Reflexive Rationality

Wolfgang Spohn

Frege Lectures

University of Tartu, June 25 – 27, 2013

Session 1: Standard Decision Models

The Basic Decision Model (in Normal Form)

The decision matrix:

	s_1	s_2	…	s_n
a_1	c_{11}	c_{12}	…	c_{1n}
a_2	c_{21}	c_{22}	…	c_{2n}
…	…	…	…	…
a_m	c_{m1}	c_{m2}	…	c_{mn}

Possible actions or options: $a_1, \ldots, a_m$
Possible states of the world: $s_1, \ldots, s_n$ with probabilities $p(s_j)$
Possible consequences: (c_{ij}) ($i = 1,\ldots,m$, $j = 1,\ldots,n$) with utilities $u(c_{ij})$
Decision rule: Maximize expected utility, i.e. choose an a_i for which

$EU(a_i) = \sum_j u(c_{ij}) \times p(s_j)$ is maximal.

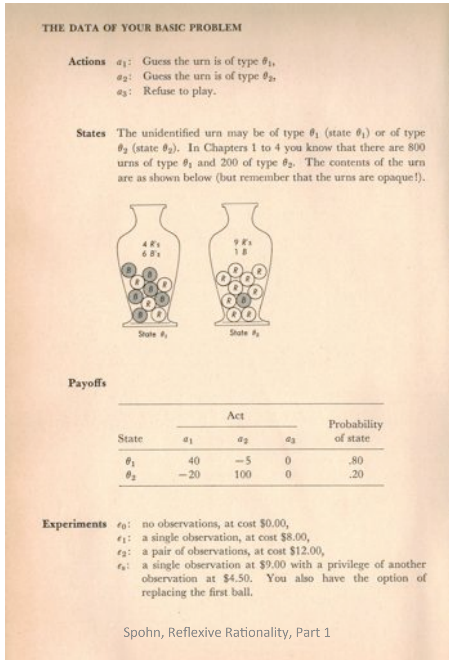

THE DATA OF YOUR BASIC PROBLEM

Actions a_1: Guess the urn is of type θ_1,
a_2: Guess the urn is of type θ_2,
a_3: Refuse to play.

States The unidentified urn may be of type θ_1 (state θ_1) or of type θ_2 (state θ_2). In Chapters 1 to 4 you know that there are 800 urns of type θ_1 and 200 of type θ_2. The contents of the urn are as shown below (but remember that the urns are opaque!).

Payoffs

State	Act a_1	a_2	a_3	Probability of state
θ_1	40	−5	0	.80
θ_2	−20	100	0	.20

Experiments e_0: no observations, at cost $0.00,
e_1: a single observation, at cost $8.00,
e_2: a pair of observations, at cost $12.00,
e_s: a single observation at $9.00 with a privilege of another observation at $4.50. You also have the option of replacing the first ball.

The following slides are from: H. Raiffa, *Decision Analysis*, Addison-Wesley 1970

Decision Problems in Extensive Form and the Roll-back Analysis

- Construct the decision tree for your decision problem with decision nodes ☐ , chance nodes ○ , and end nodes • .
- (Attention: What's the order in the tree? A temporal, causal, or epistemic order?)
- Attach utilities to the end nodes.
- Attach probabilities to all the branches starting at a chance node; the probabilities must add to 1.
- Perform the roll-back analysis on the tree by
  – calculating the expected utility of each chance node and
  – maximizing expected utility at each decision node,
  – till you reach the origin of the tree.
- Thereby you have solved your decision problem.

