Methods for eliciting preference systems with applications to decision making under severe uncertainty

Christoph Jansen, Georg Schollmeyer & Thomas Augustin
Department of Statistics, LMU Munich
IFORS 2021 Virtual

Problem and motivation

We consider the basic model of finite **Decision Theory**:

- $A = \{a_1, \dots, a_n\}$ denotes a finite set of *consequences*.
- $S = \{s_1, \dots, s_m\}$ denotes a finite set of *states*.
- $\mathcal{G} \subseteq A^S = \{X : S \to A\}$ denotes a finite set of *acts*.

Goal: Find optimal acts via some *choice function*

$$ch : 2^{\mathcal{G}} \to 2^{\mathcal{G}} \text{ with } ch(\mathcal{D}) \subseteq \mathcal{D} \text{ for all } \mathcal{D} \in 2^{\mathcal{G}}$$

that best possibly utilizes the available information.

Classical approach: *If both*

I) *preferences* on A are characterized by a cardinal utility $u : A \to \mathbb{R}$ *and*

II) *beliefs* on S are characterized by a classical probability π,

then one commonly *maximizes expected utility*, i.e. defines

$$\mathrm{ch}_{u,\pi}(\mathcal{D}) := \Big\{ Y \in \mathcal{D} : \mathbb{E}_\pi(u \circ Y) \geq \mathbb{E}_\pi(u \circ X) \text{ for all } X \in \mathcal{D} \Big\}$$

Problem and motivation, continued

Obviously: If assumptions I) and/or II) are not satisfied, then $\text{ch}_{u,\pi}(\mathcal{D})$ in general *won't be well-defined*.

Problem: In practice, this will often be the case.

(Requires *strong axiomatic assumptions*, e.g. the axioms of Savage)

Idea: Replace

- u by a set $\mathcal{U}$ of compatible utility functions on A and
- π by a set $\mathcal{M}$ of compatible probability measures on S

and generalize $\text{ch}_{u,\pi}$ to a choice function $\text{ch}_{\mathcal{U},\mathcal{M}}$ utilizing exactly the information encoded in $\mathcal{U}$ and $\mathcal{M}$ (and *nothing more* than that).

Details: There are several ways to proceed. We focus on the approach introduced in Jansen, Schollmeyer & Augustin (2018, Int. J. Approx. Reason).

(briefly summarized on the next three slides)

Modelling the set $\mathcal{U}$

Notation: Binary relation R has *strict part* P_R and *indifference part* I_R.

Preference system & Consistency

Let A denote a set of consequences. Let further

- $R_1 \subseteq A \times A$ be a binary relation on A
- $R_2 \subseteq R_1 \times R_1$ be a binary relation on R_1

The triplet $\mathcal{A} = [A, R_1, R_2]$ is called a *preference system* on A. We call $\mathcal{A}$ *consistent* if there exists $u : A \to [0, 1]$ such that for all $a, b, c, d \in A$:

- $(a, b) \in R_1 \Rightarrow u(a) \geq u(b)$ (with $=$ *iff* $\in I_{R_1}$).
- $((a, b), (c, d)) \in R_2 \Rightarrow u(a) - u(b) \geq u(c) - u(d)$ (with $=$ *iff* $\in I_{R_2}$).

The set of all representations u of $\mathcal{A}$ is denoted by $\mathcal{U}_{\mathcal{A}}$.

Interpretation of the components of $\mathcal{A}$:

- $(a, b) \in R_1$: *"a is at least as desirable as b"*
- $((a, b), (c, d)) \in R_2$: *"exchanging b by a is at least as desirable as d by c"*

Modelling the set $\mathcal{M}$

The agent's uncertainty among the elements of S is characterized by a polyhedral *credal set* of probability measures of the form

$$\mathcal{M} = \left\{ \pi \in \mathcal{P} : \underline{b}_\ell \leq \mathbb{E}_\pi(f_\ell) \leq \overline{b}_\ell \text{ for } \ell = 1, \ldots, r \right\}$$

where $\mathcal{P}$ denotes the set of all probability measures on $(S, 2^S)$ and

- $f_1, \ldots, f_r : S \to \mathbb{R}$ are real-valued mappings and
- $\underline{b}_\ell \leq \overline{b}_\ell$, $\ell = 1, \ldots, r$, are lower and upper expectation bounds.

