# Decision theory meets linear optimization beyond computation


**Christoph Jansen & Thomas Augustin**, Department of Statistics, University of Munich

**Contact:** christoph.jansen@stat.uni-muenchen.de augustin@stat.uni-muenchen.de


LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

## 1. Guideline through the poster

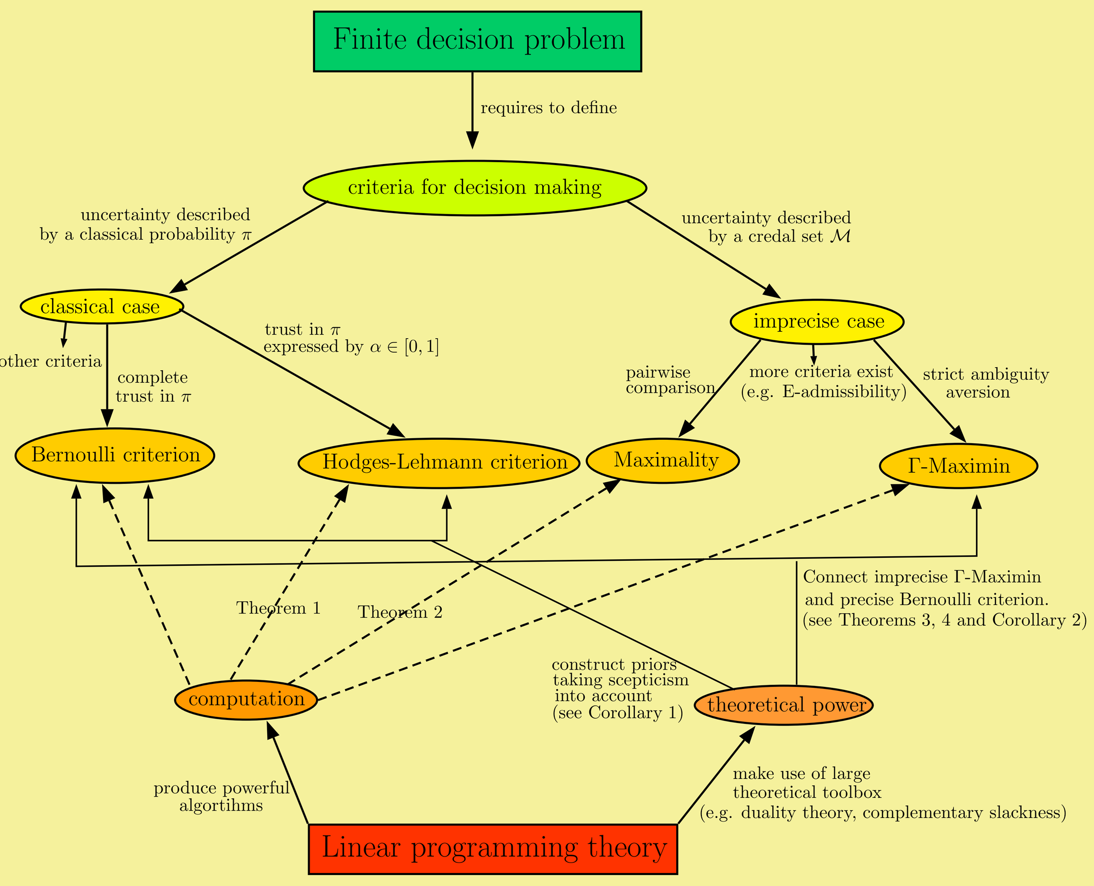


## 2. Background: Some finite decision theory

We consider the standard model of *finite* decision theory:

- $\mathbb{A} = \{a_1, \dots, a_n\}$, $n \in \mathbb{N}$ (*set of actions*)
- $\Theta = \{\theta_1, \dots, \theta_m\}$, $m \in \mathbb{N}$ (*states of nature*)
- $u : \mathbb{A} \times \Theta \to \mathbb{R}$ (*cardinal utility function*)

Naturally, the utility function associates

- every action $a \in \mathbb{A}$ with a *gamble* $u_a$ on $(\Theta, 2^\Theta)$:

$$u_a : \Theta \to \mathbb{R} \quad , \quad \theta \mapsto u(a, \theta) \tag{1}$$

