Sequential Decision Making and Imprecise Probabilities

Review Talk

Matthias C. M. Troffaes

Department of Mathematical Sciences
Durham University

22 June, 2010

Outline

Outline

Static Decision Problems

- we must make a single decision, at some point in time
- something uncertain happens, depending on decision

(example taken from Raiffa, 1968)

- an oil wildcatter must decide whether or not to drill for oil
- yield depends on the uncertain richness of the site

game-theoretic interpretation

- Skeptic decides whether to drill or not
- World decides the richness of the site
  Skeptic does not know World's decision in advance

Sequential Decision Problems

- we must make multiple decision, at different points in time
- uncertain things happen, depending on decisions
- we can learn between decision points

(example taken from Raiffa, 1968)

- oil wildcatter must decide whether or not to drill for oil
- yield depends on uncertain richness of the site
- subject may pay to test the seismic structure of the site first

game-theoretic interpretation

- Skeptic decides whether to pay for a test or not
- World decides the outcome of the test
- Skeptic decides whether to drill or not
- World decides the richness of the site

Decision Trees

- conveniently represent all available strategies of Skeptic at once
- are equivalent to a set of event trees

Event Tree

- all possible paths for World
- Skeptic's rewards for a particular decision strategy of Skeptic

Decision Tree

- all possible paths for World and Skeptic jointly
- Skeptic's rewards for any possible path
- note that Skeptic can affect World's move space and vice versa

Decision Trees

Oil Wildcatter Example

Skeptic:	decision nodes + decisions
World:	chance nodes + events
Gamble:	rewards

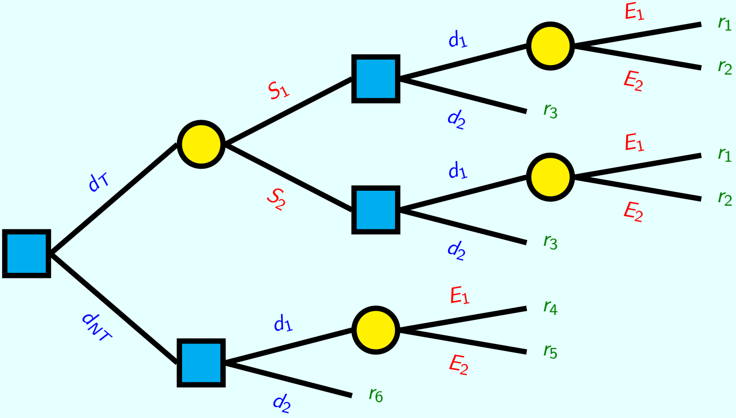

Decision Trees

Oil Wildcatter Example

Skeptic: decision nodes + decisions

World: chance nodes + events

Gamble: rewards

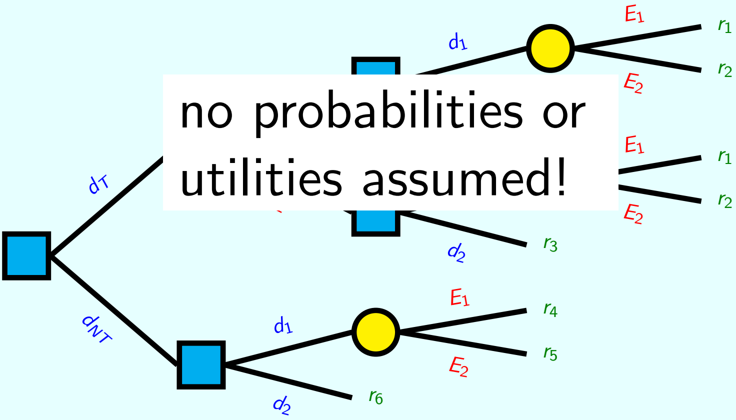

Choice Functions

- subset of good decisions?

(Conditional) (Static) Choice Function

- maps any set of decisions $\mathcal{X}$ to a subset of that set
- may depend on hypothetical information A

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

- example:

$$\mathrm{opt}(\{\text{drill site A}, \text{drill site B}, \text{do not drill}\}) = \{\text{drill site A}, \text{do not drill}\}$$

- you could imagine Forecaster specifying opt

Choice Functions

- should choice always be unique?
- if not, multi-tiered choice functions?

$$\mathrm{opt}_n(\dots\mathrm{opt}_2(\mathrm{opt}_1(\mathcal{X}))\dots)$$

- models for sequential choice?

