Decision Theory Meets Linear Optimization Beyond Computation

Christoph Jansen(✉), Thomas Augustin, and Georg Schollmeyer

Department of Statistics, LMU, Munich, Germany
{christoph.jansen,augustin,georg.schollmeyer}@stat.uni-muenchen.de

Abstract. The paper is concerned with decision making under complex uncertainty. We consider the Hodges and Lehmann-criterion relying on uncertain classical probabilities and Walley's maximality relying on imprecise probabilities. We present linear programming based approaches for computing optimal acts as well as for determining least favorable prior distributions in finite decision settings. Further, we apply results from duality theory of linear programming in order to provide theoretical insights into certain characteristics of these optimal solutions. Particularly, we characterize conditions under which randomization pays out when defining optimality in terms of the Gamma-Maximin criterion and investigate how these conditions relate to least favorable priors.

Keywords: Linear programming · Decision making · Least favorable prior · Duality · Maximality · Imprecise probabilities · Gamma-maximin · Hodges & Lehmann

1 Introduction

Many problems arising in modern sciences, e.g. estimation and hypothesis testing in statistics or modeling an agent's preferences in economics, can be embedded in the formal framework of *decision theory under uncertainty*. However, as the specification of a *precise* (i.e. classical) probability measure on the space of uncertain states often turns out to be too restrictive from an applicational point of view, decision theory using *imprecise probabilities* (for a survey see, e.g., [12]) has become a more and more attractive modeling tool recently. For determining optimal decisions with respect to the complex decision criteria particularly (but not exclusively) arising in the context of the theory of imprecise probabilities, *linear programming theory* (see, e.g., [15]) often turns out to be well-suited: By embedding decision problems into this general optimization framework, one can draw on the whole theoretical toolbox of this well-investigated mathematical discipline. Particularly, this allows for a computational treatment of complex decision making problems in standard software (e.g. MATLAB or for statisticians R) and, therefore, helps in order to make the abstract theory applicable for practitioners. Accordingly, there exists plenty of literature on linear optimization driven algorithms for facing complex decision problems. Examples include [6,13]. A survey is given in [5].

A. Antonucci et al. (Eds.): ECSQARU 2017, LNAI 10369, pp. 329–339, 2017.
DOI: 10.1007/978-3-319-61581-3_30

However, quite similar to characterizations of imprecise probabilities and natural extensions in [17, Chap. 4] and [14], the opportunities of using linear programming in decision theory are by far not exhausted by producing powerful algorithms (see [18, p. 402]). Instead, applying basic results on duality from linear programming theory (such as, e.g., the *complementary slackness* property, see, e.g., [15, Sect. 5.5]) can often provide theoretical insights on both the connection between different decision criteria and the specific properties shared by all optimal solutions with respect to a certain criterion.

The paper is structured as follows: In Sect. 2, we recall the classical model of finite decision theory as well as the extended version of the model allowing for randomized acts. In Sect. 3, we give a linear program for determining optimal randomized acts with respect to a decision criterion of Hodges and Lehmann which tries to cope with uncertain prior probabilistic information and investigate the corresponding dual programming problem. In Sect. 4, we consider the case of decision making under imprecise probabilistic information. Particularly, we present an algorithm for checking maximality of pure acts in one single linear program in Sect. 4.1 and use duality theory for deriving connections between least favorable prior distributions and the Gamma-Maximin criterion in Sect. 4.2. Finally, Sect. 5 is preserved for concluding remarks.

2 The Basic Model

Throughout the paper, we consider the standard model of *finite* decision theory: An *agent* (or *decision maker*) has to decide which *act* a_i to pick from a finite set $\mathbb{A} = \{a_1, \dots, a_n\}$. However, the *utility* of the chosen act depends on which *state of nature* from a finite set $\Theta = \{\theta_1, \dots, \theta_m\}$ corresponds to the true description of reality. Specifically, we assume that the utility of every pair $(a, \theta) \in \mathbb{A} \times \Theta$ can be evaluated by a *known* real-valued *cardinal utility function* $u : \mathbb{A} \times \Theta \to \mathbb{R}$. For simplicity, we will often use the notation $u_{ij} := u(a_i, \theta_j)$, where $i = 1, \dots, n$ and $j = 1, \dots, m$. The structure of the basic model and a running example repeatedly considered throughout the paper are visualized in Table 1. For every act $a \in \mathbb{A}$, the utility function u is naturally associated with a random variable $u_a : (\Theta, 2^{\Theta}) \to \mathbb{R}$ defined by $u_a(\theta) := u(a, \theta)$ for all $\theta \in \Theta$. Similarly, for every $\theta \in \Theta$, we can define a random variable $u^{\theta} : (\mathbb{A}, 2^{\mathbb{A}}) \to \mathbb{R}$ by setting $u^{\theta}(a) := u(a, \theta)$ for all $a \in \mathbb{A}$.