- every state $\theta \in \mathbb{A}$ with a variable $u^\theta$ on $(\mathbb{A}, 2^\mathbb{A})$:

$$u^\theta : \mathbb{A} \to \mathbb{R} \quad , \quad a \mapsto u(a, \theta) \tag{2}$$

With $u_{ij} := u(a_i, \theta_j)$, we represent the model by its *utility matrix*:

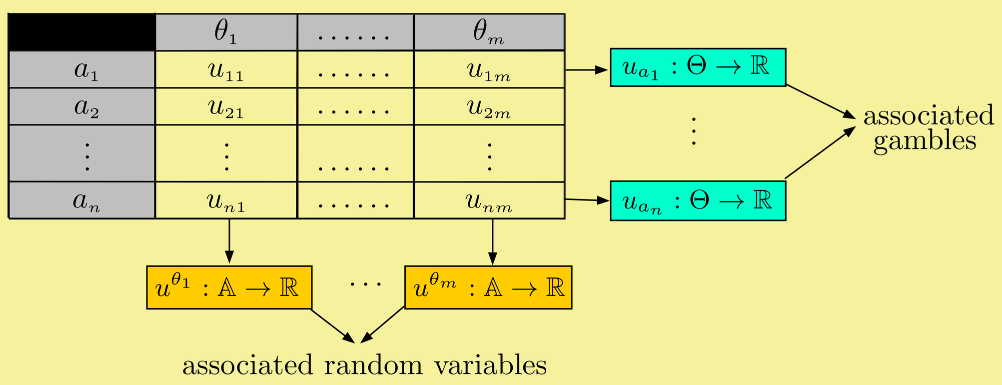


Depending on the context, we also allow for choosing *randomized actions*, i.e. classical probability measures on $(\mathbb{A}, 2^\mathbb{A})$. We denote the set of all randomized actions by $G(\mathbb{A})$.

The utility function $u$ is then extended to a utility function $G(u)$ on $G(\mathbb{A}) \times \Theta$ by assigning each pair $(\lambda, \theta)$ the expectation of the random variable $u^\theta$ under the measure $\lambda$, i.e. $\mathbb{E}_\lambda[u^\theta]$.

Every *pure* action $a \in \mathbb{A}$ then can uniquely be identified with the *Dirac-measure* $\delta_a \in G(\mathbb{A})$ and we have $u(a, \theta) = G(u)(\delta_a, \theta)$ for all $(a, \theta) \in \mathbb{A} \times \Theta$. Further, also (1) can easily be extended to randomized actions by defining, for every $\lambda \in G(\mathbb{A})$ fixed, $G(u)_\lambda(\theta) := G(u)(\lambda, \theta)$ for all $\theta \in \Theta$.

## 3. A criterion from classical decision theory

Apart from the border cases of *maximizing expected utility* w.r.t. a precise prior and the *maximin-criterion*, classical decision theory tries to cope with decision making under vague information, too: The criterion of *Hodges and Lehmann* allows the decision maker to model his *degree of trust* in the prior by a parameter $\alpha \in [0, 1]$.

Specifically, if $\pi$ is a probability measure on $(\Theta, 2^\Theta)$, a randomized action $\lambda^* \in G(\mathbb{A})$ is said to be *Hodges-Lehmann*-optimal w.r.t. $\pi$ and $\alpha$ (short: $\Phi_{\pi,\alpha}$-optimal), if $\Phi_{\pi,\alpha}(\lambda^*) \geqslant \Phi_{\pi,\alpha}(\lambda)$ for all $\lambda \in G(\mathbb{A})$, where

$$\Phi_{\pi,\alpha}(\lambda) := (1 - \alpha) \cdot \min_\theta G(u)(\lambda, \theta) + \alpha \cdot \mathbb{E}_\pi \Big[G(u)_\lambda\Big] \tag{3}$$

Theorem 1 describes an algorithm determining a randomized Hodges-Lehmann-actions for arbitrary pairs $(\pi, \alpha)$.

