Resolute choice in sequential decision problems with multiple priors

Hélène Fargier, Gildas Jeantet et Olivier Spanjaard
IMPRECISE PROBABILITIES WORKSHOP
Toulouse, May 29th, 2015

UPMC SORBONNE UNIVERSITÉS
UNIVERSITÉ TOULOUSE III PAUL SABATIER

Sequential decision-making under uncertainty

- Consequences of actions are dependent on states of the world
- Applications : medical diagnosis, troubleshooting under uncertainty, poker-playing program, etc.
- Graphical models :
  - decision trees (Raiffa, 1968)
  - influence diagrams (Shachter, 1986)
  - MDPs (Dean *et al.*, 1993; Kaebling et al., 1999)
- Sometimes hard to elicit sharp probabilities (several experts, missing data)

Need for models and algorithms
for dealing with imprecise probabilities

Expected utility model

Given an act $f : \Theta \to X$
where Θ is the set of states of the world
X is the set of consequences
and $u: R \to R$ a nondecreasing function:

$$EU_{P,u}(f) = \sum_\theta P(\theta).u(f(\theta))$$

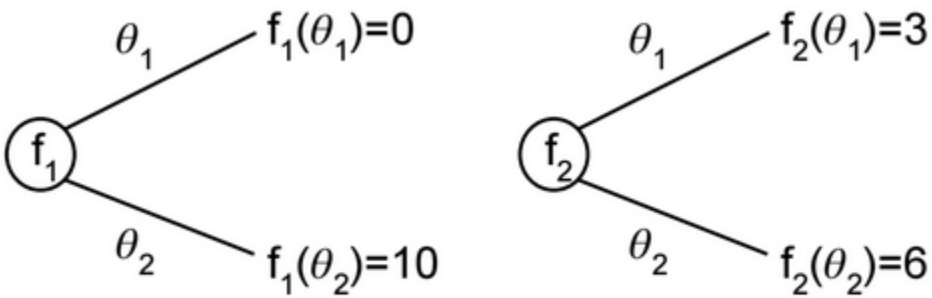

If $u(x)=\sqrt{x}$ then:

$EU(f_1)=½.u(0)+½.u(10) \approx 1.6 < 2.1 \approx 1/2.u(3)+1/2u(6)=EU(f_2)$

$$f_2 \succ f_1$$

Ambiguity: Ellsberg's example

Ellsberg's urn: ⅓ of red balls, ⅔ of black or yellow balls.

Lottery	Red	Black	Yellow
f_R	1	0	0
f_B	0	1	0
f_{RY}	1	0	1
f_{BY}	0	1	1

Usually: $f_R \succ f_B$ and $f_{BY} \succ f_{RY}$

There exists no probability P and utility function u such that:

$$EU_{P,u}(f_R) > EU_{P,u}(f_B) \text{ and } EU_{P,u}(f_{BY}) > EU_{P,u}(f_{RY})$$

Min expected utility model

Multiple priors

The underlying probability measure P could be any probability in the set:

$$\mathscr{P} = \{ P : P(\text{Red})=1/3,\ P(\text{Black or Yellow})=2/3 \}$$

Min expected utility (Gilboa & Schmeidler, 1989)

Most DM do use the EU model, but on the basis of the whole *set* of priors. They maximize the min, over $\mathscr{P}$, of the possible values of EU:

$$\underline{EU}_{\mathscr{P},u}(f) = \min_{P\in\mathscr{P}} EU_{P,u}(f)$$

Back to Ellsberg's example

Ellsberg's urn: ⅓ of red balls, ⅔ of black or yellow balls.

Lottery	Red	Black	Yellow
f_R	1	0	0
f_B	0	1	0
f_{RY}	1	0	1
f_{BY}	0	1	1

$\underline{EU}_{\mathscr{P},u}(f_R) = 1/3$

$\underline{EU}_{\mathscr{P},u}(f_B) = 1/3$

$\underline{EU}_{\mathscr{P},u}(f_{BY}) = 0$

$\underline{EU}_{\mathscr{P},u}(f_{RY}) = 2/3$

A sequential game with ambiguity

Consider the following game:

1. toss a coin;
2. a ball is drawn from an Ellsberg's urn ($\frac{1}{3}$ of red balls, $\frac{2}{3}$ of black or yellow balls):
   a. if the coin comes up heads, then bet on red or black
   b. if the coin comes up tails, then bet on red or yellow
3. if the guess is wrong, then win 0, otherwise win 1-ε if red, 1 if another color.

