Factuality and Locality

with Arbitrary Choice Functions

Matthias C. M. Troffaes Nathan Huntley Ricardo Shirota

Durham University

11th September 2009

The crazy three

- Matthias C. M. Troffaes
  - lecturer in statistics at Durham University
  - foundations of statistics
  - sequential decision making, dynamic programming
  - imprecise probabilities
  - reliability, fault trees

The crazy three

- Matthias C. M. Troffaes
  - lecturer in statistics at Durham University
  - foundations of statistics
  - sequential decision making, dynamic programming
  - imprecise probabilities
  - reliability, fault trees
- Nathan Huntley
  - PhD student at Durham University
  - supervised by Matthias Troffaes
  - sequential decision making
  - imprecise probabilities

The crazy three

- Matthias C. M. Troffaes
  - lecturer in statistics at Durham University
  - foundations of statistics
  - sequential decision making, dynamic programming
  - imprecise probabilities
  - reliability, fault trees
- Nathan Huntley
  - PhD student at Durham University
  - supervised by Matthias Troffaes
  - sequential decision making
  - imprecise probabilities
- Ricardo Shirota
  - PhD student at the University of Sao Paulo
  - supervised by Prof. Fabio G. Cozman
  - currently visiting Durham
  - Markov decision processes (sequential decision making)
  - imprecise probabilities
  - artificial intelligence planning
  - algorithms

Outline

Outline

1. Introduction
   - Problem Description
   - Gambles and Choice Functions
   - Decision Trees
2. Factuality
   - Definition
   - Necessary and Sufficient Conditions
   - Implications and Examples
3. Locality
   - Definition
   - Necessary and Sufficient Conditions
   - Implications and Examples

Problem Description

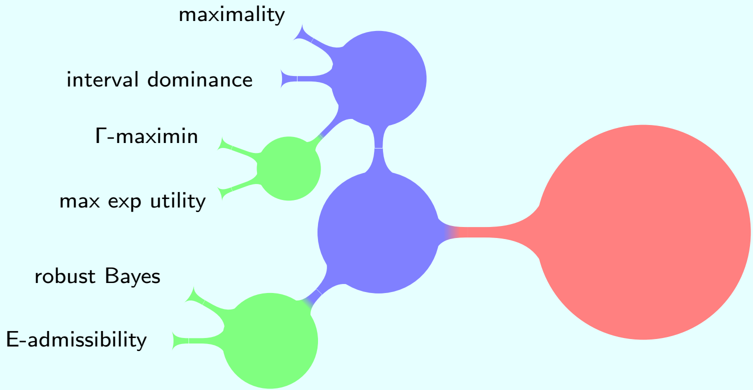

Problem Description

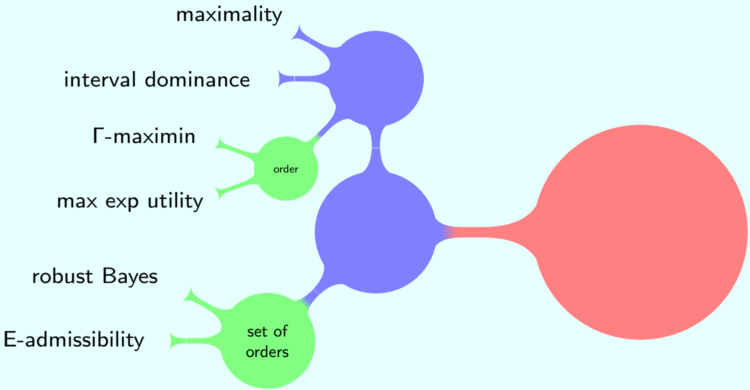

Problem Description

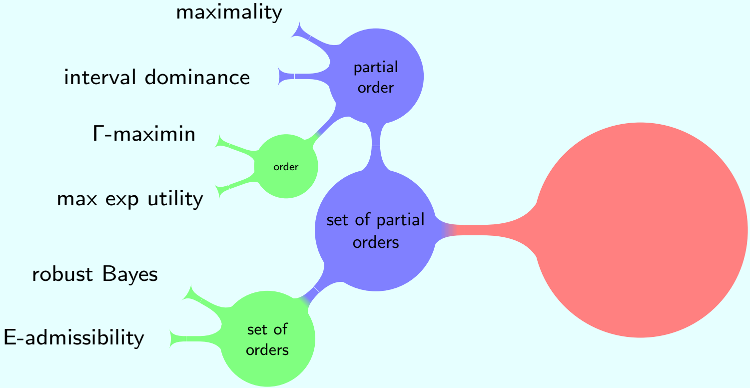

Problem Description

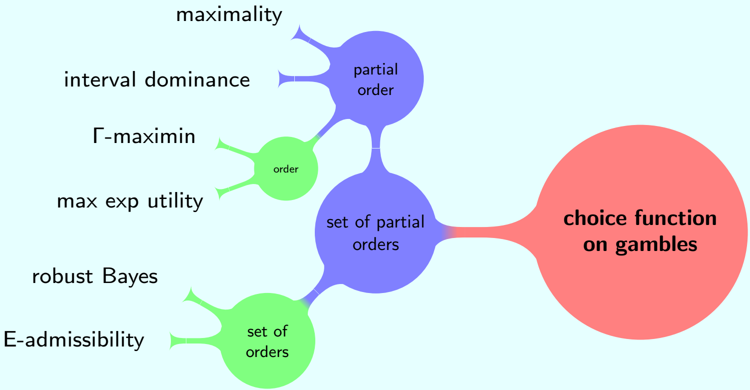

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

- common aspects of all such generalizations?