Depending on the context, we also allow for *randomized acts*, i.e. classical probability measures λ on $(\mathbb{A}, 2^{\mathbb{A}})$. Choosing λ is then interpreted as leaving your final decision to a random experiment which yields act a_i with probability $\lambda(\{a_i\})$. We denote the set of randomized acts on $(\mathbb{A}, 2^{\mathbb{A}})$ by $G(\mathbb{A})$.

The utility function u on $\mathbb{A} \times \Theta$ is then extended to a utility function $G(u)$ on $G(\mathbb{A}) \times \Theta$ by assigning each pair (λ, θ) the expectation of the random variable u^{θ} under the measure λ, i.e. $G(u)(\lambda, \theta) := \mathbb{E}_{\lambda}\big[u^{\theta}\big]$, which corresponds to the expectation of utility that choosing the randomized act λ will lead to, given θ is the true description of reality. Every *pure* act $a \in \mathbb{A}$ then can uniquely be identified with the *Dirac-measure* $\delta_a \in G(\mathbb{A})$, and we have $u(a, \theta) = G(u)(\delta_a, \theta)$

Table 1. Basic model (left) and running example with acts $\mathbb{A} = \{a_1, a_2, a_3\}$, states $\Theta = \{\theta_1, \ldots, \theta_4\}$ (right) and the credal set $\mathcal{M} := \{\pi : 0.3 \leqslant \pi(\{\theta_2\}) + \pi(\{\theta_3\}) \leqslant 0.7\}$ additionally considered in the Sects. 4.1 and 4.2.

$u(a_i, \theta_j)$	θ_1	$\cdots$	θ_m
a_1	$u(a_1, \theta_1)$	$\cdots$	$u(a_1, \theta_m)$
$\vdots$	$\vdots$	$\ldots$	$\vdots$
a_n	$u(a_n, \theta_1)$	$\cdots$	$u(a_n, \theta_m)$

$u(a_i, \theta_j)$	θ_1	θ_2	θ_3	θ_4
a_1	20	15	10	5
a_2	30	10	10	20
a_3	20	40	0	20

for all $(a, \theta) \in \mathbb{A} \times \Theta$. Again, for every $\lambda \in G(\mathbb{A})$ fixed, the extended utility function $G(u)$ is associated with a random variable $G(u)_\lambda$ on $(\Theta, 2^\Theta)$ by setting $G(u)_\lambda(\theta) := G(u)(\lambda, \theta)$ for all $\theta \in \Theta$. Finally, we refer to the triplet $(\mathbb{A}, \Theta, u)$ as the *(finite) decision problem* and to the triplet $(G(\mathbb{A}), \Theta, G(u))$ as the corresponding *randomized extension*.

Within this framework, our goal is to determine an *optimal* act (depending on the context, either randomized or pure). However, any appropriate definition of optimality depends on (what we assume about) the *mechanism generating the states of nature*. Here, traditional decision theory mainly covers two *extremes*: The mechanism follows a *known* probability measure π on $(\Theta, 2^\Theta)$ or it can be compared to a *game against an omniscient enemy*. In this cases optimality is almost unanimously defined by either *maximizing expected utility* with respect to π (also known as *Bayes-criterion*) or applying the *Maximin-criterion* (i.e. choosing an act that has maximal utility under the worst possible state of nature).

In contrast, defining optimality of acts becomes less obvious if the prior π is only *partially* known (case of *imprecise probabilities*) or there is uncertainty about the complete appropriateness of it (case of *uncertainty about precise probabilities*). The following sections are concerned with these two situations.