Sequential Choice

reduce sequential choice to conditional static choice

Normal Form

- list all decision strategies for Skeptic (i.e. all possible event trees)
- choose among them as if each decision strategy is a single decision

Extensive Form

- at each stage, Skeptic considers his future self has part of World
- then, at each stage, he only needs to consider a static problem
- (usually employed with backward induction)

in general, normal and extensive forms may disagree!

Outline

Lower Previsions: Notation and Assumptions

- finite sample space Ω
  (**set of all paths of an event tree**)
  example:
  $$\Omega = \{\text{no oil}, \text{some oil}, \text{plenty oil}\}$$
- gamble X = a random quantity on Ω interpreted as reward
  ($\simeq$ **game-theoretic strategy for Skeptic, process**)
  example:
  $$X(\text{no oil}) = \text{-£7,000}$$
  $$X(\text{some oil}) = \text{£5,000}$$
  $$X(\text{plenty oil}) = \text{£20,000}$$
- event A = a subset of Ω
  example: sounding test indicates at least some oil
  $$A = \{\text{some oil}, \text{plenty oil}\}$$

Lower Previsions: Notation and Assumptions

- credal set $\mathcal{M}$ = set of probability mass functions
  (for convenience, assume $p(A) > 0$ for all $A \neq \emptyset$ and all p in $\mathcal{M}$)
- expectation of X w.r.t. some p in $\mathcal{M}$

$$E_p(X|A) = \frac{\sum_{\omega \in A} p(\omega)X(\omega)}{\sum_{\omega \in A} p(\omega)}$$

- lower prevision = lower expectation w.r.t. $\mathcal{M}$

$$\underline{P}(X|A) = \inf_{p \in \mathcal{M}} E_p(X|A) \qquad \overline{P}(X|A) = \sup_{p \in \mathcal{M}} E_p(X|A)$$

- betting interpretation

Outline

Choice Functions for Lower Previsions

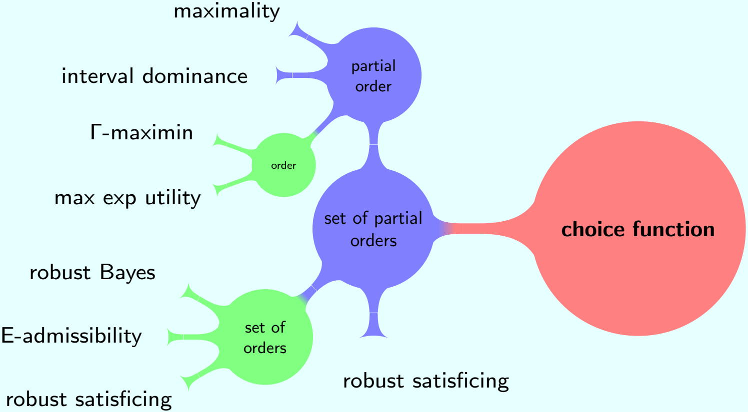

Maximality

(Condorcet 1785 [6], Sen 1977 [29], Walley 1991 [36])

- a lower prevision determines a partial ordering between gambles

$$X >_{\underline{P}|A} Y \text{ whenever } \underline{P}(X - Y|A) > 0$$

(willing to pay a small amount in order to trade X for Y)
($X - Y + \epsilon$ is an acceptable transaction for some $\epsilon > 0$)

Choose any gamble X in $\mathcal{X}$ which is maximal (i.e. undominated) with respect to $>_{\underline{P}|A}$.

$$\text{opt}(\mathcal{X}|A) = \max_{>_{\underline{P}|A}} \mathcal{X}$$

E-admissibility / Robust Bayes

(Pascal 1662 [25], Levi 1980 [22], Berger 1984 [5])

Choose any gamble X in $\mathcal{X}$ whose expectation $E_p(X|A)$ is maximal for at least one probability measure p in the credal set $\mathcal{M}$.

$$\mathrm{opt}(\mathcal{X}|A) = \bigcup_{p \in \mathcal{M}} \arg\max_{X \in \mathcal{X}} E_p(X|A)$$

- employed by Pascal's wager, among other criteria (Pascal, 1662)
- interpretation in terms of maximality if $\mathcal{M}$ is convex:
  X is E-admissible in $\mathcal{X}$ if and only if
  it is not $>_{\underline{P}}$-dominated by any mixture of gambles in $\mathcal{X}$

Γ-maximin

(Wald 1945 [35], Gilboa & Schmeidler 1989 [10])

Choose any gamble X in $\mathcal{X}$ whose lower expectation $\underline{P}(X|A)$ is maximal.

$$\mathrm{opt}(\mathcal{X}|A) = \arg\max_{X \in \mathcal{X}} \underline{P}(X|A)$$

- every Γ-maximin gamble is also maximal (Walley 1991 [36])
- a Γ-maximin gamble is not always E-admissible (Seidenfeld 2004 [28])
- only worst case is considered; overly conservative?