**Theorem 1.** *Consider the linear programming problem*

$$(1 - \alpha) \cdot (w_1 - w_2) + \alpha \cdot \sum_{i=1}^{n} \mathbb{E}_\pi(u_{a_i}) \cdot p_i \longrightarrow \max_{(w_1, w_2, p_1, \dots, p_n)} \tag{4}$$

*with constraints* $(w_1, w_2, p_1, \dots, p_n) \geqslant 0$ *and*

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

*Then the following holds:*

*i) Every optimal solution* $(w_1^*, w_2^*, p_1^*, \dots, p_n^*)$ *to (4) induces a* $\Phi_{\pi,\alpha}$*-optimal randomized action* $\lambda^* \in G(\mathbb{A})$ *by setting* $\lambda^*(\{a_i\}) := p_i^*$.

*ii) There always exists an* $\Phi_{\pi,\alpha}$*-optimal randomized action.*

By applying *duality* theory, we receive the following Corollary. Its proof can be interpreted as a method to construct priors that take the actor's *scepticism about* $\pi$ (expressed by $\alpha$) into account.

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

$$\Phi_{\pi,\alpha}(\lambda^*) = \mathbb{E}_{\mu_{\pi,\alpha}}[u_{a^*}] \tag{5}$$

## 4. Linear partial information

*Kofler* and *Menges'* theory of *linear partial information* (see [4]) assumes the uncertainty underlying the decision situation to be expressable by a *convex credal set* $\mathcal{M}$ on $(\Theta, 2^\Theta)$ of the form

$$\mathcal{M} := \big\{\pi | \ \underline{b}_s \leqslant \mathbb{E}_\pi(f_s) \leqslant \overline{b}_s \ \forall s = 1, \dots, r\big\} \tag{6}$$

where, for all $s = 1, \dots, r$, we have $(\underline{b}_s, \overline{b}_s) \in \mathbb{R}^2$ such that $\underline{b}_s \leqslant \overline{b}_s$ and $f_s : \Theta \to \mathbb{R}$. Note that these sets correspond to the credal sets induced by finite sets of gambles $\mathcal{K}$ from Walley's theory.

Here, criteria for decision making strongly depend on the actor's *attitude towards ambiguity*, i.e. the non-stochastic uncertainty between the measures contained in $\mathcal{M}$. Accordingly, many concurring criteria exist (see for instance [3]). Linear programming based results for a selection of them are presented in the following sections.

## 5. Checking maximality of pure actions

An action $a^* \in \mathbb{A}$ is said to be $\mathcal{M}$*-maximal*, if

$$\forall \, a \in \mathbb{A} \ \exists \, \pi_a \in \mathcal{M} : \quad \mathbb{E}_{\pi_a}(u_{a^*}) \geqslant \mathbb{E}_{\pi_a}(u_a) \tag{7}$$

Naturally, the above definition extends to randomized actions. For randomized actions, $\mathcal{M}$-maximality and E($\mathcal{M}$)-admissibility coincide. A algorithm for determining the set of all randomized E($\mathcal{M}$)-admissible actions has been introduced in [1, section 5.2].

However, for finite $\mathbb{A}$, being $\mathcal{M}$-maximal is a strictly weaker condition and, therefore, needs to be checked seperatly from E($\mathcal{M}$)-admissibility. Theorem 2 describes a linear programming based algorithm for checking $\mathcal{M}$-maximality of a pure $a^* \in \mathbb{A}$.

**Theorem 2.** *Let* $(\mathbb{A}, \Theta, u(\cdot))$ *denote a finite decision problem and let* $\mathcal{M}$ *be of the form (6). Consider the linear program*

$$\sum_{i=1}^{n} \Big(\sum_{j=1}^{m} {}^i\gamma_j\Big) \longrightarrow \max_{({}^1\gamma_1, \dots, {}^n\gamma_m)} \tag{8}$$

*with constraints* $({}^1\gamma_1, \dots, {}^n\gamma_m) \geqslant 0$ *and*

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

*Then* $a^* \approx (u_{*1}, \dots, u_{*m}) \in \mathbb{A}$ *is* $\mathcal{M}$*-maximal iff the optimal outcome of (8) equals* $n$.