Decision tree with ambiguity

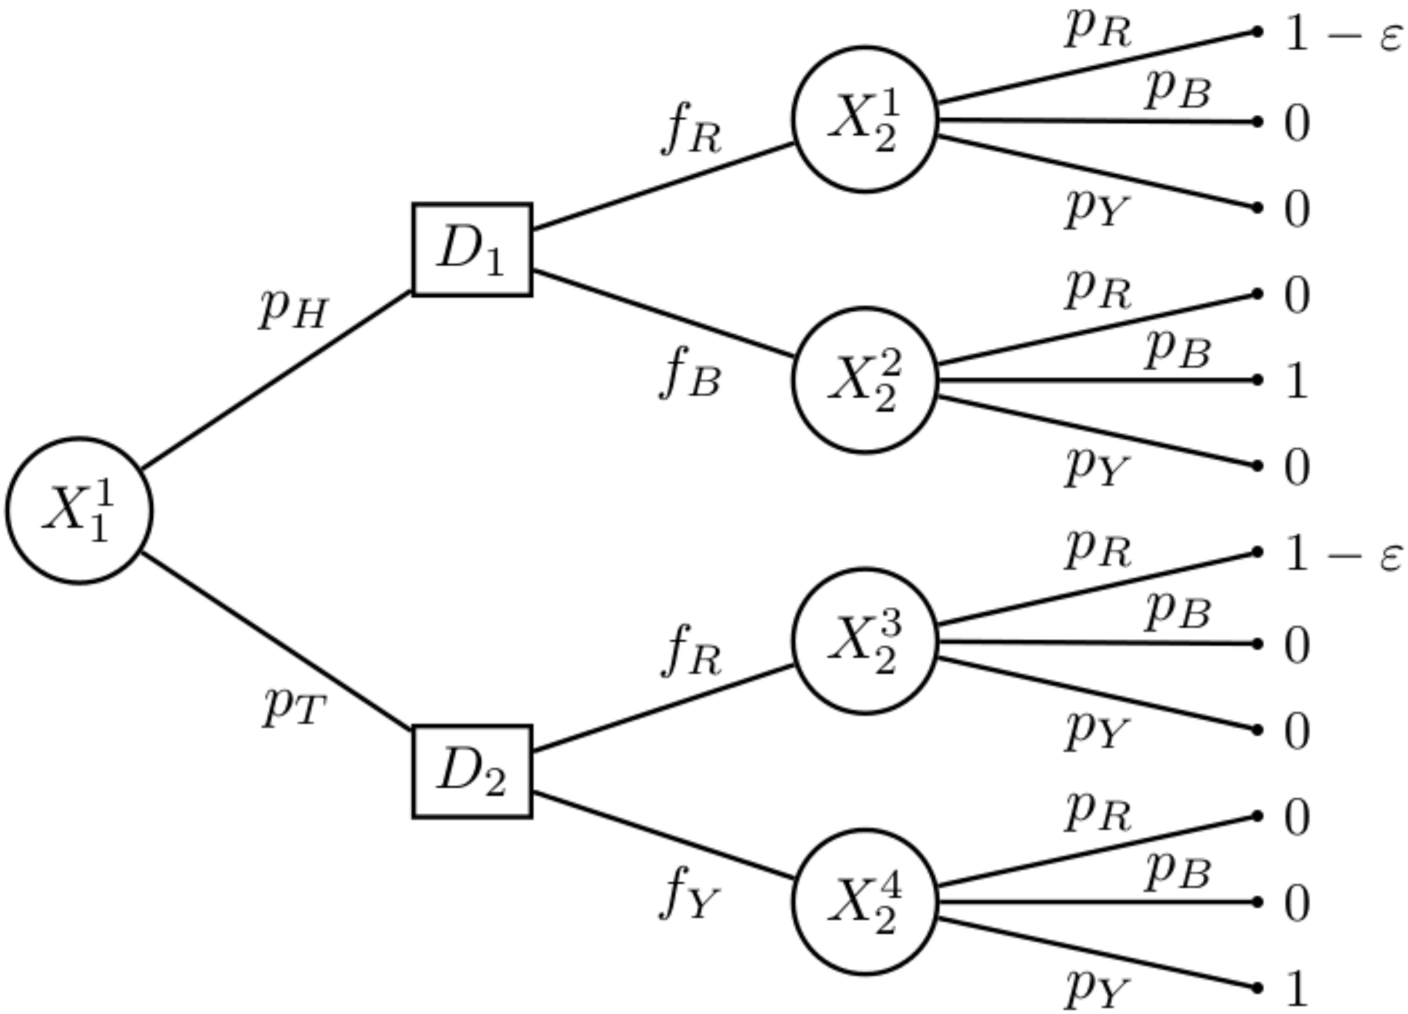

Related works

Two research directions:

- assume dynamic feasibility [Kikuti et al., 2011] (seeking a strategy returned by rolling back the decision tree): strategy followed by consequentialist decision maker, i. e. a DM whose present decision does not depend on the past nor on what she planned to do when making her first decision.

Pros: appealing from an algorithmic viewpoint

Cons: it may return a dominated strategy [Hammond, 1988]