Variants:

- Γ-maximax (Satia and Lave 1973 [26])
- Hurwicz criterion (Hurwicz 1950)

Interval Dominance

(Satia and Lave 1973 [26], Kyburg 1983 [21], *many* others)

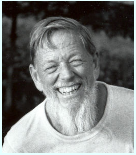

Choose any gamble X in $\mathcal{X}$ whose upper expectation $\underline{P}(X|A)$ is not dominated by the lower expectation of any other gamble.

$$\mathrm{opt}(\mathcal{X}|A) = \{X \in \mathcal{X} \colon \overline{P}(X|A) \geq \max_{Y \in \mathcal{X}} \underline{P}(Y|A)\}$$

- every maximal gamble is interval dominant (Troffaes 2007 [32])
- an interval dominant gamble can be $>_{\underline{P}|A}$-dominated!
- interval dominance is computationally very fast
  can help eliminating some options quickly

Outline

Satisficing vs. Optimizing

(Simon 1955 [30], Gigerenzer & Goldstein 1996 [9], Ben-Haim 2001 [4])

- often, we *do not have to*, or we simply *cannot*, optimize
- simply need to satisfy a critical requirement E_c

robust satisficing = maximize the satisficing event

$$\{\omega \in \Omega \colon X(\omega) \geq E_c\}$$

- using a choice function
- or by measuring $C(X)$ in a reasonable way (e.g. Ben-Haim 2001 [4])

- trade-off between robustness and minimal requirements
- need to think carefully about E_c

Outline

Updating vs. Conditioning

(Goldstein 1985 [11] [12])

Conditioning

(some canonical system for determining. . .)
Skeptic's hypothetical beliefs about World
after having followed a particular hypothetical path

Updating

Skeptic's actual beliefs about World
after having followed a particular actual path

Skeptic may have actual beliefs
about his hypothetical beliefs after following a particular hypothetical path

- how should this be part of normative decision theory?

Outline

Act-State Dependence

(James Joyce 1999 [19], Dréze 2000 [8])

Example

I, World, will raise either my left right hand, or my right hand. Before I do so, you, Skeptic, have to bet on which hand I will raise. You win £10 if your guess is correct, and you lose £10 otherwise.

- Skeptic's moves influence World's moves
- happens commonly (for example, Markov decision processes)
- hard to model under imprecision except in special cases (Troffaes [31])

Outline

Sequential Issues

Many important issues when using normal and extensive form ideas under imprecision!

- Seidenfeld 1988 [27], 2004 [28]
- Jaffray 1999 [18]
- Augustin 2001 [3]
- De Cooman & Troffaes 2005 [7]

For this review talk: only the most important highlights.

Subtree Perfectness: Example

Suppose opt 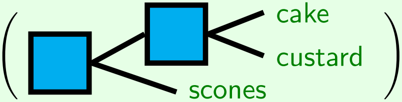

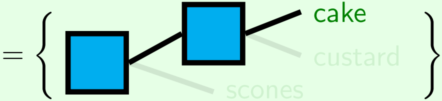

Subtree Perfectness: Example

Suppose opt (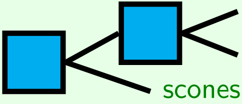

)

= { 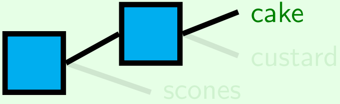

}

but opt (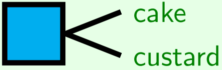

) = { 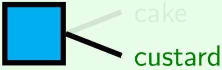

}

Subtree Perfectness: Example

Suppose opt (
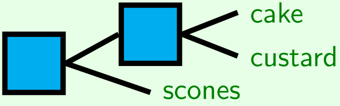

)

= {
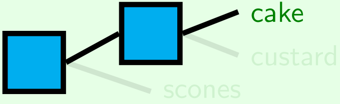

}

but opt (
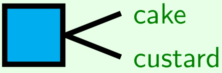

) = {
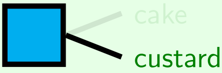

}

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

Subtree Perfectness: Definition

Definition

opt is called subtree perfect 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 game theory this principle is called subgame perfection.