3 Handling Uncertain Precise Probabilistic Information: The Hodges and Lehmann-Criterion

Apart from the border cases of maximizing expected utility with respect to a precise prior π in the presence of perfect probabilistic information and the Maximin-criterion in complete absence of probabilistic information, classical decision theory tries to cope with decision making under uncertain probabilistic information, too: Anticipating ideas of *robust statistics*, Hodges and Lehmann proposed applying the Bayes-criterion only to such acts, whose worst possible utility does not fall below a certain amount of the Minimax utility (see [4]). Their idea is to utilize probabilistic information from previous experience while simultaneously distrusting the complete appropriateness of this information and restricting analysis to acts that are not too bad under the worst state. They also give the following alternative representation of their approach that has a different, intuitively more accessible, interpretation[1]: The decision maker is allowed to model his *degree of trust* in the prior by a parameter $\alpha \in [0,1]$. Specifically, if π is a probability measure on $(\Theta, 2^{\Theta})$, a randomized act $\lambda^* \in G(\mathbb{A})$ is said to be *Hodges and Lehmann*-optimal w.r.t. π and α (short: $\Phi_{\pi,\alpha}$-optimal), if $\Phi_{\pi,\alpha}(\lambda^*) \geqslant \Phi_{\pi,\alpha}(\lambda)$ for all $\lambda \in G(\mathbb{A})$, where

$$\Phi_{\pi,\alpha}(\lambda) := (1-\alpha) \cdot \min_{\theta} G(u)(\lambda, \theta) + \alpha \cdot \mathbb{E}_{\pi}\Big[G(u)_{\lambda}\Big] \tag{1}$$

Thus, the parameter α in (1) controls how the linear trade-off between expectation maximization w.r.t. π and applying the Maximin-criterion is actually made. The following Proposition 1 describes an algorithm for determining a randomized Hodges and Lehmann-optimal act for arbitrary pairs (π, α).[2]

Proposition 1. *Consider the linear programming problem*

$$(1-\alpha) \cdot (w_1 - w_2) + \alpha \cdot \sum_{i=1}^{n} \mathbb{E}_{\pi}(u_{a_i}) \cdot \lambda_i \longrightarrow \max_{(w_1, w_2, \lambda_1, \ldots, \lambda_n)} \tag{2}$$

with constraints $(w_1, w_2, \lambda_1, \ldots, \lambda_n) \geqslant 0$ *and*

- $\sum_{i=1}^{n} \lambda_i = 1$
- $w_1 - w_2 \leqslant \sum_{i=1}^{n} u_{ij} \cdot \lambda_i$ *for all* $j = 1, \ldots, m$.

Then the following holds:

(i) Every optimal solution $(w_1^*, w_2^*, \lambda_1^*, \ldots, \lambda_n^*)$ *to (2) induces a* $\Phi_{\pi,\alpha}$*-optimal randomized act* $\lambda^* \in G(\mathbb{A})$ *by setting* $\lambda^*(\{a_i\}) := \lambda_i^*$.
(ii) There always exists an $\Phi_{\pi,\alpha}$*-optimal randomized act.* □

By computing the dual linear program of the linear program given in Proposition 1, we receive the following Corollary. It can be interpreted as a method to construct priors that take the agent's *scepticism* about the prior probability π (expressed by the parameter α) into account.

Corollary 1. *Let* $\lambda^* \in G(\mathbb{A})$ *denote a* $\Phi_{\pi,\alpha}$*-optimal randomized act. Then, there exists a probability measure* $\mu_{\pi,\alpha}$ *on* $(\Theta, 2^{\Theta})$ *and a pure act* $a^* \in \mathbb{A}$ *such that*

$$\Phi_{\pi,\alpha}(\lambda^*) = \mathbb{E}_{\mu_{\pi,\alpha}}[u_{a^*}] \tag{3}$$

Proof. The dual of the optimization problem (2) is given by:

$$z_1 - z_2 \longrightarrow \min_{(z_1, z_2, \sigma_1, \ldots, \sigma_m)} \tag{4}$$

with constraints $(z_1, z_2, \sigma_1, \ldots, \sigma_m) \geqslant 0$ and

[1] A further mathematical characterization from the viewpoint of Gamma-Maximinity for certain imprecise probabilities is given in Footnote 3.

[2] The proofs of Propositions 1, 2 and 3 are straightforward and therefore left out.

- $\sum_{j=1}^{m} \sigma_j = 1 - \alpha$
- $z_1 - z_2 \geq \sum_{j=1}^{m} u_{ij} \cdot \sigma_j + \alpha \cdot \mathbb{E}_\pi(u_{a_i})$ for all $i = 1, \ldots, n$.