- follow a resolute choice approach [McClennen, 1990]: commit to an initial strategy and never deviate later

Huntley and Troffaes [2008] proposed a generic method.

Evaluating a strategy

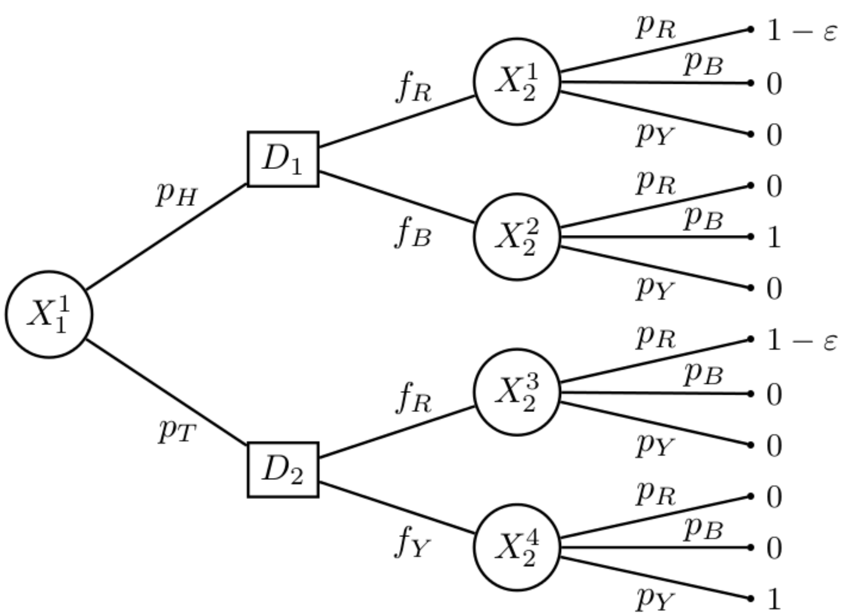

Evaluating strategy

($D_1 = f_B$, $D_2 = f_Y$)

amounts to evaluate compound lottery:

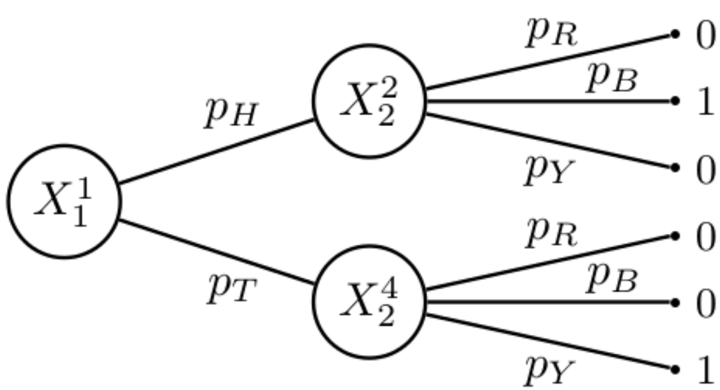

with $p_H=½$, $p_T=½$,
$p_R=⅓$ and $p_B+p_Y=⅔$.

Evaluating a strategy

Proposition. Evaluating a strategy according to its min expected utility is an NP-hard problem, even if all non-degenerated probability intervals are [0,1].