$\Rightarrow$ Very general uncertainty model *capturing special cases* such as

- *Classical probability*
- *Interval probability*
- *Lower previsions*
- *Linear partial information*
- *Contamination models*

Decision making based on $\mathcal{U}_{\mathcal{A}}$ and $\mathcal{M}$

Jansen, Schollmeyer & Augustin (2018, Int. J. Approx. Reason) proposes several choice functions based on the sets $\mathcal{U}_{\mathcal{A}}$ and $\mathcal{M}$ and provide *linear programming based algorithms* for their evaluation.

We focus on one specific choice function, namely $\text{ch}_{\mathcal{A},\mathcal{M}} : 2^{\mathcal{G}} \to 2^{\mathcal{G}}$ with

$$\text{ch}_{\mathcal{A},\mathcal{M}}(\mathcal{D}) := \Big\{ Y \in \mathcal{D} : \nexists X \in \mathcal{D} \text{ s.t. } \mathbb{E}_{\pi}(u \circ X) \geq \mathbb{E}_{\pi}(u \circ Y) \text{ for all } u \in \mathcal{U}_{\mathcal{A}}, \pi \in \mathcal{M} \Big\}$$

The choice function $\text{ch}_{\mathcal{A},\mathcal{M}}$ ...

- ... selects acts that are *not expectation dominated* by any other act for arbitrary compatible pairs $(u, \pi) \in \mathcal{U}_{\mathcal{A}} \times \mathcal{M}$.
- ... can be evaluated by using *linear programming* theory.
- ... can be thought of as a generalization of *first order stochastic dominance* to partially cardinal and partially ordinal scaled spaces.

Main focus today: Eliciting $\mathcal{A}^* = [A, R_1^*, R_2^*]$ efficiently

Goal: Elicit an agent's true preference system

$$\mathcal{A}^* = [A, R_1^*, R_2^*]$$

by asking as few as possible ranking questions only about R_1^*.

Two different elicitation procedures:

- **Procedure 1:** For every presented pair $\{a_i, a_j\}$ with $(a_i, a_j) \in R_1^*$, we measure the agent's **consideration time** $t_{ij} > 0$ and use these times for constructing R_2 (hopefully matching R_2^*).
- **Procedure 2:** For every presented pair $\{a_i, a_j\}$ with $(a_i, a_j) \in R_1^*$, we collect a label of **preference strength** and utilize the collected labels for constructing R_2 (hopefully matching R_2^*).

Question: Under which conditions do Procedures 1 and 2 produce the agent's true preference system $\mathcal{A}^* = [A, R_1^*, R_2^*]$?

Procedure 1: Time elicitation

Time elicitation

Input: $A = \{a_1, \ldots, a_n\}$; $R_1 = \emptyset$; $C = \emptyset$;

Output: $\mathcal{A} = [A, R_1, R_2]$;

Procedure: Present all pairs $\{a_i, a_j\}$ from $A_{\{2\}} := \{\{a, b\} : a \neq b \in A\}$.

i) Agent judges a_i and a_j incomparable. Set $C = C \cup \{(a_j, a_i), (a_i, a_j)\}$ and $t_{ij} = t_{ji} = 0$.

ii) Agent ranks a_i strictly better than a_j. Set $R_1 = R_1 \cup \{(a_i, a_j)\}$ and measure consideration time $t_{ij} > 0$. Set $t_{ji} = 0$.

iii) Agent ranks a_j strictly better than a_i. Set $R_1 = R_1 \cup \{(a_j, a_i)\}$ and measure consideration time $t_{ji} > 0$. Set $t_{ij} = 0$.

iv) Agent is indifferent between a_j and a_i. Set $R_1 = R_1 \cup \{(a_i, a_j), (a_j, a_i)\}$ and $t_{ij} = t_{ji} = c_\infty > \max\{t_{pq} : (a_p, a_q) \in P_{R_1^*}\}$.

Define R_2 by setting $((a_i, a_j), (a_k, a_l)) \in R_2 \quad :\Leftrightarrow \quad t_{kl} - t_{ij} \geq 0 \ \wedge \ t_{ij} > 0$.