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

- common aspects of all such generalizations?
  - events E on some possibility space Ω

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

- common aspects of all such generalizations?
  - events E on some possibility space Ω
  - rewards r in some reward set $\mathcal{R}$

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

- common aspects of all such generalizations?
  - events E on some possibility space Ω
  - rewards r in some reward set $\mathcal{R}$
  - decisions d in some decision set D

Problem Description

Aim: study sequential decision processes for all of these typical generalizations of maximizing expected utility.

- common aspects of all such generalizations?
  - events E on some possibility space Ω
  - rewards r in some reward set $\mathcal{R}$
  - decisions d in some decision set D
  - some means of selecting decisions based on uncertain rewards

choice function opt on sets of gambles

Gambles and Choice Functions

Definition

A gamble is an uncertain reward, i.e. a mapping from the possibility space Ω to the reward set $\mathcal{R}$.

"probabilityless (horse-)lottery"

Gambles and Choice Functions

Definition

A gamble is an uncertain reward, i.e. a mapping from the possibility space Ω to the reward set $\mathcal{R}$.

"probabilityless (horse-)lottery"

Definition

A choice function opt selects, for any set of gambles $\mathcal{X}$ and event A, a subset of $\mathcal{X}$:

$$\emptyset \neq \text{opt}(\mathcal{X}|A) \subseteq \mathcal{X}$$

Gambles and Choice Functions

Definition

A gamble is an uncertain reward, i.e. a mapping from the possibility space Ω to the reward set $\mathcal{R}$.

"probabilityless (horse-)lottery"

Definition

A choice function opt selects, for any set of gambles $\mathcal{X}$ and event A, a subset of $\mathcal{X}$:

$$\emptyset \neq \text{opt}(\mathcal{X}|A) \subseteq \mathcal{X}$$

How to solve sequential decision problems
with a choice function?