Proof. Reduction from 3-SAT:

$$(x_1 \vee x_2 \vee x_3) \wedge (\overline{x_1} \vee x_3 \vee x_4)$$

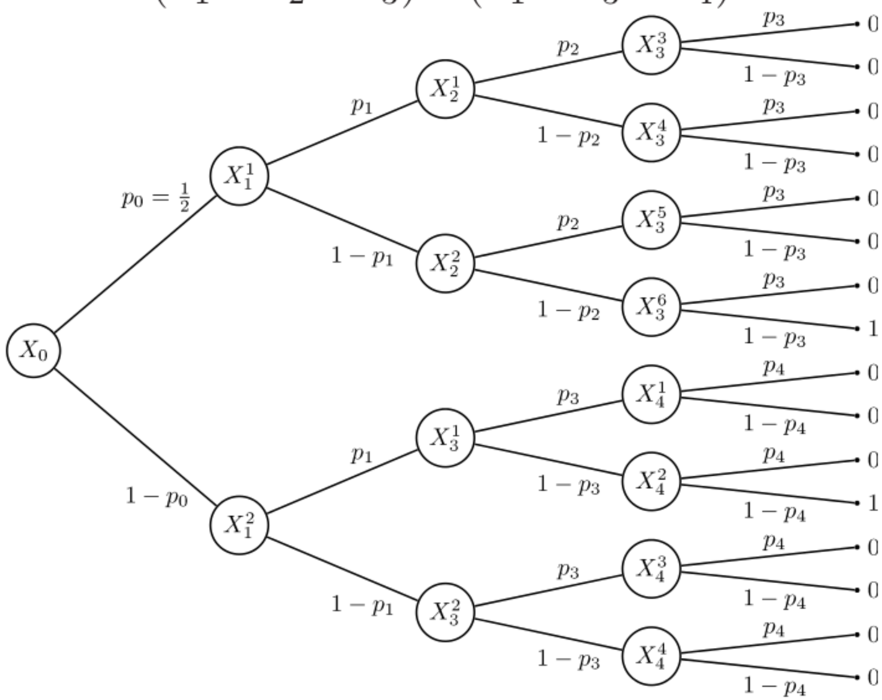

The 3-SAT formula is satisfiable iff min expected utility = 0.

Evaluating a strategy

The evaluation of a compound lottery can be done via a mathematical programming formulation, with one variable for each instanciation of $X=\langle X_1,\ldots,X_n\rangle$:

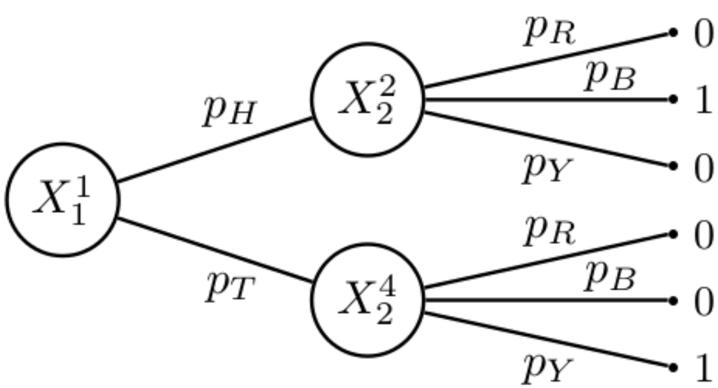

$$\min_{P\in\mathscr{P}} P(X_1=H, X_2=B) + P(X_1=T, X_2=Y)$$

where $\mathscr{P}$ denotes the set of possible probability measures over the considered decision tree.