Let $(z_1^*, z_2^*, \sigma_1^*, \ldots, \sigma_m^*)$ denote an optimal solution to (4). Then the constraints guarantee that assigning $\mu_{\pi,\alpha}(\{\theta_j\}) := \alpha \cdot \pi(\{\theta_j\}) + \sigma_j^*$ for all $j = 1, \ldots, m$ induces a probability measure on $(\Theta, 2^\Theta)$ and that for all expectation maximal acts $a^* \in \mathbb{A}$ with respect to $\mu_{\pi,\alpha}$ it holds that $z_1^* - z_2^* = \mathbb{E}_{\mu_{\pi,\alpha}}[u_{a^*}]$. Further, by duality, we know that $z_1^* - z_2^*$ coincides with the optimal value of program (2) and, therefore, with $\Phi_{\pi,\alpha}(\lambda^*)$ where $\lambda^* \in G(\mathbb{A})$ denotes an Hodges and Lehmann-optimal randomized act. Thus, $\Phi_{\pi,\alpha}(\lambda^*) = \mathbb{E}_{\mu_{\pi,\alpha}}[u_{a^*}]$, as desired. □

Running Example (Table 1): Let π denote the prior on $(\Theta, 2^\Theta)$ induced by $(0.2, 0.7, 0.05, 0.05)$ and let our trust in π be expressed by $\alpha = 0.35$. Resolving the linear programming problem from Proposition 1 gives the optimal solution $(8, 0, 0.8, 0, 0.2)$. Thus, a $\Phi_{\pi,0.35}$-optimal randomized act $\lambda^* \in G(\mathbb{A})$ is induced by $(0.8, 0, 0.2)$. Next, we can use Corollary 1 to compute $\mu_{\pi,0.35}$. An optimal solution of problem (4) is given by the vector $(11.78, 0, 0, 0, 0.6385, 0.0115)$, and thus the measure $\mu_{\tau,0.35}$ is induced by the vector $(0.070, 0.245, 0.656, 0.029)$.

4 Handling Imprecise Probabilistic Information: The Gamma-Maximin View

We now turn to decision criteria taking into account the uncertainty in the prior information in a more direct way: For modeling prior knowledge, instead of one classical probability, we consider polyhedral sets of probability measures that are a common tool in different theories of imprecise probabilities, like e.g. *linear partial information* ([7]), *credal sets* ([8]), *lower previsions* ([16]) or *interval probability* ([17]) as well as in robust statistics, like e.g. *ε-contamination models* (see [3, p. 12]). Particularly, we assume probabilistic information is expressed by a polyhedrical set $\mathcal{M}$ of probability measures on $(\Theta, 2^\Theta)$ of the form

$$\mathcal{M} := \{\pi | \ \underline{b}_s \leqslant \mathbb{E}_\pi(f_s) \leqslant \overline{b}_s \ \forall s = 1, ..., r\} \tag{5}$$

where, for all $s = 1, ..., r$, we have $(\underline{b}_s, \overline{b}_s) \in \mathbb{R}^2$ such that $\underline{b}_s \leqslant \overline{b}_s$ and $f_s : \Theta \to \mathbb{R}$. Specifically, the available information is assumed to be describable by lower and upper bounds for the expected values of a finite number of random variables on the space of states. Clearly, if uncertainty is described by a set of probability measures, defining meaningful criteria for decision making strongly depend on the agent's *attitude towards ambiguity*, i.e. towards the non-stochastic uncertainty between the measures contained in $\mathcal{M}$. Accordingly, many competing criteria exist (see [12] for a survey or [2,8,16] for original sources). In the following sections, we present linear programming based results for a selection of such criteria, namely *Walley's maximality* and the *Gamma-Maximin* criterion. For the latter, we also investigate some connections to least favorable priors.

4.1 Checking Maximality of Pure Acts

The idea behind maximality of an act $a^* \in \mathbb{A}$ is quite simple: One repeatedly compares an act a^* pairwise to all other acts and checks whether there exists an element of the set $\mathcal{M}$ with respect to which u_{a^*} dominates the corresponding other act in expectation. Formally, an act $a^* \in \mathbb{A}$ is said to be $\mathcal{M}$-*maximal*, if

$$\forall\, a \in \mathbb{A} \;\; \exists\, \pi_a \in \mathcal{M}: \quad \mathbb{E}_{\pi_a}(u_{a^*}) \geqslant \mathbb{E}_{\pi_a}(u_a) \tag{6}$$

Naturally, the above definition extends to randomized acts. However, when also considering randomized acts, the criterion of $\mathcal{M}$-Maximality coincides (see [16, p. 163]) with another well-investigated criterion known from IP decision theory contributed to *Levi*: E-admissibility. For a detailed discussion of connections between the two criteria see [11]. An algorithm for determining the set of all randomized E-admissible acts has been introduced in [13]. However, for finite $\mathbb{A}$, being $\mathcal{M}$-Maximal is a strictly weaker condition and, therefore, needs to be checked separately from E-admissibility. Other approaches for doing so have already been proposed in [6]. Proposition 2 describes an algorithm for checking $\mathcal{M}$-Maximality of a pure act $a_z \in \mathbb{A}$ by solving one single linear program.