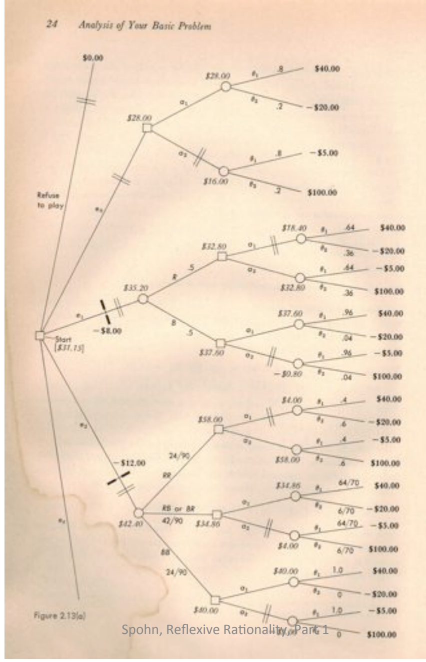

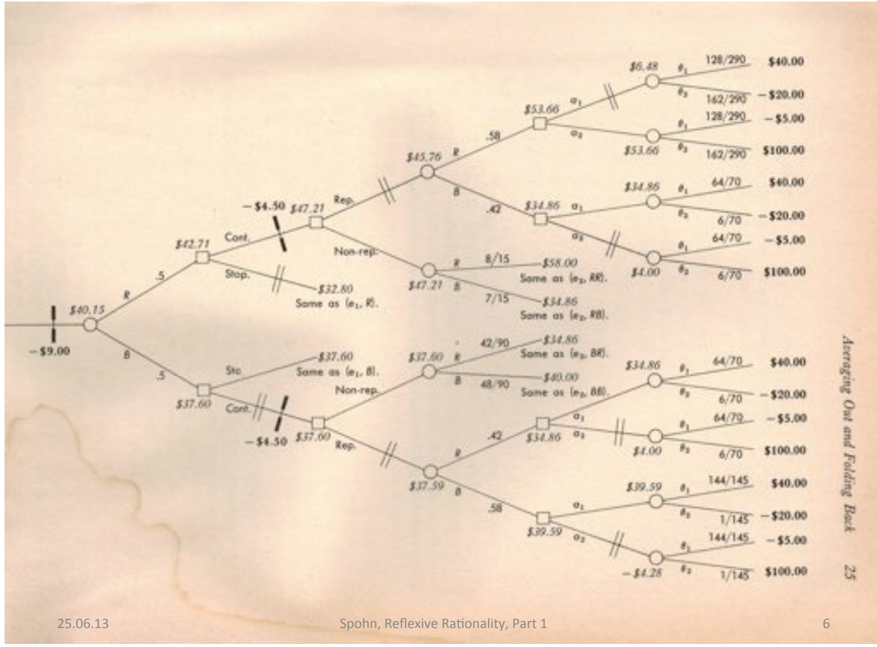

Strategies: From the Extensive Form Back to the Normal Form

Strategies are sub-trees of the decision tree which start at the origin and take only one branch at each decision node they reach and all branches at each chance node they reach.

For instance σ_5:

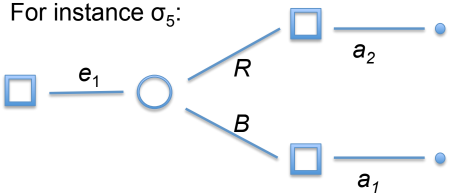

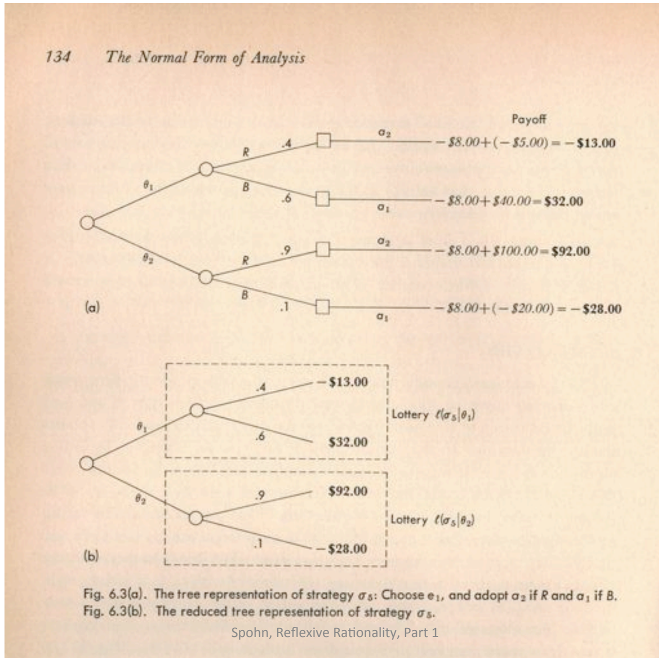

Strategies 135

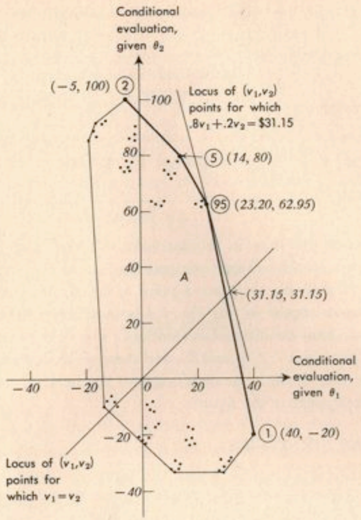

Fig. 6.4 The joint conditional evaluation of all possible strategies for an EMV'er.

In this way, by conceiving options not as single actions, but as entire strategies, we reduce the extensive form of decision problems to the normal form:

	θ_1	θ_2
σ_0	0	0
…	…	…
σ_5	14	80
…	…	…
σ_{114}		

Act-Independent States of the World versus Act-Dependent Consequences

	Disarms	Doesn't disarm
Disarm	Peace	Surrender
Don't disarm	Victory	War

Hence, a more general representation of decision problems is this:
Each option a_i has a set of possible consequences c_{ij} ($j = 1,\dots,n$).
Each possible consequence c_{ij} has a utility $u(c_{ij})$ and a conditional probability $p(c_{ij} \mid a_i)$.
Then the decision rule is to maximize conditional expected utility, i.e., to choose an option for which $CEU(a_i) = \sum_j u(c_{ij}) \times p(c_{ij} \mid a_i)$ is maximal.
(**Attention**: It is not obvious how the reduction of the extensive form to the normal form works then, because it it is not obvious what a probability $p(c_{ij} \mid \sigma_i)$ of a consequence conditional on a strategy is. However, the problem can be solved.)

A Third Representation: Influence Diagrams and Decision Graphs

- There is another representation of decision situations, which I prefer, though it is rarely found in textbooks.
- The normal form (= decision matrix) hides the complexity of decision situations in the possibly very complex options (which may be entire strategies) and in the possibly very complex states of the world (which may be long conjunctions of singular states of affairs).
- Conversely, the extensive form (= decision tree) can be overly complex, since one and the same action or state of affairs is often multiply represented in many branches of the tree.
- Moreover, both forms do not reveal the causal structure of the decision situation. The matrix doesn't do so, anyway, and the tree doesn't, either (because its branches are not causally ordered).
- Finally, the normal form may be causally overdemanding (by requiring act-independent states of the world).

A Third Representation: Influence Diagrams and Decision Graphs

- I prefer the language of variables (in the sense of random variables as understood in probability theory).
- According to it, we deal with a set *U* of variables. Each variable *X* can be simply characterized by the set of its values *x*. The set *W* of possible worlds, of possible evolutions of the decision situation, is represented by $\times U$.
- There is special subset *H* of *U* containing the action variables (standing for possible actions, not strategies). The variables in *U* – *H* may be called chance or occurrence variables.
- The variables are arranged in a causal graph (representing the direct causal dependencies among them – see exemplifications).

A Third Representation: Influence Diagrams and Decision Graphs

- The probability function *p* has to respect the causal graph, i.e., the graph and the probability function together form (something like) a Bayesian net.
- More specifically, *p* provides probabilities $p(x \mid v)$ of the values *x* of each chance variable *X* conditional on each value *v* of the set *V* of the parents or causal predecessors of *X*, i.e., *p* provides only conditional probabilities of chance events and no probabilities of actions. (Concerning probabilities of actions see below the discussion of the difference between influence diagrams and decision graphs.)
- Finally, the utility function *u* is defined on $W = \times U$, the set of possible worlds.