Subtree Perfectness Theorem

(early discussion 1950's: Houthakker [15], Ville [34], Arrow [2] [1])
(partial results: Hammond '88 [13], Machina '89 [23], McClennen '90 [24])
(full result: Huntley & Troffaes 2009 [17])

Theorem

opt is subtree perfect if and only if it satisfies:

- Conditioning property. If $\{X, Y\} \subseteq \mathcal{X}$ and $AX = AY$, then

$$X \in \mathrm{opt}(\mathcal{X}|A) \iff Y \in \mathrm{opt}(\mathcal{X}|A).$$

- Intersection property. If $\mathcal{Y} \subseteq \mathcal{X}$ and $\mathrm{opt}(\mathcal{X}|A) \cap \mathcal{Y} \neq \emptyset$, then

$$\mathrm{opt}(\mathcal{Y}|A) = \mathrm{opt}(\mathcal{X}|A) \cap \mathcal{Y}.$$

- Mixture property.

$$\mathrm{opt}(A\mathcal{X} \oplus A^c Z|B) = A\,\mathrm{opt}(\mathcal{X}|A \cap B) \oplus A^c Z.$$

Note: some technical details omitted.

Subtree Perfectness: 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

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

So it is impossible to be at the same time

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

Subtree Perfectness: What Choice Functions are Subtree Perfect?

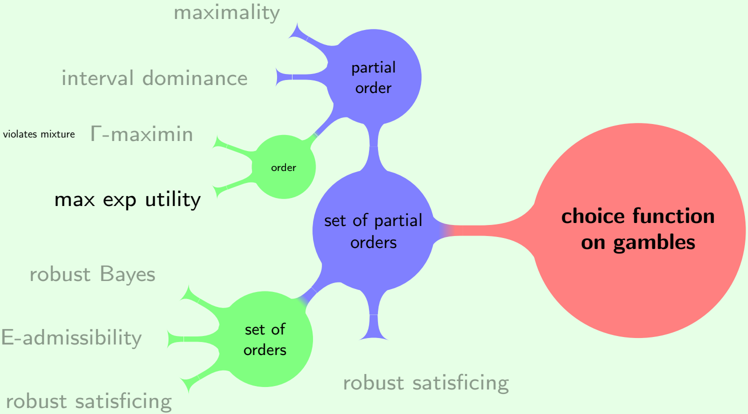

Subtree Perfectness: What Can Be Done?

- Some types of subtree imprefectness 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 subtree perfectness can be obtained under substantially weaker assumptions.

Outline

Locality: Markov Decision Processes

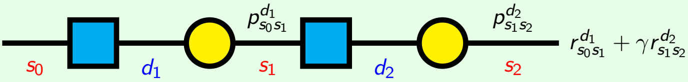

Solution for an n-stage process:

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

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

$$\iff$$

sequence of solutions to local problems

"locality"

Locality: Definition

in the definition below, sequential decision process =

- Skeptic and World move alternatively
- perfect information
- rewards are added up throughout the game

(i.e. the usual game-theoretic process, not an arbitrary decision tree)

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

opt (

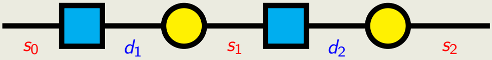

$r_1(s_0 d_1 s_1) + r_2(s_0 s_1 d_2 s_2) \,\Big|\, s_0 \Big)$

= opt (

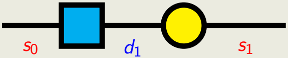

$r_1(s_0 d_1 s_1) \,\Big|\, s_0 \Big)$

× opt (

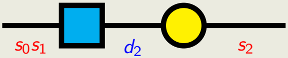

$r_2(s_0 s_1 d_2 s_2) \,\Big|\, s_0 s_1 \Big)$

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

Locality Theorem

(Troffaes & Huntley & Shirota [33])

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 / backward concatentation (future conditioned on the past!)