Procedure 1: Assumptions

Assumption 1

For $(a_i, a_j), (a_k, a_l) \in R_1^*$ the following holds:

i) $t_{kl} > t_{ij} > 0$ if and only if $((a_i, a_j), (a_k, a_l)) \in P_{R_2^*}$

ii) $t_{kl} = t_{ij} > 0$ if and only if $((a_i, a_j), (a_k, a_l)) \in I_{R_2^*}$

iii) $t_{ij} = t_{ji} = c_\infty$ if and only if $(a_i, a_j) \in I_{R_1^*}$

Assumption 2

For $(a_i, a_j), (a_j, a_k) \in P_{R_1^*}$ we have $\frac{1}{t_{ij}} + \frac{1}{t_{jk}} = \frac{1}{t_{ik}}$, whenever $(a_i, a_k) \in P_{R_1^*}$.

Assumption 3

For $(a_i, a_j) \in I_{R_1^*}$ we have

i) $t_{ki} = t_{kj}$ whenever $(a_k, a_i), (a_k, a_j) \in P_{R_1^*}$ and

ii) $t_{ik} = t_{jk}$ whenever $(a_i, a_k), (a_j, a_k) \in P_{R_1^*}$.

Procedure 1: Findings

Proposition 1

Under Assumption 1, time elicitation produces the agents's true preference system $\mathcal{A}^* = [A, R_1^*, R_2^*]$.

Proposition 2

Under Assumptions 1, 2 and 3 the true preference system $\mathcal{A}^* = [A, R_1^*, R_2^*]$ is consistent if and only if R_1^* is transitive.

Procedure 1*: Suppose that after k steps of Procedure 1 we have elicited R_1^k and C^k. Sample the next pair to present from

$$A_{\{2\}} \setminus \left\{\{a,b\} : (a,b) \in H_{R_1^k} \vee (b,a) \in H_{R_1^k} \vee (a,b) \in C^k\right\}$$

and *compute* the missing times by using Assumption 2.

Proposition 3

Under Assumptions 1, 2 and 3, Procedure 1* terminates in $\mathcal{A}^*$ if and only if R_1^* is transitive. By Proposition 2 we know $\mathcal{A}^*$ is consistent in this case.

Procedure 2: Label elicitation

Setup: Agent assigns a label $\ell_r^{ij} \in \mathcal{L}_r := \{\mathbf{n}, \mathbf{c}, 0, 1, \ldots, r\}$ to every (a_i, a_j) by some *labelling function* $\ell_r : A \times A \to \mathcal{L}_r$:

$\mathbf{n}$:	non-comparable
$\mathbf{c}$:	strict preference of unknown strength
0 :	indifferent
$1, \ldots, r$:	strict preference of increasing strength

Label elicitation

Input: $A = \{a_1, \ldots, a_n\}$; $R_1 = \emptyset$; number of labels r;

Output: $\mathcal{A} = [A, R_1, R_2]$;

Procedure: Present all pairs $(a_i, a_j) \in A \times A$.

i) If $\ell_r^{ij} \in \mathcal{L}_r \setminus \{\mathbf{n}, 0\}$, set $R_1 = R_1 \cup \{(a_i, a_j)\}$.

ii) If $\ell_r^{ij} = 0$, set $R_1 = R_1 \cup \{(a_i, a_j), (a_j, a_i)\}$.

iii) If $\ell_r^{ij} = \mathbf{n}$, set $R_1 = R_1$.

Define R_2 by setting $((a_i, a_j), (a_k, a_l)) \in R_2 \quad :\Leftrightarrow \quad \ell_r^{ij} > \ell_r^{kl} \quad \vee \quad \ell_r^{ij} = \ell_r^{kl} = 0$

Procedure 2: Assumptions

Assumption 4

i) $(a_i, a_j) \in I_{R_1^*} \quad \Leftrightarrow \quad \ell_r^{ij} = 0$

ii) $(a_i, a_j) \in P_{R_1^*} \quad \Leftrightarrow \quad \ell_r^{ij} \in \mathcal{L}_r \setminus \{\mathbf{n}, 0\} \ \wedge \ \ell_r^{ji} = \mathbf{n}$

iii) $(a_i, a_j) \in C_{R_1^*} \quad \Leftrightarrow \quad \ell_r^{ij} = \ell_r^{ji} = \mathbf{n}$