Influence Diagram

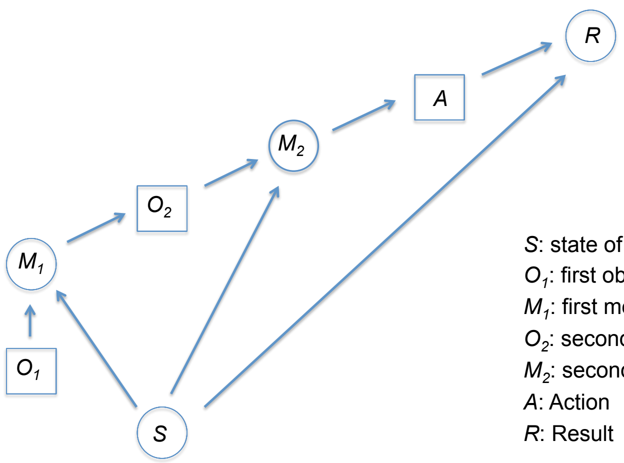

S: state of the world
O_1: first observation
M_1: first measurement
O_2: second observation
M_2: second measurement
A: Action
R: Result

Decision Graph

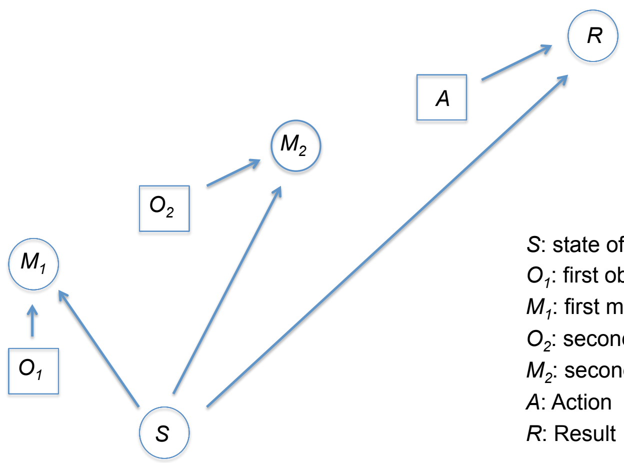

S: state of the world
O_1: first observation
M_1: first measurement
O_2: second observation
M_2: second measurement
A: Action
R: Result

What are the differences between influence diagrams and decision graphs?

- The arrows in influence diagrams have two different meanings. The arrows starting at *S* or at action nodes express genuine causal dependence, whereas the arrows arriving at action nodes express an informational influence (which in the end has also causal influence on the actions).
- However, the arrows should represent not some objective causal picture, but the causal pictue of the agent. And then those arrows violate the „no probabilities for acts“ principle and the ensuing „acts are exogenous“ principle.
- The decision graph respects both principles by eliminating all arrows ending at action nodes.

Dependency Schemes

- We might give influence diagrams an intentional reading: by interpreting the arrows ending at action nodes as intentions about how to make the action probabilistically and causally dependent on ist parents or causal predecessors.
- Such an intention is expressed in a dependency scheme: i.e., a probability function $q(a \mid x)$ over the values a of the action variable $A \in H$ conditional on possible values x of the set of its causal predecessors, for each $A \in H$.
- Each dependency scheme q has an expected utility, which again can be maximized.
- A dependency scheme seems to be the same as a strategy. The crucial question, however, is: Can each dependency scheme be intentionally realized? And the crucial answer is: generally not!
- Only those dependency schemes are strategies which can be intentionally realized.

Reflexive Rationality

Wolfgang Spohn

Frege Lectures

University of Tartu, June 25 – 27, 2013

Session 2: Endogenous Preference Change

What's Endogenous Preference Change?

Examples:

- Ulysses and the Sirens
- Getting Addicted
- How to get to a party and how to get back home?
- Writing down a telephone number

The general point:

- So far, probabilities may change through information, but only in this way. Hence expected utilities may change as well, but intrinsic utilities don't change. These are the assumptions of the subject herself, not our external assumptions.
- Now we intend to consider any kind of probability change (through forgetting, alcohol, etc.) and any kind of intrinsic and expected utility change (through addiction, bewitching, maturing, aging) – again as foreseen or envisaged by the subject herself.