## 6. Γ-Maximin and least favourable priors

For a probability measure $\pi$ on $(\Theta, 2^\Theta)$, let $B(\pi)$ denote the Bayes-utility w.r.t. $\pi$ (that is $B(\pi) = \mathbb{E}_\pi(u_{a^*})$, where $a^* \in \mathbb{A}$ denotes an arbitrary Bayes-action w.r.t. $\pi$). The set of all Bayes-actions w.r.t. $\pi$ is denoted by $\mathbb{A}_\pi$.

If $\mathcal{M}$ is a credal set of the form defined in (6), we call $\pi^- \in \mathcal{M}$ a *least favourable prior (lfp)* from $\mathcal{M}$ iff $B(\pi^-) \leqslant B(\pi)$ holds for all $\pi \in \mathcal{M}$. Theorem 3 describes a linear programming approach for determining a least favourable prior from $\mathcal{M}$.

**Theorem 3.** *Let* $(\mathbb{A}, \Theta, u(\cdot))$ *denote a decision problem and let* $\mathcal{M}$ *be of the form (6). Consider the linear program*

$$w_1 - w_2 \longrightarrow \min_{(w_1, w_2, \pi_1, \dots, \pi_m)} \tag{9}$$

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

- $\sum_{j=1}^m \pi_j = 1$
- $\underline{b}_s \leqslant \sum_{j=1}^m f_s(\theta_j) \cdot \pi_j \leqslant \overline{b}_s$ *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 the following holds:*

*i) Every optimal solution* $(w_1^*, \dots, \pi_m^*)$ *to (9) induces a least favourable prior* $\pi^- \in \mathcal{M}$ *by setting* $\pi^-(\{\theta_j\}) := \pi_j^*$.

*ii) There always exists a least favourable prior.*

Next, we show some connections between least favourable priors and randomized Γ-Maximin actions. We start by recalling the Γ-Maximin criterion: A randomized action $\lambda^* \in G(\mathbb{A})$ is said to be $\mathcal{M}$*-Maximin optimal* iff for all $\lambda \in G(\mathbb{A})$:

$$\min_{\pi \in \mathcal{M}} \mathbb{E}_\pi\big[G(u)_{\lambda^*}\big] \geqslant \min_{\pi \in \mathcal{M}} \mathbb{E}_\pi\big[G(u)_\lambda\big] \tag{10}$$

It turns out that the linear program from Theorem 3 is *dual* to the one for determining a $\mathcal{M}$-Maximin optimal randomized action described in [1, section 3.2]. Together with the *complementary slackness property* from linear optimization theory, this allows to derive deep connections between least favourable priors and the Γ-Maximin criterion.

**Theorem 4.** *Let* $(\mathbb{A}, \Theta, u(\cdot))$ *denote a finite decision problem and let* $\mathcal{M}$ *be of the form (6). Then the following holds:*

*i) If* $\pi^-$ *is a lfp from* $\mathcal{M}$*, then for all optimal randomized* $\mathcal{M}$*-Maximin actions* $\lambda^* \in G(\mathbb{A})$ *we have* $\lambda^*(\{a\}) = 0$ *for all* $a \in \mathbb{A} \setminus \mathbb{A}_{\pi^-}$.

*ii) Let* $\lambda^* \in G(\mathbb{A})$ *be an optimal randomized* $\mathcal{M}$*-Maximin action. If, for* $a \in \mathbb{A}$*, we have* $\lambda^*(\{a\}) > 0$*, then* $a \in \mathbb{A}_{\pi^-}$ *for all least favourable priors* $\pi^-$ *from* $\mathcal{M}$.

*iii) Let* $\pi^-$ *denote a lfp from* $\mathcal{M}$ *and let* $\lambda^* \in G(\mathbb{A})$ *denote a randomized* $\mathcal{M}$*-Maximin action. Then for all* $a \in \mathbb{A}_{\pi^-}$ *we have*

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

As an immediate consequence of Theorem 4, we can specify conditions under which randomization cannot improve utility, if optimality is defined in terms of the Γ-maximin criterion.

