Decision Theory Meets Linear Optimization (Beyond Computation)

C. Jansen T. Augustin G. Schollmeyer

Department of Statistics, LMU Munich

ECSQARU 2017, 13th of July, Lugano

LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Scope of the talk

- linear programming theory has been shown to be a powerful tool for (imprecise) decision theory regarding both
  - efficient computation of optimal acts w.r.t. complex criteria (cf., e.g., Kikuti et al. (2012) or Utkin and Augustin (2005))
  - providing theoretical insights on properties of optimal acts (cf., e.g., Weichselberger (1996))
- our paper presents some new results concerning both regards including
  - linear programs for Hodges and Lehmann and Walley's maximality
  - connection between least favorable priors and Gamma-Maximin

Setup and notation

We consider the standard model of (finite) cardinal decision theory:

- $\mathbb{A} = \{a_1, \ldots, a_n\}$: set of acts
- $\Theta = \{\theta_1, \ldots, \theta_m\}$: set of states of the world
- $u : \mathbb{A} \times \Theta \to \mathbb{R}$: utility function, where $u_{ij} := u(a_i, \theta_j)$ is the utility of choosing act a_i given θ_j is the true state of the world

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

- for every $a \in \mathbb{A}$, define $u_a : \Theta \to \mathbb{R}$ by $u_a(\theta) := u(a, \theta)$ for all $\theta \in \Theta$
- for every $\theta \in \Theta$, define $u^\theta : \mathbb{A} \to \mathbb{R}$ by $u^\theta(a) := u(a, \theta)$ for all $a \in \mathbb{A}$

Setup and notation, continued

Depending on the context, we also allow for randomized acts:

- call every probability measure λ on $(\mathbb{A}, 2^{\mathbb{A}})$ a randomized act and denote by $G(\mathbb{A})$ the set of all randomized acts
- choosing λ is interpreted as leaving the final decision to a random experiment which yields act a_i with probability $\lambda(\{a_i\})$
- evaluate choosing λ given θ by $G(u)(\lambda, \theta) := \mathbb{E}_\lambda\left[u^\theta\right]$
- for $\lambda \in G(\mathbb{A})$, define $G(u)_\lambda : \Theta \to \mathbb{R}$ by $G(u)_\lambda(\theta) := G(u)(\lambda, \theta)$
- identify $a \in \mathbb{A}$ with $\delta_a \in G(\mathbb{A})$ and observe $u(a, \theta) = G(u)(\delta_a, \theta)$

Randomization: A toy example

- Consider a game between two players: Pinky (rows) and Brain (columns)
- Pinky chooses moves $P = \{p_1, p_2\}$, Brain reacts by moves $B = \{b_1, b_2\}$
- Pinky's utility $u_p : P \times B \to \mathbb{R}$ is given by the below table
- Brain's utility $u_b : B \times P \to \mathbb{R}$ is given by $u_b(b, p) := -u_p(p, b)$

$u_p(\cdot)$	b_1	b_2	Pinky's reward
p_1	10	20	10
p_2	30	5	5

- Pinky tosses a (fair) coin, i.e. chooses randomized act $\lambda \approx \begin{pmatrix} p_1 & p_2 \\ 0.5 & 0.5 \end{pmatrix}$.
- He receives reward of $\min_b G(u_p)(\lambda, b) = 12.5$.

Two ways of incorporating imperfect prior knowledge

Considered here: Decision problems with prior information on the states Θ.

If prior information is precisely given by an (undoubted) probability on the state space, acts are most commonly ranked with respect to their expected utility values.

Otherwise (of interest here), we distinguish two different cases:

(1) Uncertainty about precise probabilities: There is a precise prior probability π on $(\Theta, 2^\Theta)$ available, however, there is some doubt about its full appropriateness.

Example: Prior available for an experiment; slight modification of the setup

(2) Imprecise probabilities: A prior probability measure π on the state space Θ cannot be fully specified. Instead a credal set $\mathcal{M}$ of prior probabilities is compatible with the available information

Example: Event E_1 is at least as likely as E_2, i.e. $\mathcal{M} = \{\pi | \pi(E_1) \geq \pi(E_2)\}$

(1) Uncertainty about precise priors: Hodges & Lehmann

One common way to deal with situation (1) is the decision criterion of Hodges & Lehmann, which linearly trades of between maximin and expected utility.

Hodges & Lehmann optimality

Let π denote some prior on $(\Theta, 2^\Theta)$ and let $\alpha \in [0,1]$ express the agent's trust in its appropriateness. The function $\Phi_{\pi,\alpha} : G(\mathbb{A}) \to \mathbb{R}$ defined by

$$\Phi_{\pi,\alpha}(\lambda) = (1-\alpha) \cdot \underbrace{\min_{\theta} G(u)(\lambda, \theta)}_{\textit{Maximin utility}} \; + \; \alpha \cdot \underbrace{\mathbb{E}_\pi \left[G(u)_\lambda \right]}_{\textit{Expected utility}}$$

is called Hodges & Lehmann-criterion w.r.t. (π, α). Any randomized act $\lambda^* \in G(\mathbb{A})$ maximizing the criterion is then called $\Phi_{\pi,\alpha}$-optimal.

Natural question: How to determine/compute $\Phi_{\pi,\alpha}$-optimal acts?

Determining optimal acts under (1)

Optimal randomized acts with respect to the criterion of Hodges and Lehmann can be obtained by the following linear programming problem:

Hodges and Lehmann acts

Consider the linear programming problem

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

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

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

Then every optimal solution $(w_1^*, w_2^*, \lambda_1^*, \ldots, \lambda_n^*)$ induces a $\Phi_{\pi,\alpha}$-optimal randomized act $\lambda^* \in G(\mathbb{A})$ by setting $\lambda^*(\{a_i\}) := \lambda_i^*$.