Characterizing set $\mathscr{P}$: main difficulty

Oil wildcatter problem

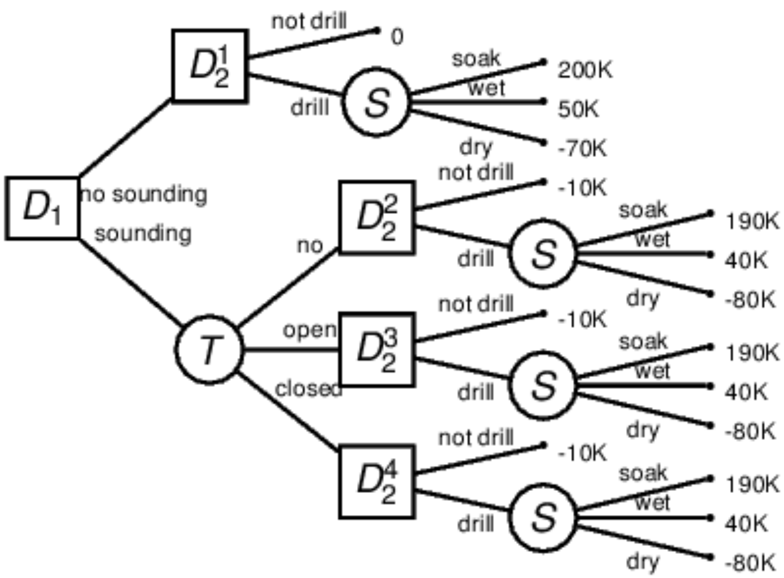

$P(S\|T)$	dry	wet	soak
no	[0.500,0.666]	[0.222,0.272]	[0.125,0.181]
open	[0.222,0.333]	[0.363,0.444]	[0.250,0.363]
closed	[0.111,0.166]	[0.333,0.363]	[0.454,0.625]

T	no	open	closed
$P(T)$	[0.181,0.222]	[0.333,0.363]	[0.444,0.454]

S	dry	wet	soak
$P(S)$	[0.214,0.344]	[0.309,0.386]	[0.307,0.456]

Characterizing set $\mathscr{P}$: main difficulty

Oil wildcatter problem

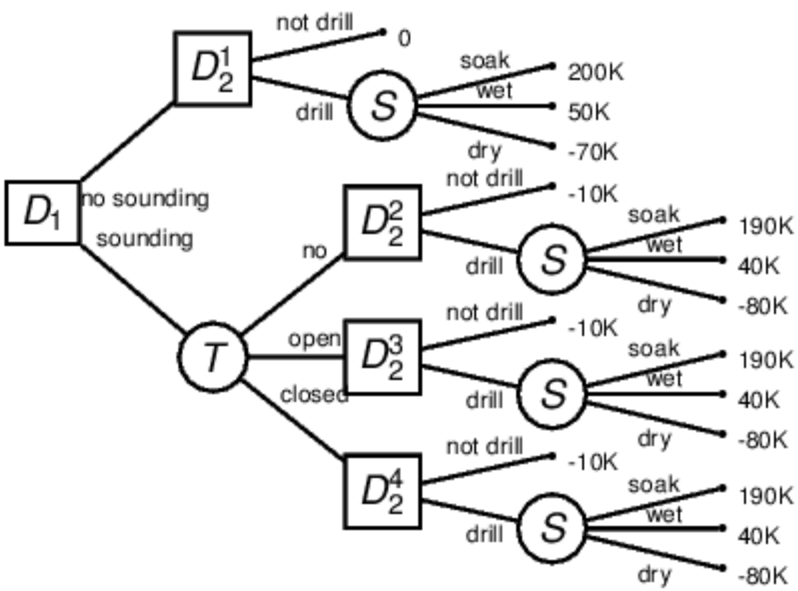

$P(S\|T)$	dry	wet	soak
no	0.55	[0.222,0.272]	[0.125,0.181]
open	0.33	[0.363,0.444]	[0.250,0.363]
closed	0.12	[0.333,0.363]	[0.454,0.625]

T	no	open	closed
$P(T)$	0.20	0.35	0.45

S	dry	wet	soak
$P(S)$	0.22	[0.309,0.386]	[0.307,0.456]

The total probability theorem does not hold:

$P(S = dry|T = no)P(T = no) + P(S = dry|T = open)P(T = open) + P(S = dry|T = closed)P(T = closed) = 0.2795 \neq 0.22 = P(S = dry)$

Selecting a strategy

The optimality principle does not hold:

In D_1: $EU(f_R) = (1-\varepsilon)/3 > 0 = EU(f_B) \Rightarrow$ the DM prefers f_R

In D_2: $EU(f_R) = (1-\varepsilon)/3 > 0 = EU(f_Y) \Rightarrow$ the DM prefers f_R