Proposition 2. *Let* $(\mathbb{A}, \Theta, u)$ *denote a finite decision problem and let* $\mathcal{M}$ *be of the form (5). Further, let* $a_z \in \mathbb{A}$ *be any act. Consider the linear program*

$$\sum_{i=1}^{n}\Big(\sum_{j=1}^{m} \gamma_{ij}\Big) \longrightarrow \max_{(\gamma_{11},\dots,\gamma_{nm})} \tag{7}$$

with constraints $(\gamma_{11}, \dots, \gamma_{nm}) \geqslant 0$ *and*

- $\sum_{j=1}^{m} \gamma_{ij} \leqslant 1 \quad$ *for all* $i = 1, \dots, n$
- $\underline{b}_s \leqslant \sum_{j=1}^{m} f_s(\theta_j) \cdot \gamma_{ij} \leqslant \overline{b}_s \quad$ *for all* $s = 1, ..., r$, $i = 1, \dots, n$
- $\sum_{j=1}^{m} (u_{ij} - u_{zj}) \cdot \gamma_{ij} \leqslant 0 \quad$ *for all* $i = 1, \dots, n$.

Then $a_z \in \mathbb{A}$ *is* $\mathcal{M}$*-Maximal iff the optimal outcome of (7) equals* n. □

If $(\gamma_{11}^*, \dots, \gamma_{nm}^*)$ is an optimal solution to problem (7) yielding an value of n, we can construct $\pi_{a_i} \in \mathcal{M}$ for which act a_z dominates act a_i in expectation by setting $\pi_{a_i}(\{\theta_j\}) := \gamma_{ij}$. The problem possesses $n(3 + r)$ constraints and nm decision variables. Determining the set of all maximal acts requires to solve n such linear programs. Compared to this, the algorithm based on pairwise comparisons of acts proposed in [6] here translates to solving $n^2 - n$ linear programs with m decision variables, however, with only $r + 2$ constraints.

Running Example (Table 1): Resolving the linear programming problem from Proposition 2 for every act a_1, a_2 and a_3 separately gives optimal value 3 for each of them. Thus, all available acts are $\mathcal{M}$-Maximal.

4.2 Gamma-Maximin and Least Favorable Priors

In this section, we first present a linear program for identifying a *least favorable prior distribution* from the credal set $\mathcal{M}$ under consideration. Afterwards, we investigate the dual of this linear program and, in this way, provide a connection between pure acts $a \in \mathbb{A}$ that maximize expected utility with respect to a least favorable prior and randomized acts $\lambda \in G(\mathbb{A})$ that are optimal with respect to the Gamma-Maximin criterion.

Before we proceed, some additional notation is needed: For a credal element $\pi \in \mathcal{M}$, let $B(\pi)$ denote the maximal expectation with respect to π that an act from $\mathbb{A}$ can yield (that is $B(\pi) = \mathbb{E}_\pi(u_{a^*})$, where $a^* \in \mathbb{A}$ maximizes expected utility with respect to π). The set of all acts $a \in \mathbb{A}$ that maximize expected utility with respect to π is denoted by $\mathbb{A}_\pi$. Further, we call a credal element $\pi^- \in \mathcal{M}$ a *least favorable prior (lfp)* from $\mathcal{M}$ iff $B(\pi^-) \leqslant B(\pi)$ holds for all $\pi \in \mathcal{M}$. Specifically, π^- is a lfp, if it yields the minimal best possible expected utility under all concurring elements on the credal set. Proposition 3 describes a linear program for determining a lfp from $\mathcal{M}$.

Proposition 3. *Let* $(\mathbb{A}, \Theta, u)$ *denote a decision problem and let* $\mathcal{M}$ *be of the form (5). Consider the linear program*

$$w_1 - w_2 \longrightarrow \min_{(w_1, w_2, \pi_1, \ldots, \pi_m)} \tag{8}$$

with constraints $(w_1, w_2, \pi_1, \ldots, \pi_m) \geqslant 0$ *and*

- $\sum_{j=1}^{m} \pi_j = 1$
- $\underline{b}_s \leqslant \sum_{j=1}^{m} f_s(\theta_j) \cdot \pi_j \leqslant \overline{b}_s$ *for all* $s = 1, ..., r$
- $w_1 - w_2 \geqslant \sum_{j=1}^{m} u_{ij} \cdot \pi_j$ *for all* $i = 1, \ldots n$.