Decision Trees: Example

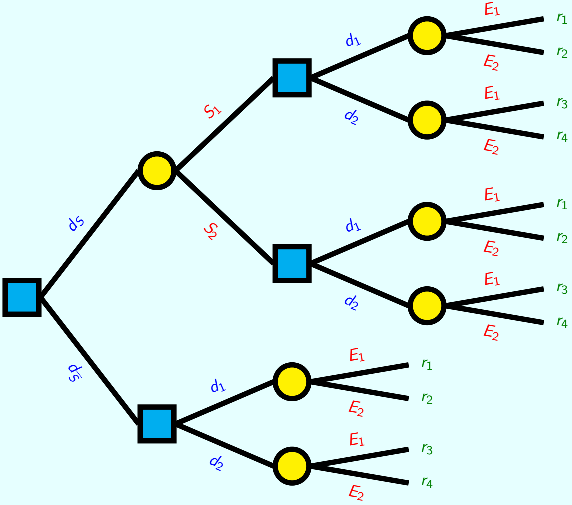

Decision Trees: Example

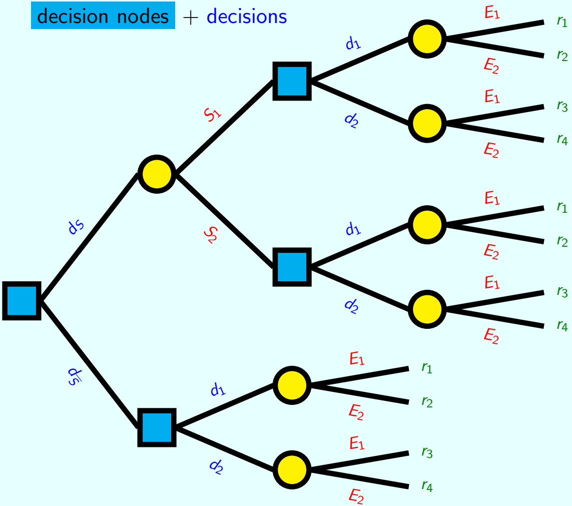

Decision Trees: Example

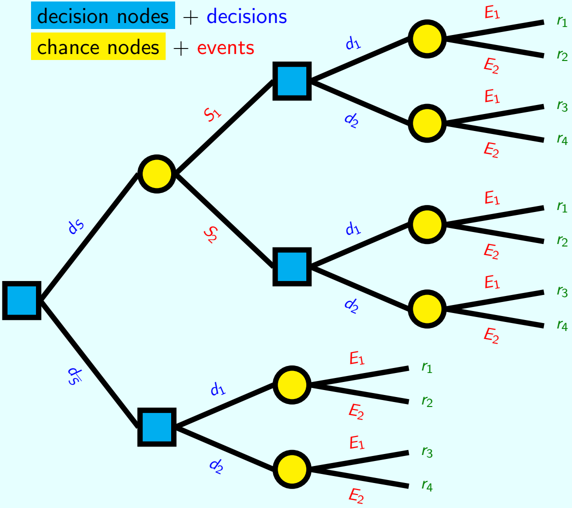

Decision Trees: Example

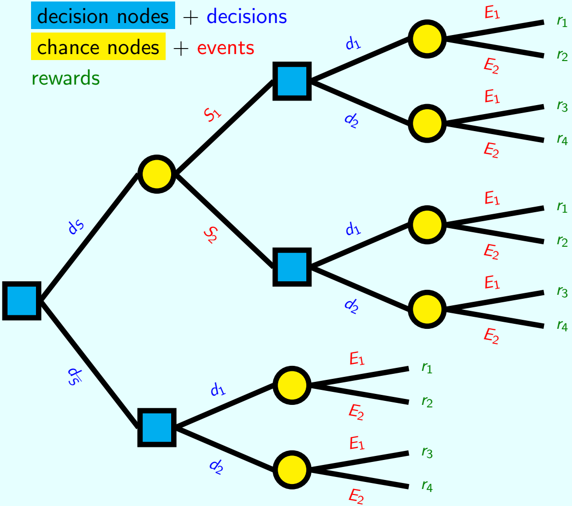

Decision Trees: Example

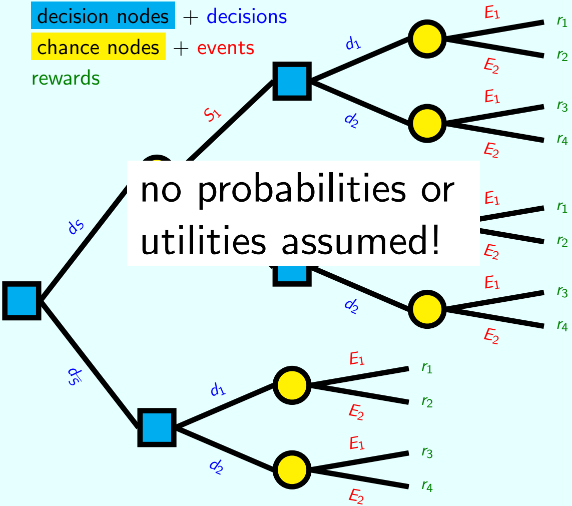

no probabilities or utilities assumed!

Decision Trees: Normal Form Decisions

Definition

A normal form decision fixes at every decision node exactly one decision.

Decision Trees: Normal Form Decisions

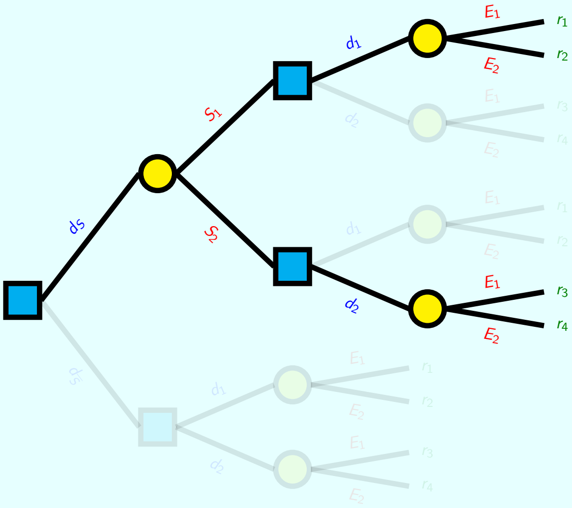

Decision Trees: Gambles

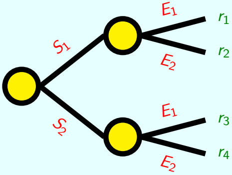

Decision Trees: Gambles

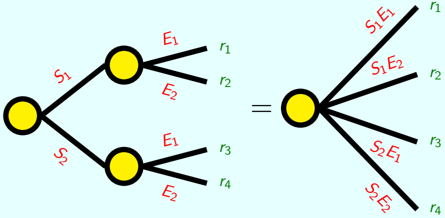

Decision Trees: Gambles

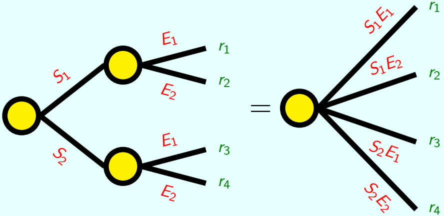

$$= X(\omega) = \begin{cases} r_1 & \text{if } \omega \in S_1E_1 \\ r_2 & \text{if } \omega \in S_1E_2 \\ r_3 & \text{if } \omega \in S_2E_1 \\ r_4 & \text{if } \omega \in S_2E_2 \end{cases}$$