Locality: E-admissibility

Proposition

E-admissibility with respect to $\underline{P}$ satisfies sequential distributivity (and hence, locality) on $S_0, \dots, S_n$, 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 / backward concatentation (future conditioned on the past!)

Locality: Γ-maximin and Interval Dominance

Locality fails for Γ-maximin and interval dominance...
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.

Outline

Backward Induction

(again investigating arbitrary decision trees)

idea: use solutions of subtrees to eliminate options in the full tree

For general choice functions several algorithms have been proposed:

- Seidenfeld 1988 [27] [28] (extensive form)
- Harmanec 1999 [14] (extensive form)
- De Cooman & Troffaes 2005 [7] (normal form)
- Kikuti et. al 2005 [20] (apparently, normal form)
- Huntley & Troffaes 2008 [16] (normal form)

Normal Form Backward Induction: Example

(Huntley & Troffaes, adapted from Kikuti et al.)

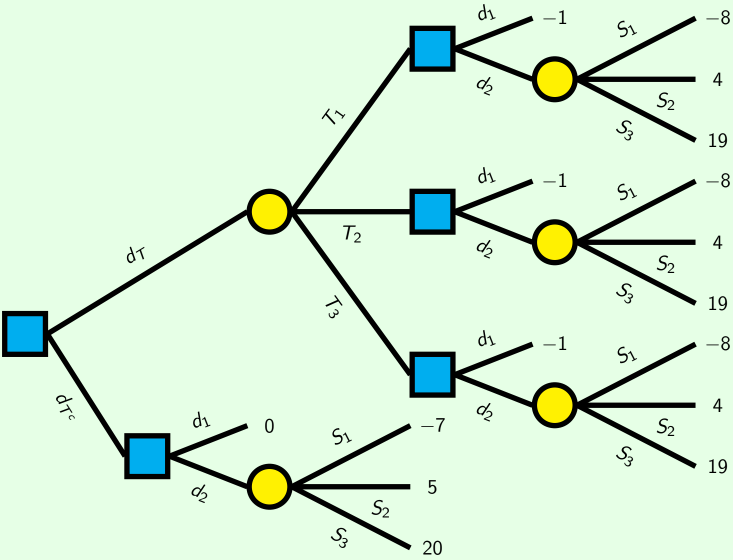

Normal Form Backward Induction: Example

(Huntley & Troffaes, adapted from Kikuti et al.)

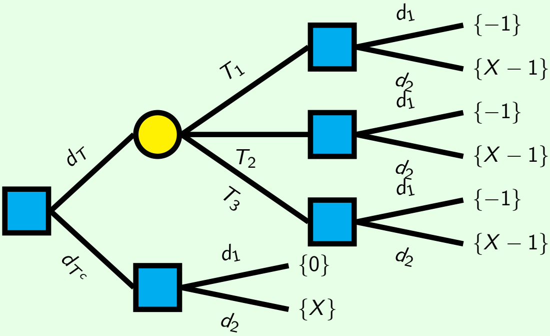

Normal Form Backward Induction: Example

(Huntley & Troffaes, adapted from Kikuti et al.)

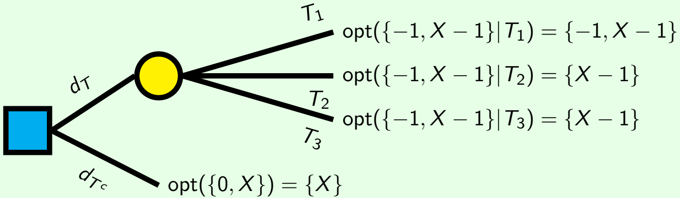

Normal Form Backward Induction: Example

(Huntley & Troffaes, adapted from Kikuti et al.)

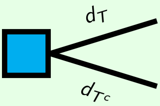

$\text{opt}(\{T_1(-1)+T_2(X-1)+T_3(X-1), T_1(X-1)+T_2(X-1)+T_3(X-1)\})$
$=\{(T_2+T_3)X-1, X-1\}$

$\{X\}$

Normal Form Backward Induction: Example

(Huntley & Troffaes, adapted from Kikuti et al.)

$$\mathrm{opt}(\{(T_2 + T_3)X - 1, X - 1, X\}) = \{X\}$$

Necessary and Sufficient Conditions

(Huntley & Troffaes 2008 [16])

Theorem

Backward induction works with opt if and only if it satisfies:

- Backward conditioning property. If $AX = AY$ and $\{X, Y\} \subseteq \mathcal{X}$, then $X \in \text{opt}(\mathcal{X}|A) \iff Y \in \text{opt}(\mathcal{X}|A)$ (subject to some technicalities)
- Path independence.

$$\text{opt}\left(\bigcup_{i=1}^{n} \mathcal{X}_i \middle| A\right) = \text{opt}\left(\bigcup_{i=1}^{n} \text{opt}(\mathcal{X}_i|A) \middle| A\right)$$

- Backward mixture property.

$$\text{opt}\left(\{AX + A^c Z \colon X \in \mathcal{X}\}|B\right) \subseteq A\,\text{opt}(\mathcal{X}|A \cap B) \oplus A^c Z$$

Note: some technical details omitted.

Backward Induction: What Choice Functions Work?

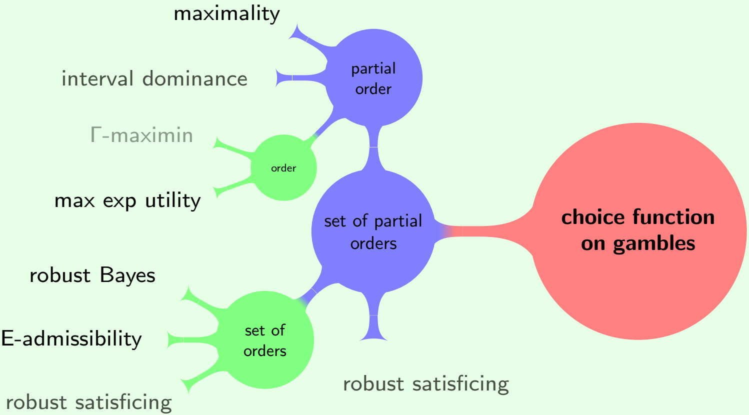

Normal Form Backward Induction vs. Subtree Perfectness vs. Locality

- locality $\implies$ subtree perfectness
  (but not the opposite implication)
- subtree perfectness $\implies$ normal form backward induction works
  (but not the opposite implication)
- subtree perfectness $\iff$ (normal form $=$ extensive form)

Conclusion

- the importance of imprecision in decision making has been recognized for a very long time (as early as 1662)
- sequentially, not everything “just works”
- subtree perfectness seems possible for at least a restricted (yet not completely useless) class of decision problems
- game-theoretic thinking could potentially be very valuable in reasoning about sequential decision making
(act-state dependence, satisficing)

References I

K. Arrow.
Exposition of the theory of choice under uncertainty.
Synthese, 16(3–4):253–269, Dec 1966.

Kenneth J. Arrow.
Rational choice functions and orderings.
Economica, 26(102):121–127, May 1959.

T. Augustin.
On decision making under ambiguous prior and sampling information.
In G. de Cooman, T.L. Fine, S. Moral, and T. Seidenfeld, editors, *ISIPTA 01: Proceedings of the Second International Symposium on Imprecise Probabilities*, 2001.

Yakov Ben-Haim.
Info-Gap Decision Theory: Decisions Under Severe Uncertainty.
Academic Press, 2001.

References II

James O. Berger.
The robust Bayesian viewpoint.
In J. B. Kadane, editor, *Robustness of Bayesian Analyses*, pages 63–144. Elsevier Science, Amsterdam, 1984.

Marquis de Condorcet.
Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix.
L'Imprimerie Royale, Paris, 1785.

Gert de Cooman and Matthias C. M. Troffaes.
Dynamic programming for deterministic discrete-time systems with uncertain gain.
International Journal of Approximate Reasoning, 39(2–3):257–278, Jun 2005.

References III

Jacques H. Dréze and Aldo Rustichini.
State-dependent utility and decision theory.
CORE Discussion Papers 2000007, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), February 2000.

Gert Gigerenzer and Daniel G. Goldstein.
Reasoning the fast and frugal way: Models of bounded rationality.
Psychological Review, 103(4):650–669, 1996.

Itzhak Gilboa and David Schmeidler.
Maxmin expected utility with non-unique prior.
Journal of Mathematical Economics, 18(2):141–153, 1989.

Michael Goldstein.
Temporal coherence.
In J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, editors, *Bayesian Statistics 2*. 1985.

References IV

Michael Goldstein and David A. Wooff.
Bayes Linear Statistics: Theory and Methods.
Wiley, Chichester, 2007.