Then the following holds:

(i) Every optimal solution $(w_1^*, \ldots, \pi_m^*)$ *to (8) induces a least favorable prior* $\pi^- \in \mathcal{M}$ *by setting* $\pi^-(\{\theta_j\}) := \pi_j^*$.
(ii) There always exists a least favorable prior. □

A lfp can be understood as a kind of "pignistic" probability, representing the decision problem under complex uncertainty in a way that is specific to the problem and the criterion under consideration, but in return gives the exact criterion value. This contrasts lfps from pignistic probabilities in Smets' spirit, who argued that a decision problem under complex uncertainty could be approached by distinguishing between a *credal level*, where the uncertain beliefs are to be expressed with all their ambiguity and scarceness by an imprecise probability (belief function in Smets' context), and a *decision level*, where eventually the imprecise probability is condensed into a traditional probability on which expected utility theory could be applied (see, e.g., [9,10], as well as, e.g., [1] for geometric techniques to represent belief functions by a single precise probability).

We now show some connections between least favorable priors and randomized Gamma-Maximin acts w.r.t. $\mathcal{M}$ ($\mathcal{M}$-Maximin). Recalling its definition, a randomized act $\lambda^* \in G(\mathbb{A})$ is said to be *$\mathcal{M}$-Maximin optimal* iff for all $\lambda \in G(\mathbb{A})$:

$$\underline{\mathbb{E}}_{\mathcal{M}}\big[G(u)_{\lambda^*}\big] \geqslant \underline{\mathbb{E}}_{\mathcal{M}}\big[G(u)_{\lambda}\big] \tag{9}$$

where $\underline{\mathbb{E}}_{\mathcal{M}}(X) := \min_{\pi \in \mathcal{M}} \mathbb{E}_{\pi}(X)$ for random variables $X : (\Theta, 2^{\Theta}) \to \mathbb{R}$.[3] It turns out that the linear program from Proposition 3 is *dual* to the one for determining a randomized $\mathcal{M}$-Maximin act described in [13, Sect. 3.2]. Together with *complementary slackness* (see, e.g., [15, Sect. 5.5]) from linear optimization theory, this allows to derive connections between lfps and the Gamma-Maximin.

Proposition 4. *Let $(\mathbb{A}, \Theta, u)$ denote a finite decision problem and let $\mathcal{M}$ be of the form (5). Then the following holds:*

(i) If π^- is a lfp from $\mathcal{M}$, then for all optimal randomized $\mathcal{M}$-Maximin acts $\lambda^* \in G(\mathbb{A})$ we have $\lambda^*(\{a\}) = 0$ for all $a \in \mathbb{A} \backslash \mathbb{A}_{\pi^-}$.
(ii) Let π^- denote a lfp from $\mathcal{M}$ and let $\lambda^* \in G(\mathbb{A})$ denote a randomized $\mathcal{M}$-Maximin act. Then for all $a \in \mathbb{A}_{\pi^-}$ we have

$$\mathbb{E}_{\pi^-}\big[u_a\big] = \underline{\mathbb{E}}_{\mathcal{M}}\big[G(u)_{\lambda^*}\big]$$

Proof. The dual programming problem of problem (8) is given by:

$$z_1 - z_2 + \sum_{s=1}^{r} (\underline{b}_s x_s - \overline{b}_s y_s) \longrightarrow \max_{(z_1, z_2, x_1, \dots, x_r, y_1, \dots, y_r, \lambda_1, \dots, \lambda_n)} \tag{10}$$

with constraints $(z_1, z_2, x_1, \dots, x_r, y_1, \dots, y_r, \lambda_1, \dots, \lambda_n) \geqslant 0$ and

- $\sum_{i=1}^{n} \lambda_i = 1$
- $z_1 - z_2 + \sum_{s=1}^{r} f_s(\theta_j)(x_s - y_s) \leq \sum_{i=1}^{n} u_{ij} \cdot \lambda_i$ for all $j = 1, \dots, m$.