Decision Trees: Gambles

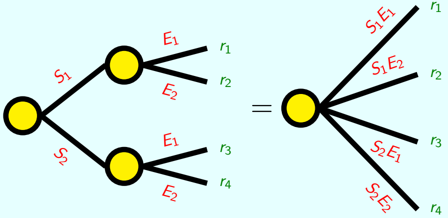

$$= X(\omega) = \begin{cases} r_1 & \text{if } \omega \in S_1 E_1 \\ r_2 & \text{if } \omega \in S_1 E_2 \\ r_3 & \text{if } \omega \in S_2 E_1 \\ r_4 & \text{if } \omega \in S_2 E_2 \end{cases}$$

Observation

Every normal form decision induces a gamble.

Decision Trees: Normal Form Solution

Definition

A normal form solution of a decision tree is a set of these normal form decisions.

Decision Trees: Normal Form Solution

Definition

A normal form solution of a decision tree is a set of these normal form decisions.

for example

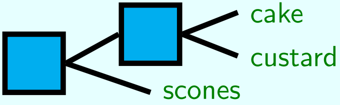

could have as normal form solution

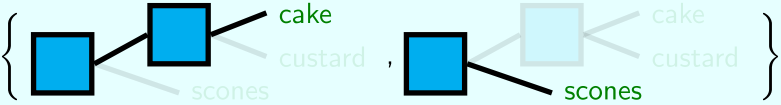

Decision Trees: Normal Form Solution Induced By opt

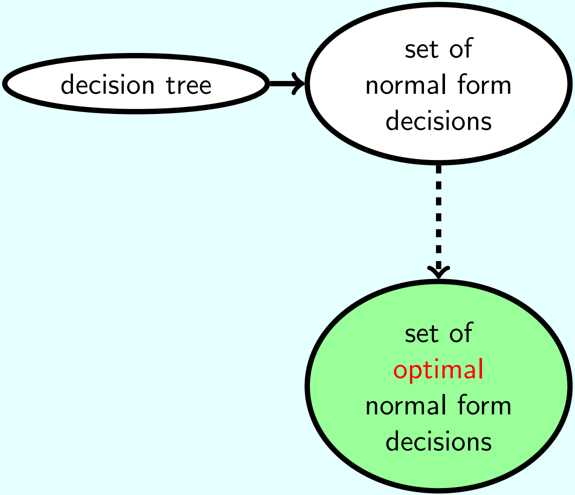

Decision Trees: Normal Form Solution Induced By opt

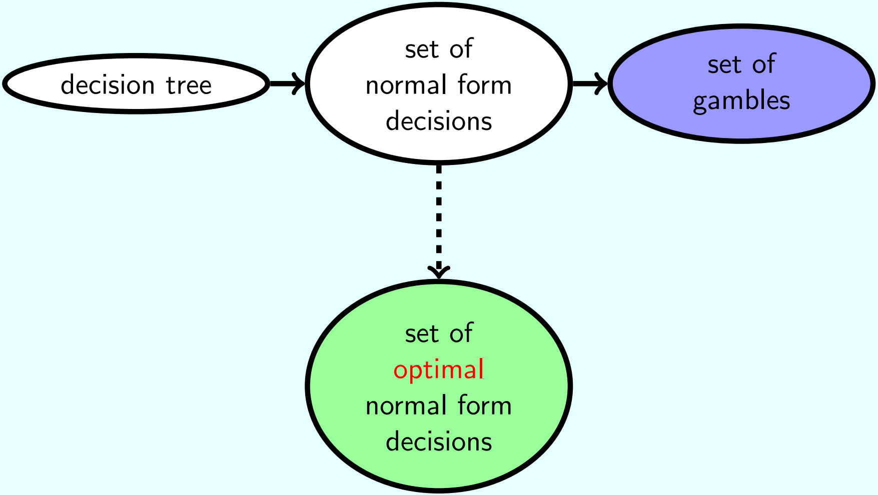

Decision Trees: Normal Form Solution Induced By opt

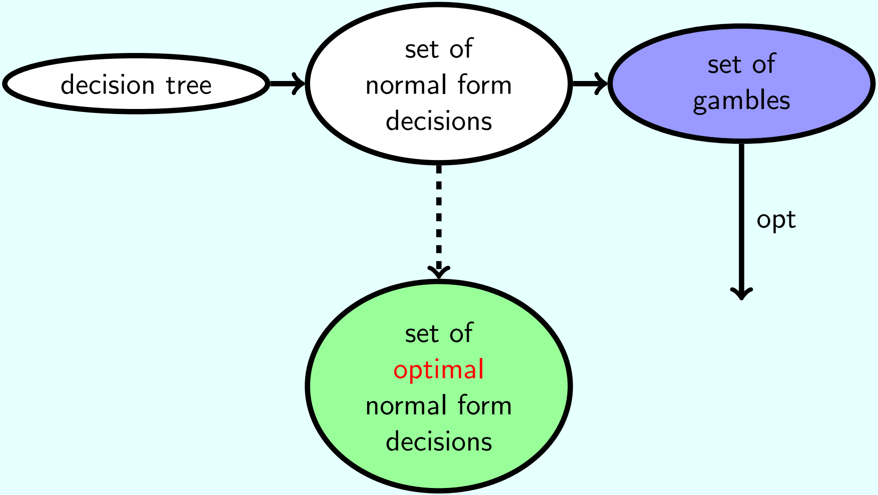

Decision Trees: Normal Form Solution Induced By opt

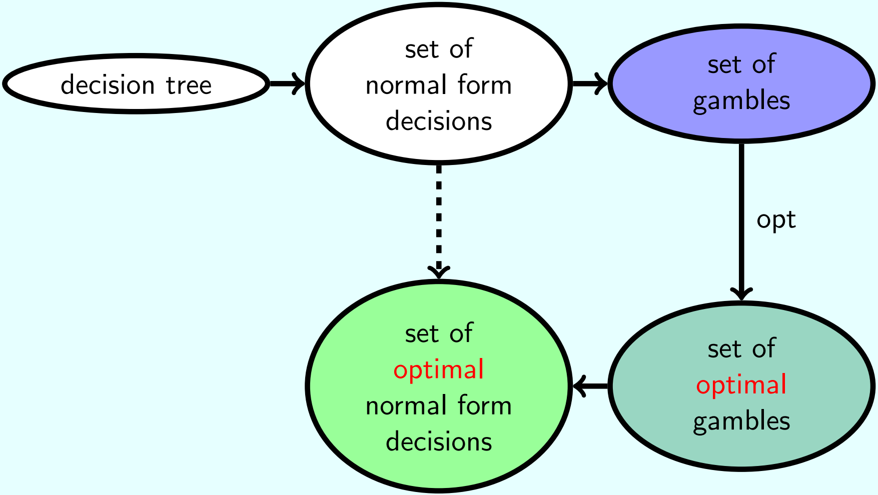

Decision Trees: Normal Form Solution Induced By opt

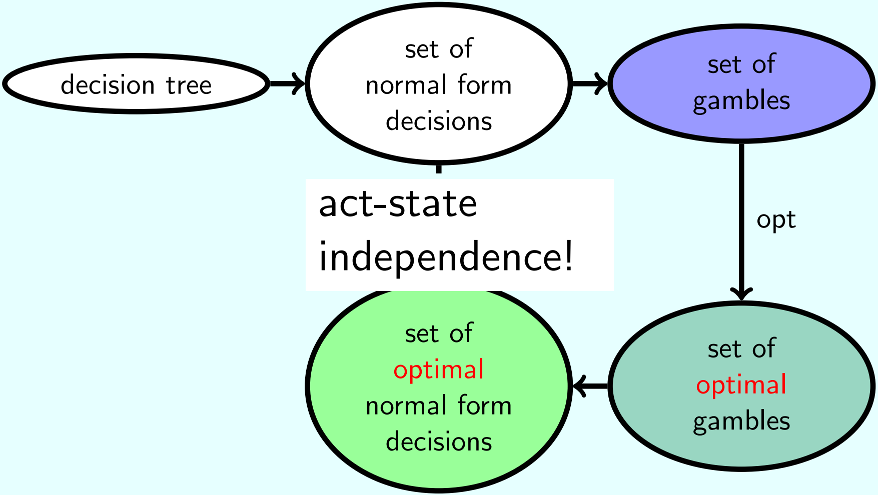

Decision Trees: Normal Form Solution Induced By opt

If we have a choice function opt
then we can in principle solve arbitrary decision trees!