Assumption 5

For all $(a_i, a_j), (a_k, a_l) \in R_1^*$ the following holds:

i) $\ell_r^{ij} > \ell_r^{kl} \quad \Rightarrow \quad ((a_i, a_j), (a_k, a_l)) \in P_{R_2^*}$

ii) $\ell_r^{ij} = \ell_r^{kl} = 0 \quad \Rightarrow \quad ((a_i, a_j), (a_k, a_l)) \in I_{R_2^*}$

iii) $\ell_r^{ij} = \mathbf{c} \ \vee \ \ell_r^{kl} = \mathbf{c} \quad \Leftrightarrow \quad ((a_i, a_j), (a_k, a_l)) \in C_{R_2^*}$

Assumption 6

For all $((a_i, a_j), (a_k, a_l)) \in P_{R_2^*}$ the statement $\ell_r^{ij} = \ell_r^{kl} = x \notin \{0, \mathbf{n}, \mathbf{c}\}$ implies that $\{1, \ldots, r\} \subset \ell_r(A \times A)$.

Procedure 2: Findings

Proposition 4

The following two statements hold true:

i) If, for some $r \in \mathbb{N}$, $\ell_r : A \times A \to \mathcal{L}_r$ satisfies Assumptions 4 and 5, then Procedure 2 produces a sub-system of the decision maker's true preference system $\mathcal{A}^*$. Particularly, the procedure produces a consistent preference system whenever $\mathcal{A}^*$ is consistent.

ii) There exists $r_0 \in \mathbb{N}$ such that if $\ell_{r_0} : A \times A \to \mathcal{L}_{r_0}$ satisfies Assumptions 4, 5 and 6, then Procedure 2 produces the true $\mathcal{A}^*$.

Challenge: Although Prop. 4 ii) guarantees that Procedure 2 reproduces the agent's true preference system for some number of labels r^*, labelling in accordance with the assumptions might be *too demanding* if r^* is large.

Solution: Use a *relatively small* number of labels and restart elicitation on those pairs with equal label. Stop as soon as you know that equal labels *originate from indifference*.

Graphical intuition:

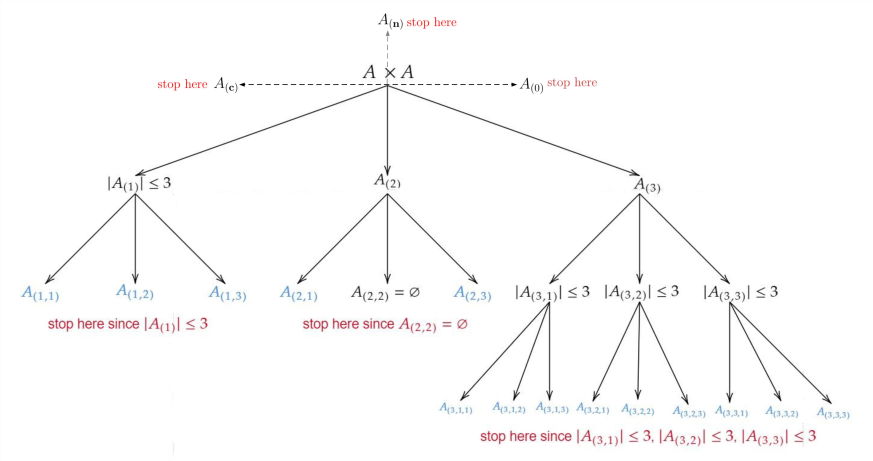

Hierarchical version: Findings

For the hierarchical version of label elicitation to work, we need to assume that the agent is able to *adapt* the labelling function to arbitrary subsets.

Formally, we arrive at:

Assumption 7

For every $N \subseteq A \times A$ the labels on the restricted set of pairs N are given w.r.t. a labelling function $\ell_{(N,r)} : N \to \mathcal{L}_r$ satisfying Assumptions 4, 5 and 6.

This indeed allows the following Proposition:

Proposition 5

Let Assumption 7 hold true. For $n = |A|$ consequences and $r \geq 2$ labels, the hierarchical version of Procedure 2 terminates in $\mathcal{A}^*$ after at most $\max\{1, \lceil \frac{n^2-r}{r-1} \rceil + 1\}$ elicitation rounds.