The resulting linear program (10) is exactly the one for determining a randomized act $\lambda^* \in G(\mathbb{A})$ which is optimal with respect to the $\mathcal{M}$-Maximin criterion as proposed and proven in [13, Sect. 3.2]. We now can use standard results on duality and complementary slackness (see, e.g., [15, Chap. 5]) to proof the proposition:

[3] For the special case of an ε-contamination model (a.k.a. *linear-vacuous model*) of the form $\mathcal{M}_{(\pi_0, \varepsilon)} := \{(1-\varepsilon)\pi_0 + \varepsilon\pi : \pi \in \mathcal{P}(\Theta)\}$, where $\mathcal{P}(\Theta)$ denotes the set of all probability measures on $(\Theta, 2^{\Theta})$, $\varepsilon > 0$ is a fixed contamination parameter and $\pi_0 \in \mathcal{P}(\Theta)$ is the central distribution, Gamma-Maximin is mathematically closely related to the Hodges and Lehmann-criterion: For fixed $X : (\Theta, 2^{\Theta}) \to \mathbb{R}$ we have $\underline{\mathbb{E}}_{\mathcal{M}_{(\pi_0, \varepsilon)}}(X) = \min_{\pi \in \mathcal{P}(\Theta)}((1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon\mathbb{E}_{\pi}(X)) = (1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon \min_{\pi \in \mathcal{P}(\Theta)} \mathbb{E}_{\pi}(X) = (1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon \min_{\theta \in \Theta} X(\theta)$. Thus, maximizing the lower expectation w.r.t. the ε-contamination model is equivalent to maximizing the Hodges and Lehmann-criterion with trust parameter $(1-\varepsilon)$ and prior π_0.

Part (i): Let $\pi^- \in \mathcal{M}$ denote a lfp and let $a_z \in \mathbb{A} \backslash \mathbb{A}_{\pi^-}$. Then

$$(\max\{B(\pi^-), 0\}, -\min\{B(\pi^-), 0\}, \pi^-(\{\theta_1\}), \ldots, \pi^-(\{\theta_m\})) \tag{11}$$

defines an optimal solution to (8) for which it holds that $B(\pi^-) > \mathbb{E}_{\pi^-}(u_{a_z})$. Thus, there exists an optimal solution to (8), for which the constraint $w_1 - w_2 \geq \sum_{j=1}^{m} u_{zj} \cdot \pi_j$ holds strictly and, therefore, the corresponding slack variable is strictly greater 0. Hence, by complementary slackness, the corresponding variable in the dual problem (10), that is λ_z, equals 0 for every optimal solution of problem (10). Finally, note that $\{\lambda_z^* : \lambda_z^* \text{ appears in optimal solution}\} = \{\lambda^*(\{a_z\}) : \lambda^* \in G(\mathbb{A})\ \mathcal{M}\text{-Maximin optimal}\}$, since, as (implicitly) shown in [13, Sect. 3.2], every $\mathcal{M}$-Maximin optimal $\lambda^* \in G(\mathbb{A})$ induces an optimal solution to (10), namely

$$(z_1^*, z_2^*, x_1, \ldots, x_r^*, y_1^*, \ldots, y_r^*, \lambda^*(\{a_1\}), \ldots, \lambda^*(\{a_n\})) \tag{12}$$

where $(z_1^*, z_2^*, x_1, \ldots, x_r^*, y_1^*, \ldots, y_r^*)$ denotes an optimal solution to a reduced version of problem (10) with $(\lambda_1, \ldots, \lambda_n) := (\lambda^*(\{a_1\}), \ldots, \lambda^*(\{a_n\}))$ fixed.

Part (ii): Let $\pi^- \in \mathcal{M}$ denote an lfp and $\lambda^* \in G(\mathbb{A})$ denote an $\mathcal{M}$-Maximin act. Use (11) and (12) to construct optimal solutions to (8) and (10). As the optimal value of (8) equals $B(\pi^-)$ and the optimal value of (10) equals $\underline{\mathbb{E}}_{\mathcal{M}}\big[G(u)_{\lambda^*}\big]$, the result follows by the duality theorem. □

As an immediate consequence of Proposition 4 (i), we can specify a condition under which randomization cannot improve utility, if optimality is defined in terms of the Gamma-Maximin criterion. Specifically, we have the following corollary.

Corollary 2. *If there exists a lfp π^- from $\mathcal{M}$ such that $\mathbb{A}_{\pi^-} = \{a_z\}$ for some $z \in \{1, \ldots, n\}$, then $\delta_{a_z} \in G(\mathbb{A})$ is the unique randomized $\mathcal{M}$-Maximin act. Specifically, considering randomized acts is unnecessary in such situations.* □

Running Example (Table 1): Algorithm 8 leads to the optimal solution vector $(13, 0, 0, 0, 0.7, 0.3)$. Thus, a lfp π^- from $\mathcal{M}$ is induced by $(0, 0.7, 0.3, 0)$. Simple computation gives $\mathbb{A}_{\pi^-} = \{a_2\}$. Hence, according to Corollary 2, a_2 is the unique $\mathcal{M}$-Maximin act (even compared to randomized acts) with utility 13.