Outline

Factuality: A Counterfactual Example

Suppose opt (

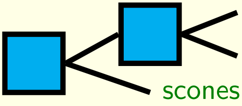

) = {

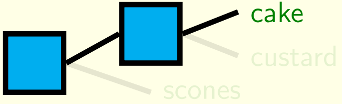

}

Factuality: A Counterfactual Example

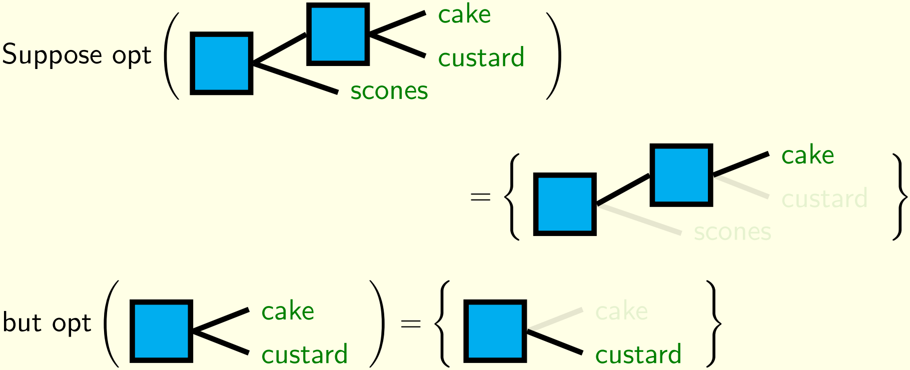

Factuality: A Counterfactual Example

Suppose opt (

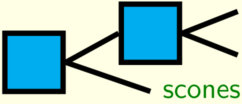

)

= {

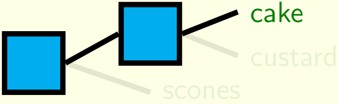

}

but opt (

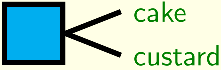

) = {

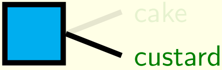

}

The choice between cake and custard
depends on the tree in which the decision is embedded.

Factuality: Definition

Definition

opt is called factual whenever for every decision tree

$$\text{restriction}(\text{opt}(\text{tree})) = \text{opt}(\text{restriction}(\text{tree}))$$

whenever restriction(tree)'s root node is in opt(tree).

In bargaining theory this principle is called subgame perfection.

Factuality Theorem

Theorem

opt is factual if and only if it satisfies:

- Conditioning property. If $\{X, Y\} \subseteq \mathcal{X}$ and $AX = AY$, then
$$X \in \text{opt}(\mathcal{X}|A) \iff Y \in \text{opt}(\mathcal{X}|A).$$
- Intersection property. If $\mathcal{Y} \subseteq \mathcal{X}$ and $\text{opt}(\mathcal{X}|A) \cap \mathcal{Y} \neq \emptyset$, then
$$\text{opt}(\mathcal{Y}|A) = \text{opt}(\mathcal{X}|A) \cap \mathcal{Y}.$$
- Mixture property.
$$\text{opt}(A\mathcal{X} \oplus \overline{A}Z|B) = A\,\text{opt}(\mathcal{X}|A \cap B) \oplus \overline{A}Z.$$

Note: some technical details omitted.

Factuality: No Imprecision

Total Preorder Theorem

The intersection property is equivalent to:

- Total preorder property. For every event $A \neq \emptyset$, there is a total preorder $\succeq_A$ on gambles such that

$$\mathrm{opt}(\mathcal{X}|A) = \max_{\succeq_A}(\mathcal{X})$$

So it is impossible to be at the same time

- factual, and
- optimal with respect a non-total preference ordering (such as for instance a partial preference ordering)

Factuality: What Choice Functions are Factual?

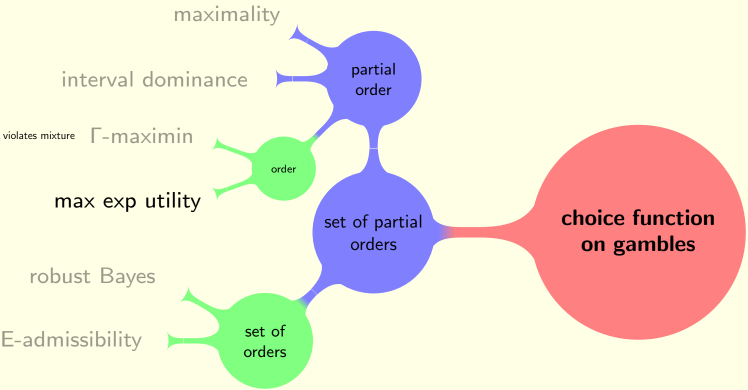

Factuality: What Can Be Done?

- Some types of counterfactuality may not be so bad, for instance those where backward induction still works (such as maximality and E-admissibility).
- Restrict type of decision trees that you are interested in: there are sequential decision processes where factuality can be obtained under substantially weaker assumptions.

Outline

Locality: Markov Decision Processes

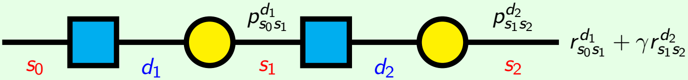

Locality: Markov Decision Processes

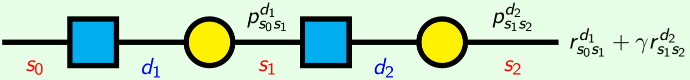

Solution for an n-stage process:

$$V^n_{n+1}(s) = 0 \qquad V^n_{k+1}(s) = \max_d \sum_t p^d_{st}(r^d_{st} + \gamma V^n_k(t))$$