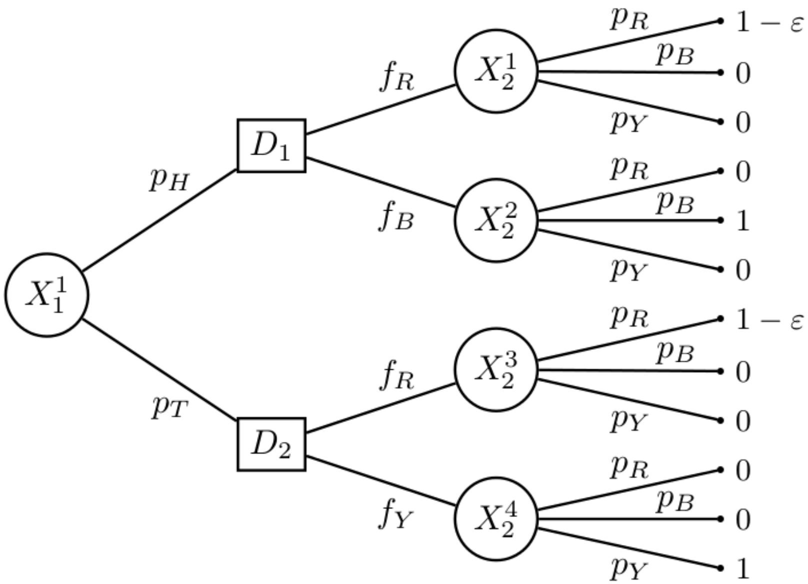

Strategy returned by rolling back: $(D_1=f_R, D_2=f_R)$

Min expected utility of $(D_1=f_R, D_2=f_R)$: $(1-\varepsilon)/3$

Min expected utility of $(D_1=f_B, D_2=f_Y)$: $1/3$

Selecting: separable case

Separable decision tree. For each chance node X_i^j, we denote by $\mathcal{P}_i^j$ the set of conditional probability distributions over $X_i | past(X_i^j)$ that satisfies constraints C_i^j. A decision tree $\mathcal{T}$ is called *separable* (or *separately specified*, Kikuti *et al.*, 2011) if

$$\mathcal{P}_{\mathcal{T}} = \prod_{X_i^j \in \mathcal{N}_C} \mathcal{P}_i^j.$$

Example: sequential variant of Ellsberg's urn with two distinct Ellsberg's urns.

The optimal strategy can be computed by rolling back the decision tree.
It involves the solution of a (small) linear program at each chance node, where the variables are the conditional probabilities.

Selecting: non-separable case

Dominance relation. A strategy s dominates s' if:

$$\forall\ P \in \mathscr{P},\ EU_{P,u}(s) \leq EU_{P,u}(s').$$

Dominance test: mathematical programming.

If s dominates s', then $\underline{EU}_{\mathscr{P},u}(s) \leq \underline{EU}_{\mathscr{P},u}(s')$.

Two-phases approach:

1. Compute the set ND of non-dominated strategies by rolling back the decision tree [Huntley & Troffaes, 2008]
2. Determine an optimal strategy in ND

Numerical tests

Separable decision trees (times in sec.)

$h \setminus d$	3		4	
	# nodes	time	# nodes	time
8	2073	< 1	5851	< 1
10	12441	< 1	46811	< 1
12	74639	1.2	374491	5.7
14	447897	7.8	2995931	65.5
16	2687385	57.3	X	X
18	X	X	X	X

h: depth of the decision tree;
d: outdegree of chance nodes.

Non-separable decision trees (times in sec.)

h	8		10	
$n \setminus w$	0.05	0.1	0.05	0.1
7	< 1	< 1	22.0	23.0
8	1.4	2.0	51.6	58.0
9	2.2	2.4	114.9	142.3
10	4.3	4.6	253.6	328.1
11	7.4	7.7	590.7	X
12	14.3	17.8	X	X
ND	5	8	29	32

n: number of random variables; w: imprecision degree;
ND: average size of the non-dominated set.

Algorithms implemented in C++.
CPLEX solved used to solve the mathematical programs.
Numerical tests performed on a Pentium IV 2.13Ghz CPU computer, 3GB RAM

Research directions

- Use credal networks to define set $\mathscr{P}$.
- Extend to influence diagrams with imprecise probabilities
- Resolute choice with selves [Jaffray & Nielsen, 2006]:

Consider each decision node as a self and search for a compromise between the selves

More specifically: define a regret for each self, and compute a strategy that optimizes an aggregation of the regrets.

Thank you