**Corollary 2.** *If there exists a least favourable prior* $\pi^-$ *from* $\mathcal{M}$ *such that* $\mathbb{A}_{\pi^-} = \{a_z\}$ *for some* $z \in \{1, \dots, n\}$*, then* $\delta_{a_z} \in G(\mathbb{A})$ *is the unique randomized* $\mathcal{M}$*-Maximin action. Particularly, randomization is unnecessary in such situations.*

## 7. A toy example

Consider the decision problem given by the table

| $u_{ij}$ | $\theta_1$ | $\theta_2$ | $\theta_3$ | $\theta_4$ |
|---|---|---|---|---|
| $a_1$ | 20 | 15 | 10 | 5 |
| $a_2$ | 30 | 10 | 10 | 20 |
| $a_3$ | 20 | 40 | 0 | 20 |

and assume that uncertainty is described by the credal set

$$\mathcal{M} := \{\pi | \ 0.3 \leqslant \pi_2 + \pi_3 \leqslant 0.7\}$$

- **Section 6:** Applying the algorithm from Theorem 3 gives the optimal solution $(13, 0, 0, 0, 0.7, 0.3)$. Thus, a least favourable prior $\pi^-$ from $\mathcal{M}$ is induced by the vector $(0, 0.7, 0.3, 0)$. Simple computation gives $\mathbb{A}_{\pi^-} = \{a_2\}$. Therefore, according to Corollary 2, $a_2$ is the unique $\mathcal{M}$-Maximin action (even compared to randomized actions) with utility 13.
- **Section 5:** Resolving the linear programming problem from Theorem 2 for actions $a_1, a_2$ and $a_3$ gives optimal value 3 for each of them. Thus, all available actions are $\mathcal{M}$-maximal.
- **Section 3:** Let $\tau$ denote the prior on $(\Theta, 2^\Theta)$ induced by $(0.2, 0.7, 0.05, 0.05)$ and let our trust in $\tau$ be expressed by $\alpha = 0.3$. Resolving the linear programming problem from Theorem 1 then gives the optimal solution $(8, 0, 0.8, 0, 0.2)$. Thus, a $\Phi_{\tau,0.3}$-optimal randomized action $\lambda^* \in G(\mathbb{A})$ is induced by $(0.8, 0, 0.2)$.

  Next, we can use the *constructive* proof of Corollary 1 to compute the measure $\mu_{\tau,0.3}$ on $(\Theta, 2^\Theta)$ defined in Corollary 1. The measure $\mu_{\tau,0.3}$ is induced by the vector $(0.070, 0.245, 0.656, 0.029)$.

**Implementation:** The R-code for the toy example is available on http://www.statistik.lmu.de/~ cjansen/index.html

## Outlook: Future research

Investigating further consequences of Theorem 4: What can we learn by restricting the set $\mathcal{M}$ to special cases (for instance *comparative probability* or *non-degenerated* credal sets)?

## References


[1] L.V. Utkin, T. Augustin. Powerful algorithms for decision making under partial prior information and general ambiguity attitudes. In: F.G. Cozman, R. Nau, T. Seidenfeld (eds.): *ISIPTA '05*, 2005, pp. 349-358.

[2] D. Kikuti, F.G. Cozman, C.P. de Campos. Partially ordered preferences in decision trees: computing strategies with imprecision in probabilities. In: R. Brafman, U. Junker (eds.): *Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, 2005, pp. 118-123.

[3] N. Huntley, R. Hable, M.C.M. Troffaes. Decision making. In: *Introduction to imprecise probabilities*. Ed. by T. Augustin, F.P.A. Coolen, G. de Cooman, M.C.M. Troffaes. Chichester: Wiley, 2014, pp. 190-206.

[4] E. Kofler, G. Menges. *Entscheiden bei unvollständiger Information*. Springer. Berlin (Lecture Notes in Economics and Mathematical Systems, 136), 1976.

[5] Hodges, Joseph L. and Lehmann, Erich L. The use of Previous Experience in Reaching Statistical Decisions. In: *The Annals of Mathematical Statistics* 23.3 (1952), pp. 396407.