Application to decision making under uncertainty

We now return to *decision under uncertainty*:

- Consider the decision problem $\mathcal{G}$ under uncertainty model $\mathcal{M}$.
- Suppose $\mathcal{A}^*$ is elicited by either Procedure 1 or 2 (or some variant).
- Let $\mathcal{A}_1, \mathcal{A}_2, \dots$ be the preference system after elicitation step $1, 2, \dots$.

Proposition 6

Let the assumptions of the used procedure be satisfied. Then, for any k :

$$X \in ch_{\mathcal{A}_k,\mathcal{M}}(\mathcal{G}) \quad \Rightarrow \quad X \in ch_{\mathcal{A}^*,\mathcal{M}}(\mathcal{G})$$

Why is this good?

If an act is optimal w.r.t. the preference system $\mathcal{A}_k$ *elicited so far*, we can terminate elicitation and *conclude* that it is optimal also w.r.t. the agent's *true preference system* $\mathcal{A}^*$.

A small example

- Consider the following decision problem:

	s_1	s_2	s_3	s_4
X_1	a_8	a_5	a_2	a_3
X_2	a_7	a_6	a_4	a_1

Decision problem

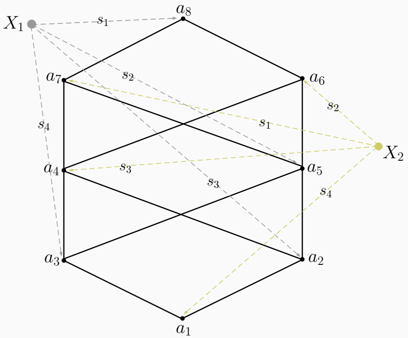

Hasse diagram of R_1^*

- The relation R_2^* is given as the transitive hull of (where $e_{ij} := (a_i, a_j)$):

$$e_{31} P_{R_2^*} e_{52} P_{R_2^*} e_{74} P_{R_2^*} e_{21} I_{R_2^*} e_{64} I_{R_2^*} e_{42} I_{R_2^*} e_{86} P_{R_2^*} e_{87} P_{R_2^*} e_{53} P_{R_2^*} e_{75} P_{R_2^*} e_{65} P_{R_2^*} e_{43}$$

- Let $\mathcal{M} = \{\pi\}$, where π is the uniform distribution on S.

A small example, continued

Assume elicitation is done by using Procedure 2 with $\ell_5 : A \times A \to \mathcal{L}_5$.

Moreover, assume the first four elicitation steps look as follows:

Elicitation step	Presented pair	Label of the pair
1	(a_8, a_7)	$\ell_5^{87} = 2$
2	(a_6, a_5)	$\ell_5^{65} = 1$
3	(a_3, a_1)	$\ell_5^{31} = 3$
4	(a_4, a_2)	$\ell_5^{42} = 2$

Then, for every $u \in \mathcal{U}_{\mathcal{A}_4}$ we can go on computing (where $u_i := u(a_i)$):

$$4 \cdot (\mathbb{E}_\pi(u \circ X_1) - \mathbb{E}_\pi(u \circ X_2)) = \underbrace{(u_8 - u_7) - (u_6 - u_5)}_{>0,\ since\ (e_{87}, e_{65}) \in P_{R_2}} + \underbrace{(u_3 - u_1) + (u_4 - u_2)}_{>0,\ since\ (e_{31}, e_{42}) \in P_{R_2}} > 0$$

Thus $X_1 \in ch_{\mathcal{A}_4, \mathcal{M}}(\mathcal{G})$. Thus $X_1 \in ch_{\mathcal{A}^*, \mathcal{M}}(\mathcal{G})$ by Prop. 6.

!! *We concluded that X_1 is optimal by asking four simple ranking questions.* **!!**

Ongoing and future research

There are several promising perspectives for future research:

- Finding data-driven methods for presenting the *most promising* pair of consequences in each elicitation step. (*Learn from previous rounds.*)
- Develop methods that *flexibly mix hierarchical and non-hierarchical procedures* to speed up elicitation.
- Investigate stopping properties of the procedures for choice functions other than $ch_{\mathcal{A},\mathcal{M}}$.