1 -dist. ≤ 0.1	ℓ_1 -dist. ≤ 0.05	ℓ_1 -dist. ≤ 0.02
0	0.5000	-	-	-	-	-
1	0.9082	-	-	-	-	-
2	0.9933	-	-	-	-	-
3	0.9997	0.0304	-	-	-	-
4	1.0000	0.1035	-	-	-	-
5	1.0000	0.1340	-	-	-	-
6	1.0000	0.1612	0.0089	-	-	-
7	1.0000	0.1783	0.0166	-	-	-
8	1.0000	0.1935	0.0192	-	-	-
9	1.0000	0.2042	0.0216	0.0738	-	-
10	1.0000	0.2133	0.0274	0.1066	0.0407	-
11	1.0000	0.2202	0.0292	0.1142	0.0609	-
12	1.0000	0.2274	0.0311	0.1163	0.0653	0.0178
∞	1.0000	1.0000	0.0615	n/a	n/a	n/a

8 Conclusions

In this paper, we present first steps towards using SOS polynomial densities in distributionally robust optimization for problems that display a polynomial dependence on the uncertain parameters. The main advantages of this approach may be summarized as follows:

1. The proposed framework is tractable (in the sense of polynomial-time solvability) for SOS density functions of any fixed degree.
2. The approach offers considerable modeling flexibility. Specifically, one may conveniently encode various salient features of the unknown distribution of the uncertain parameters through linear constraints and/or linear matrix inequalities.

3. In the limit as the degree of the SOS density functions tends to infinity, one recovers the usual robust counterpart or generalized moment problem. One may therefore view the degree of the density as a tuning parameter that captures the model's 'level of conservativeness.'

The approach also suffers from shortcomings that necessitate further work and insights:

1. The approach is not applicable to objective or constraint functions that display a general (decision-dependent) piecewise polynomial dependence on the uncertain parameters as is the case for the recourse functions of linear two-stage stochastic programs.
2. The proposed distributionally robust optimization problems can be reduced to generalized eigenvalue problems or even semidefinite programs of large sizes that are often poorly conditioned. We have introduced a heuristic solution procedure in Section 3 as a practical remedy, but additional work is required to make the approach more scalable.

Acknowledgements Etienne de Klerk would like to thank Dorota Kurowicka and Jean Bernard Lasserre for valuable discussions and references. Daniel Kuhn gratefully acknowledges financial support from the Swiss National Science Foundation under grant BSCG10_157733.

References

1. Ben-Tal, A., El Ghaoui, L., Nemirovski, A., *Robust Optimization*, Princeton University Press (2009).
2. Ben-Tal, A. den Hertog, D., de Waegenaere, A., Melenberg, B., Rennen, G. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science* 59(2), 341–357 (2013).
3. Bertsimas, D., Popescu, I. On the relation between option and stock prices: a convex optimization approach, *Operations Research* 50(2), 358–374 (2002).
4. Bertsimas, D., Popescu, I. *Optimal inequalities in probability theory: a convex optimization approach*, *SIAM Journal on Optimization* 15(3), 780–804 (2005).
5. Birge, J.R., Louveaux, F. *Introduction to Stochastic Programming*, Springer (1997).
6. Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* 1, 223–236 (2001).
7. de Klerk, E., Laurent, M. Comparison of Lasserre's measure-based bounds for polynomial optimization to bounds obtained by simulated annealing. *Mathematics of Operations Research*, to appear. Preprint available at <http://arxiv.org/abs/1703.00744>
8. de Klerk, E., Laurent, M. Worst-case examples for Lasserre's measure-based hierarchy for polynomial optimization on the hypercube (2018). Preprint available at <http://arxiv.org/abs/1804.05524>
9. de Klerk, E., Laurent, M., Sun, Z. Convergence analysis for Lasserre's measure-based hierarchy of upper bounds for polynomial optimization. *Mathematical Programming Series A* 162(1), 363–392 (2017).
10. Dantzig, G.B. Linear programming under uncertainty. *Management Science* 1(3-4), 197–206 (1955).
11. Delage, E., Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* 58(3), 595–612 (2010).
12. Doan, X.V., Li, X., Natarajan, K. Robustness to dependency in portfolio optimization using overlapping marginals. *Operations Research* 63(6), 1468–1488 (2015).
13. Doan X.V., Natarajan K. On the complexity of nonoverlapping multivariate marginal bounds for probabilistic combinatorial optimization problems. *Operations Research* 60(1),138–49 (2012).
14. Genz, A., Cools, R. An adaptive numerical cubature algorithm for simplices. *ACM Transactions on Mathematical Software* 29(3), 297–308 (2003).
15. Goh, J., Sim, M. Distributionally robust optimization and its tractable approximations, *Operationis Research* 58(4), 902–917 (2010).
16. Grötschel, M., Lovász, L., Schrijver, A. *Geometric Algorithms and Combinatorial Optimization*. Springer (1988).
17. Grundmann, A., Moeller, H.M. Invariant integration formulas for the n -simplex by combinatorial methods. *SIAM Journal on Numerical Analysis* 15, 282–290 (1978).
18. Hanasusanto, G.A., Roitch, V., Kuhn, D., Wiesemann, W. A distributionally robust perspective on uncertainty quantification and chance constrained programming. *Mathematical Programming Series B* 151(1), 35–62 (2015).