(2) Imprecise probabilistic information

We assume probabilistic information is expressed by a polyhedrical credal set $\mathcal{M}$ of probability measures on $(\Theta, 2^\Theta)$ of the form

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

where, for all $s = 1, ..., r$, we have

- $f_s : \Theta \to \mathbb{R}$ is a random variables on Θ and
- $(\underline{b}_s, \overline{b}_s) \in \mathbb{R}^2$ with $\underline{b}_s \leqslant \overline{b}_s$ are lower and upper bounds for their expectation.

Least favorable priors

In the following, one additional concept will be needed:

- for $\pi \in \mathcal{M}$, let $B(\pi) := \sup\{\mathbb{E}_\pi(u_a) : a \in \mathbb{A}\}$ and $\mathbb{A}_\pi := \text{argmax}_{a \in \mathbb{A}} \mathbb{E}_\pi(u_a)$
- call $\pi^- \in \mathcal{M}$ a least favorable prior (lfp) if $B(\pi^-) \leq B(\pi)$ for all $\pi \in \mathcal{M}$

Computing least favorable priors

The following proposition describes an easy linear program for determining a least favorable prior distributions from a given credal set.

Least favorable priors

Let $(\mathbb{A}, \Theta, u)$ and $\mathcal{M}$ be as before. Consider the linear program

$$w_1 - w_2 \longrightarrow \min_{(w_1, w_2, \pi_1, \dots, \pi_m)}$$

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

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

Then every optimal solution $(w_1^*, \dots, \pi_m^*)$ induces a least favorable prior $\pi^- \in \mathcal{M}$ by setting $\pi^-(\{\theta_j\}) := \pi_j^*$.

Decision making under (2)

If uncertainty is characterized by a credal set $\mathcal{M}$, many different approaches for decision making exist. We focus on three of these, namely

Walley's maximality: An act $a^* \in \mathbb{A}$ is said to be $\mathcal{M}$-maximal, if

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

E-admissibility: An act $a^* \in \mathbb{A}$ is said to be $\mathcal{M}$-admissible, if

$$\exists\ \pi \in \mathcal{M}\ \ \forall\ a \in \mathbb{A}: \quad \mathbb{E}_{\pi}(u_{a^*}) \geqslant \mathbb{E}_{\pi}(u_a)$$

Gamma-Maximin: An randomized act $\lambda^* \in G(\mathbb{A})$ is said to be $\mathcal{M}$-Maximin optimal iff for all $\lambda \in G(\mathbb{A})$:

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

where $\underline{\mathbb{E}}_{\mathcal{M}}(X) := \min_{\pi \in \mathcal{M}} \mathbb{E}_{\pi}(X)$ for random variables $X : \Theta \to \mathbb{R}$.

A linear program for maximality

The set of maximal (non-randomized) acts can be determined by running the following linear program for every act $a \in \mathbb{A}$ under consideration:

Checking maximality of non-randomized acts

Let $a_z \in \mathbb{A}$ be any act. Consider the linear program

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

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 equals n.

A slight modification: c-constraint maximality

Checking c-constraint maximality of pure acts

Let $a_z \in \mathbb{A}$ be any act and let $c \in [0,1]$. Consider the linear program

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

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 $i = 1,\dots,n$
- $\frac{1}{2}\sum_{j=1}^{m}|\gamma_{i_1 j} - \gamma_{i_2 j}| \leqslant c \quad$ for all $i_1, i_2 \in \{1,\dots,n\}$

Then $a_z \in \mathbb{A}$ is $c\mathcal{M}$-Maximal iff the optimal outcome equals n.

Crossing the border between (1) and (2)

For the special case of an ε-contamination model of the form

$$\mathcal{M}_{(\pi_0,\varepsilon)} := \{(1-\varepsilon)\pi_0 + \varepsilon\pi : \pi \in \mathcal{P}(\Theta)\}$$

where

- $\varepsilon > 0$ 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 & Lehmann:

$$\begin{aligned} \underline{\mathbb{E}}_{\mathcal{M}_{(\pi_0,\varepsilon)}}(X) &= \min_{\pi \in \mathcal{P}(\Theta)} \left((1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon \mathbb{E}_{\pi}(X)\right) \\ &= (1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon \min_{\pi \in \mathcal{P}(\Theta)} \mathbb{E}_{\pi}(X) \\ &= (1-\varepsilon)\mathbb{E}_{\pi_0}(X) + \varepsilon \min_{\theta \in \Theta} X(\theta) \end{aligned}$$

$\mathcal{M}_{(\pi_0,\varepsilon)}$-Maximin $\approx$ Hodges & Lehmann w.r.t. $(1-\varepsilon)$ and π_0.

A result beyond computation

Gamma-Maximin and lfps

Let $(\mathbb{A}, \Theta, u)$ and $\mathcal{M}$ be as before. 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 $a \in \mathbb{A} \setminus \mathbb{A}_{\pi^-}$.

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

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

Corollary

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.

Summary and outlook

We investigated

- linear programming approaches for determining optimal randomized acts
- what can be learned by dualizing our programs

Future research includes:

- consider $\mathcal{M}$ is non-degenerated, i.e. $\pi(\{\theta\}) > 0$ for all $(\pi, \theta) \in \mathcal{M} \times \Theta$
- then every lfp π^- from $\mathcal{M}$ is non-degenerated as well
- by complementary slackness property, the constraints in the linear program for determining for Gamma-Maximin acts are binding: when sufficient?