5 Summary and Concluding Remarks

We presented linear programming based approaches for determining optimal randomized acts and investigated what can be learned by dualizing these. Future research includes the following issues: If $\mathcal{M}$ is non-degenerated, i.e. $\pi(\{\theta\}) > 0$ for all $(\pi, \theta) \in \mathcal{M} \times \Theta$, the same holds for every lfp π^-. Since every π^- induces an optimal solution to (8), complementary slackness implies that all constraints of problem (10) are binding for every optimal solution. This gives a system of linear *equations* that have to be satisfied by every randomized $\mathcal{M}$-Maximin act. A natural question is: Under which conditions is this system sufficient to identify an optimal act without solving an optimization problem at all? A further interesting point is that algorithm (7) for checking maximality of an act a_z takes into account all other acts a_i in one linear program simultaneously. This could be used to modify the algorithm for finding maximal acts that are not too far from being E-admissible in the sense that the involved probabilities π_{a_i} that establish maximality of a_z differ not too much w.r.t. the L_1-norm which can be guaranteed by imposing further *linear* constraints.

Acknowledgement. The authors would like to thank the three anonymous referees for their helpful comments and their support.

References

1. Cuzzolin, F.: Two new Bayesian approximations of belief functions based on convex geometry. IEEE T. Syst. Man. Cy. B **37**, 993–1008 (2007)
2. Gilboa, I., Schmeidler, D.: Maxmin expected utility with non-unique prior. J. Math. Econ. **18**, 141–153 (1989)
3. Huber, P.: Robust Statistics. Wiley, New York (1981)
4. Hodges, J., Lehmann, E.: The use of previous experience in reaching statistical decisions. Ann. Math. Stat. **23**, 396–407 (1952)
5. Hable, R., Troffaes, M.: Computation. In: Augustin, T., Coolen, F., de Cooman, G., Troffaes, M. (eds.) Introduction to Imprecise Probabilities, pp. 329–337. Wiley, Chichester (2014)
6. Kikuti, D., Cozman, F., Filho, R.: Sequential decision making with partially ordered preferences. Artif. Intel. **175**, 1346–1365 (2011)
7. Kofler, E., Menges, G.: Entscheiden bei unvollständiger Information. Springer, Berlin (1976)
8. Levi, I.: The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance. MIT Press, Cambridge (1983)
9. Smets, P.: Decision making in the TBM: the necessity of the pignistic transformation. Int. J. Approx. Reason. **38**, 133–147 (2005)
10. Smets, P.: Decision making in a context where uncertainty is represented by belief functions. In: Srivastava, R., Mock, T. (eds.) Belief Functions in Business Decisions, pp. 17–61. Physica, Heidelberg (2002)
11. Schervish, M., Seidenfeld, T., Kadane, J., Levi, I.: Extensions of expected utility theory and some limitations of pairwise comparisons. In: Bernard, J.-M., Seidenfeld, T., Zaffalon, M. (eds.) Proceedings of ISIPTA 2003, pp. 496–510. Carleton Scientific, Waterloo (2003)
12. Troffaes, M.: Decision making under uncertainty using imprecise probabilities. Int. J. Approx. Reason. **45**, 17–29 (2007)
13. Utkin, L., Augustin, T.: Powerful algorithms for decision making under partial prior information and general ambiguity attitudes. In: Cozman, F., Nau, R., Seidenfeld, T. (eds.) Proceedings of ISIPTA 2005, pp. 349–358 (2005)
14. Utkin, L., Kozine, I.: Different faces of the natural extension. In: de Cooman, G., Fine, T., Seidenfeld, T. (eds.) Proceedings of ISIPTA 2001, pp. 316–323 (2001)
15. Vanderbei, R.: Linear Programming: Foundations and Extensions. Springer, New York (2014)
16. Walley, P.: Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London (1991)

17. Weichselberger, K.: Elementare Grundbegriffe einer allgemeineren Wahrscheinlichkeitsrechnung I: Intervallwahrscheinlichkeit als umfassendes Konzept. Physica, Heidelberg (2001)
18. Weichselberger, K.: Interval probability on finite sample spaces. In: Rieder, H. (ed.) Robust Statistics, Data Analysis, and Computer Intensive Methods, pp. 391–409. Springer, New York (1996)