19. Hanasusanto, G.A., Roitch, V., Kuhn, D., Wiesemann, W. Ambiguous joint chance constraints under mean and dispersion information. *Operations Research* 65(3), 751–767 (2017).
20. Hanasusanto, G.A., Kuhn, D., Wallace, S.W., Zymler, S. Distributionally robust multi-item newsvendor problems with multimodal demand distributions. *Mathematical Programming Series A* 152(1), 1–32 (2015).
21. Kroo, A., Szilárd, R. On Bernstein and Markov-type inequalities for multivariate polynomials on convex bodies. *Journal of Approximation Theory* 99(1), 134–152 (1999).
22. Lasserre, J.B., Zeron, E.S. Solving a class of multivariate integration problems via Laplace techniques. *Applicaciones Mathematicae* 28(4), 391–405 (2001).
23. Lasserre, J.B. A semidefinite programming approach to the generalized problem of moments. *Mathematical Programming Series B* 112, 65–92 (2008).
24. Lasserre, J.B. A new look at nonnegativity on closed sets and polynomial optimization. *SIAM Journal on Optimization* 21(3), 864–885 (2011).
25. Lasserre, J.B. The \mathbf{K} -moment problem for continuous linear functionals. *Transactions of the American Mathematical Society* 365(5), 2489–2504 (2012).
26. Lasserre, J.B., Weisser, T. Representation of distributionally robust chance-constraints (2018). Preprint available at <http://arxiv.org/abs/1803.11500>
27. Li, B., Jiang, R., Mathieu, J.L. Ambiguous risk constraints with moment and unimodality information, *Mathematical Programming Series A*, to appear (2017). Preprint available at http://www.optimization-online.org/DB_FILE/2016/09/5635.pdf
28. McNeil, A., Frey, R., Embrechts, P. *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton University Press (2015).
29. Marichal, J.-L., Mossinghoff, M.J. Slices, slabs, and sections of the unit hypercube. *Online Journal of Analytic Combinatorics* 3, 1–11 (2008).
30. Mevissen, M., Ragnoli, E., Yu, J.Y. Data-driven distributionally robust polynomial optimization, *Advances in Neural Information Processing Systems (NIPS)* 26 (2013).
31. Mohajerin Esfahani, P., Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations, *Mathematical Programming Series A*, to appear (2017). Preprint available at <https://arxiv.org/abs/1505.05116>
32. Natarajan, K., Pachamanova, D., Sim, M. Constructing risk measures from uncertainty sets, *Operations Research* 57(5), 1129–1141 (2009).
33. Pflug, G.C., Pichler, A., Wozabal, D. The $1/N$ investment strategy is optimal under high model ambiguity, *Journal of Banking & Finance* 36(2), 410–417 (2012).
34. Pflug, G.C., Wozabal, D. Ambiguity in portfolio selection, *Quantitative Finance* 7, 435–442 (2007).
35. Popescu, I. A semidefinite programming approach to optimal-moment bounds for convex classes of distributions, *Mathematics of Operations Research* 30(3), 632–657 (2005).
36. Prékopa, A. *Stochastic Programming*, Kluwer Academic Publishers (1995).
37. Rockafellar, R.T. *Convex Analysis*, Princeton University Press (1970).
38. Rogosinski, W.W. Moments of non-negative mass, *Proceedings of the Royal Society A* 245, 1–27 (1958).
39. Scarf, H. A min-max solution of an inventory problem, in H. Scarf, K. Arrow, and S. Karlin (Eds.), *Studies in the Mathematical Theory of Inventory and Production*, Volume 10, pp. 201–209. Stanford University Press (1958).
40. Sklar, A. Fonctions de répartition à n dimensions et leurs marges, *Publications de l'Institut de Statistique de l'Université de Paris* 8, 229–231 (1959).
41. Shapiro, A. *On duality theory of conic linear problems*, Semi-Infinite Programming: Recent Advances (M.Á. Goberna and M.A. López, eds.), Springer, 135–165 (2001).
42. Shapiro, A., Dentcheva, D., Ruszczyński, A. *Lectures on Stochastic Programming: Modeling and Theory*, SIAM (2009).
43. Van Parys, B.P.G., Goulart, P.J., Embrechts, P. Fréchet inequalities via convex optimization (2016). Preprint available at http://www.optimization-online.org/DB_FILE/2016/07/5536.pdf
44. Van Parys, B.P.G., Goulart, P.J., Kuhn, D. Generalized Gauss inequalities via semidefinite programming. *Mathematical Programming Series A* 156(1-2), 271–302 (2016).
45. Wiesemann, W., Kuhn, D., Sim, M. Distributionally robust convex optimization. *Operations Research* 62(6), 1358–1376 (2014).
46. Žáčková, J. On minimax solutions of stochastic linear programming problems, *Časopis pro pěstování matematiky* 91, 423–430 (1966).
47. Zuluaga, L., Peña J.F. A conic programming approach to generalized Tchebycheff inequalities, *Mathematics of Operations Research* 30(2), 369–388 (2005).