P. Hammond.
Consequentialist foundations for expected utility.
Theory and Decision, 25(1):25–78, Jul 1988.

David Harmanec.
A generalization of the concept of Markov decision process to imprecise probabilities.
pages 175–182, Ghent, 1999. Imprecise Probabilities Project.

H. S. Houthakker.
Revealed preference and the utility function.
Economica, 17(66):159–174, May 1950.

References V

Nathan Huntley and Matthias C. M. Troffaes.
An efficient normal form solution to decision trees with lower previsions.
In Didier Dubois, M. Asunción Lubiano, Henri Prade, María Ángeles Gil, Przemyslaw Grzegorzewski, and Olgierd Hryniewicz, editors, *Soft Methods for Handling Variability and Imprecision*, Advances in Soft Computing, pages 419–426. Springer, Sep 2008.

Nathan Huntley and Matthias C. M. Troffaes.
Characterizing factuality in normal form sequential decision making.
In Thomas Augustin, Frank P. A. Coolen, Serafín Moral, and Matthias C. M. Troffaes, editors, *ISIPTA'09: Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications*, pages 239–248, Durham, UK, July 2009.

References VI

J. Jaffray.
Rational decision making with imprecise probabilities.
In *1st International Symposium on Imprecise Probabilities and Their Applications*, 1999.

James Joyce.
The Foundations of Causal Decision Theory.
Cambridge University Press, 1999.

Daniel Kikuti, Fabio G. Cozman, and Cassio P. de Campos.
Partially ordered preferences in decision trees: Computing strategies with imprecision in probabilities.
In Ronen Brafman and Ulrich Junker, editors, *Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, pages 118–123, 2005.

References VII

H. E. Kyburg.
Rational belief.
Technical report, University of Rochester, 1983.

Isaac Levi.
The Enterprise of Knowledge. An Essay on Knowledge, Credal Probability, and Chance.
MIT Press, Cambridge, 1980.

M. J. Machina.
Dynamic consistency and non-expected utility models of choice under uncertainty.
Journal of Economic Literature, 27:1622–1688, 1989.

E. F. McClennen.
Rationality and Dynamic Choice: Foundational Explorations.
Cambridge University Press, 1990.

References VIII

Blaise Pascal.
Pensées.
Maxi-Livres, Paris, 2001.
Unfinished work, published posthumously from collected fragments. First incomplete edition: Port-Royal, 1670. First complete reproduction: Michaut, Basle, 1896.

Jay K. Satia and Jr. Roy E. Lave.
Markovian decision processes with uncertain transition probabilities.
Operations Research, 21(3):728–740, 1973.

T. Seidenfeld.
Decision theory without 'independence' or without 'ordering': What is the difference?
Economics and Philosophy, 4:267–290, 1988.

References IX

Teddy Seidenfeld.
A contrast between two decision rules for use with (convex) sets of probabilities: Gamma-maximin versus E-admissibility.
Synthese, 140(1–2):69–88, May 2004.

Amartya Sen.
Social choice theory: A re-examination.
Econometrica, 45(1):53–89, January 1977.

Herbert A. Simon.
A behavioral model of rational choice.
The Quarterly Journal of Economics, 69(1):99–118, February 1955.

References X

Matthias C. M. Troffaes.
Learning and optimal control of imprecise Markov decision processes by dynamic programming using the imprecise Dirichlet model.
In Miguel Lopéz-Díaz, María Á. Gil, Przemyslaw Grzegorzewski, Olgierd Hyrniewicz, and Jonathan Lawry, editors, *Soft Methodology and Random Information Systems*, pages 141–148, Berlin, 2004. Springer.

Matthias C. M. Troffaes.
Decision making under uncertainty using imprecise probabilities.
International Journal of Approximate Reasoning, 45(1):17–29, May 2007.

Matthias C. M. Troffaes, Nathan Huntley, and Ricardo Shirota Filho.
Sequential decision processes under act-state independence with arbitrary choice functions.
Accepted for IPMU'2010.

References XI

Jean Ville and P. K. Newman.
The existence-conditions of a total utility function.
The Review of Economic Studies, 19(2):123–128, 1951.

Abraham Wald.
Statistical decision functions which minimize the maximum risk.
The Annals of Mathematics, 46(2):265–280, 1945.

Peter Walley.
Statistical Reasoning with Imprecise Probabilities.
Chapman and Hall, London, 1991.