Under act-state independence, $p^d_{st} = p_{st}$, and so:

$$V^n_{n+1}(s) = 0 \qquad V^n_{k+1}(s) = \left(\max_d \sum_t p_{st} r^d_{st}\right) + \gamma \sum_t p_{st} V^n_k(t)$$

Locality: Markov Decision Processes

Observation

Under act-state independence, the solution becomes trivial:

global solution to sequential problem

$$\Longleftrightarrow$$

sequence of solutions to local problems

"locality"

Locality: Markov Decision Processes

Observation

Under act-state independence, the solution becomes trivial:

global solution to sequential problem

$\Longleftrightarrow$

sequence of solutions to local problems

"locality"

Under what conditions on opt **do we have locality?**

Locality: Sequential Decision Process on $S_0, \ldots, S_n$

- states $S_0, \ldots, S_n$: sequence of random variables
- history $h_k = (s_0, \ldots, s_k)$: sequence of observed states
- decisions and local rewards
  - observe $s_0 \in S_0$,
  - pick $d_1 \in D_1$, see $S_1 = s_1$, get $r_1(s_0 d_1 s_1) = r_1(h_0 d_1 s_1)$,
  - pick $d_2 \in D_2$, see $S_2 = s_2$, get $r_2(s_0 s_1 d_2 s_2) = r_2(h_1 d_2 s_2)$,
  - ...
  - pick $d_n \in D_n$, see $S_n = s_n$, get $r_n(h_{n-1} d_n s_n)$.
- $+$ operator on $\mathcal{R}$ with 0: $r + 0 = r$
- combined reward (evaluate from right to left)

$$r_1(h_0 d_1 s_1) + r_2(h_1 d_2 s_2) + \cdots + r_n(h_{n-1} d_n s_n)$$

Locality: Policies

Assume $s_0 d_1 s_1 \dots d_{k-1} s_{k-1}$ has occurred
then a normal form decision, or policy π_k^n, consists of

- a decision $d_k \in D_k$,
- a decision function $d_{k+1}(s_k) \in D_{k+1}$,
- ...
- a decision function $d_n(s_k \dots s_{n-1}) \in D_n$.

Definition

Set of all policies for an n-stage process at the k-th stage:

$$\Pi_k^n = D_k \times D_{k+1}^{S_k} \times \dots \times D_n^{S_k \times \dots \times S_{n-1}}$$

Locality: Gambles and Optimality

Definition

Each policy $\pi_k^n = \Pi_k^n$ incurs a gamble $X_k^n(h_{k-1}, \pi_k^n)$.

Definition

The set of all these gambles is

$$\mathcal{X}_k^n(h_{k-1}) = \{X(h_{k-1}, \pi_k^n) \colon \pi_k^n \in \Pi_k^n\}$$

Definition

Set of all optimal policies is

$$\Pi_k^n(h_{k-1}) = \{\pi_k^n \in \Pi_k^n \colon X_k^n(h_{k-1}, \pi_k^n) \in \mathrm{opt}(\mathcal{X}_k^n(h_{k-1}) | h_{k-1})\}$$

Locality: Definition

Definition

opt is said to satisfy locality on $S_0, \ldots, S_n$ whenever for every sequential decision process on $S_0, \ldots, S_n$ and every $1 \leq k \leq n$:

$$\Pi_k^n(\cdot) = \Pi_k^k(\cdot) \times \Pi_{k+1}^{k+1}(\cdot) \times \cdots \times \Pi_n^n(\cdot)$$

Locality: Definition

Definition

opt is said to satisfy locality on S_0, S_1, S_2 whenever

Locality: Definition

Definition

opt is said to satisfy locality on S_0, S_1, S_2 whenever

opt (

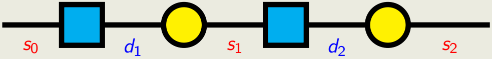

$r_1(s_0 d_1 s_1) + r_2(s_0 s_1 d_2 s_2) \,\Big|\, s_0 \Big)$

Locality: Definition

Definition

opt is said to satisfy locality on S_0, S_1, S_2 whenever

opt (

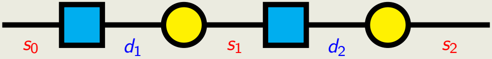

$r_1(s_0 d_1 s_1) + r_2(s_0 s_1 d_2 s_2) \Big| s_0 \Big)$

= opt (

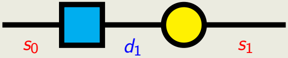

$r_1(s_0 d_1 s_1) \Big| s_0 \Big)$

Locality: Definition

Definition

opt is said to satisfy locality on S_0, S_1, S_2 whenever

opt (

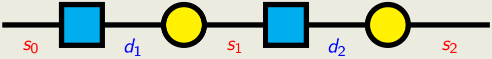

$r_1(s_0 d_1 s_1) + r_2(s_0 s_1 d_2 s_2) \,\Big|\, s_0$)

= opt (

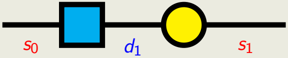

$r_1(s_0 d_1 s_1) \,\Big|\, s_0$)

× opt (

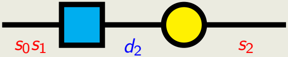

$r_2(s_0 s_1 d_2 s_2) \,\Big|\, s_0 s_1$)

Locality Theorem

Theorem

opt satisfies locality on $S_0, \dots, S_n$ if and only if it satisfies:

- Sequential distributivity. For any $1 \leq k < n$, any value h_{k-1} of H_{k-1}, all finite sets of gambles $\mathcal{X}$ on S_k, all finite sets of gambles $\mathcal{Y}(s_k)$ on F_{k+1} (one such set for each $s_k \in S_k$), and all $X \in \mathcal{X}$ and $Y(s_k) \in \mathcal{Y}(s_k)$:

$$X + \bigoplus_{s_k} E_{s_k} Y(s_k) \in \operatorname{opt}\left(\mathcal{X} + \bigoplus_{s_k} E_{s_k} \mathcal{Y}(s_k) \middle| h_{k-1}\right)$$

$$\iff$$

$$X \in \operatorname{opt}(\mathcal{X}|h_{k-1}) \text{ and } Y(s_k) \in \operatorname{opt}(\mathcal{Y}(s_k)|h_{k-1}s_k) \text{ for all } s_k.$$

Locality: Maximality

Proposition

Maximality with respect to $\underline{P}$ satisfies sequential distributivity (and hence, locality) on $S_0, \ldots, S_n$, if and only if for all $1 \leq k < n$, all $h_{k-1} \in H_{k-1} = S_0 \times \cdots \times S_{k-1}$, all $s_k \in S_k$, and all gambles Z on $F_k = S_k \times \cdots \times S_n$,

$$\underline{P}(E_{s_k}|h_{k-1}) > 0$$

and

$$\underline{P}(Z|h_{k-1}) = \underline{P}(\underline{P}(Z|h_{k-1}S_k)|h_{k-1}).$$

- strictly positive transition probabilities
- marginal extension

Locality: E-admissibility

Proposition

E-admissibility with respect to $\underline{P}$ satisfies sequential distributivity (and hence, locality) on $S_0, \ldots, S_n$, if and only if for all $1 \leq k < n$, all $h_{k-1} \in H_{k-1} = S_0 \times \cdots \times S_{k-1}$, all $s_k \in S_k$, and all gambles Z on $F_k = S_k \times \cdots \times S_n$,

$$\underline{P}(E_{s_k}|h_{k-1}) > 0$$

and

$$\underline{P}(Z|h_{k-1}) = \underline{P}(\underline{P}(Z|h_{k-1}S_k)|h_{k-1}).$$

- strictly positive transition probabilities
- marginal extension

Locality: Γ-maximin

Locality for Γ-maximin fails under these conditions. . .
However:

Locality: Γ-maximin

Locality for Γ-maximin fails under these conditions. . .
However:

Observation

Globally optimal solution could of course still be obtained by non-local means (i.e. backward induction á la Satia and Lave).

Locality: Γ-maximin

Locality for Γ-maximin fails under these conditions. . .
However:

Observation

Globally optimal solution could of course still be obtained by non-local means (i.e. backward induction á la Satia and Lave).

Observation

Any Γ-maximin solution obtained locally will be globally optimal with respect to maximality.

Conclusion and Challenges

- Factuality and locality provide rationality arguments for
  - total preference orderings: say no to imprecision
  - marginal extension
  - strictly positive transition lower probabilities

  (although some of these may be stronger than you'd like!)

Conclusion and Challenges

- Factuality and locality provide rationality arguments for
  - total preference orderings: say no to imprecision
  - marginal extension
  - strictly positive transition lower probabilities

  (although some of these may be stronger than you'd like!)
- Locality can be satisfied for non-total orderings such as
  - maximality
  - E-admissibility

Conclusion and Challenges

- Factuality and locality provide rationality arguments for
  - total preference orderings: say no to imprecision
  - marginal extension
  - strictly positive transition lower probabilities

  (although some of these may be stronger than you'd like!)
- Locality can be satisfied for non-total orderings such as
  - maximality
  - E-admissibility
- Local Γ-maximin
  - may *not* be global Γ-maximin!
  - but it is not entirely stupid either

Conclusion and Challenges

- Factuality and locality provide rationality arguments for
  - total preference orderings: say no to imprecision
  - marginal extension
  - strictly positive transition lower probabilities

  (although some of these may be stronger than you'd like!)
- Locality can be satisfied for non-total orderings such as
  - maximality
  - E-admissibility
- Local Γ-maximin
  - may *not* be global Γ-maximin!
  - but it is not entirely stupid either

Challenges:

- is the definition of factuality too strong?
- how to accomodate **act-state dependence**?!

Thanks for your attention!

questions? comments? discussion?