McClennen's Basic Example

Consider various gambles:

g_1 = [\$ 2400, 1]

g_2 = [\$ 2500, 33/34 ; \$ 0, 1/34]

g_3 = [\$ 2400, 34/100 ; \$ 0, 66/100]

g_4 = [\$ 2500, 33/100 ; \$ 0, 67/100]

g_3^+ = [\$ 2401, 34/100 ; \$ 1, 66/100]

Assume the subject has the following (plausible) preferences:

$g_1 \succ g_2$, $g_4 \succ g_3$, and indeed $g_4 \succ g_3^+ \succ g_3$.

- This is a variation of the famous Allais paradox.
- The first two preferences look incoherent, they violate Savage's sure-thing principle, The third preference may be assumed given the second, and the fourth follows from the strict dominance principle.

The incoherence may be unfolded as follows (the slides are taken from E.F. McClennen, *Rationality and Dynamic Choice*, CUP 1990):

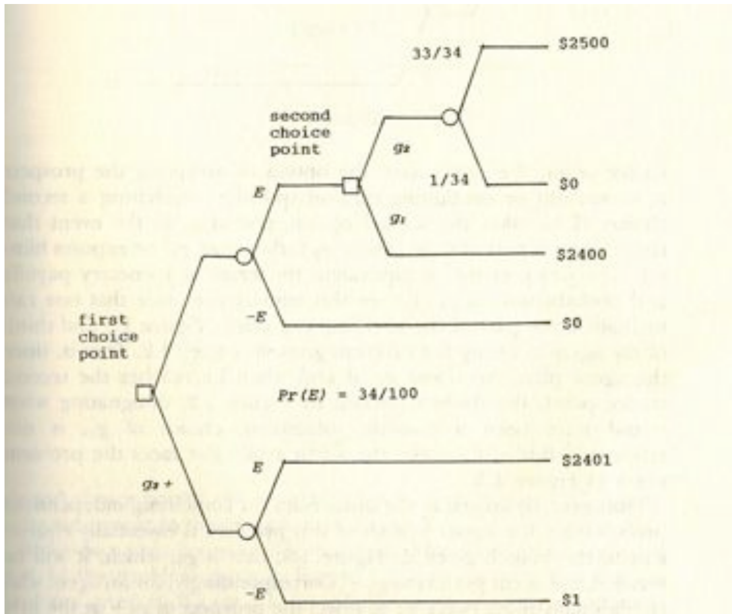

Figure 1.1

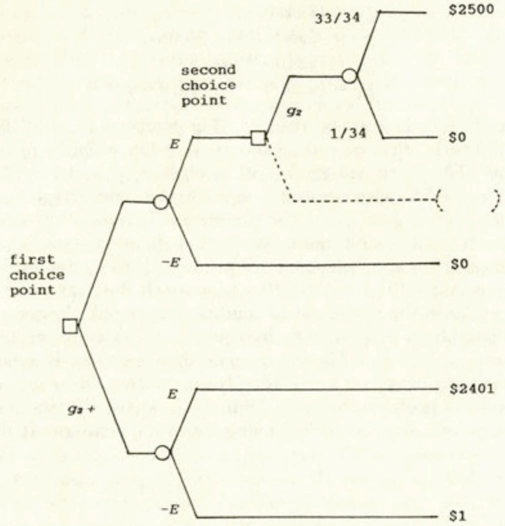

Figure 1.2

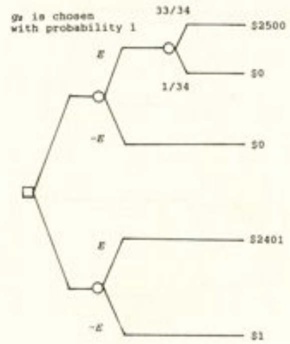

Figure 1.3

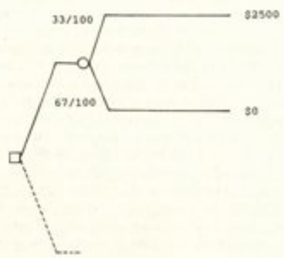

Figure 1.4

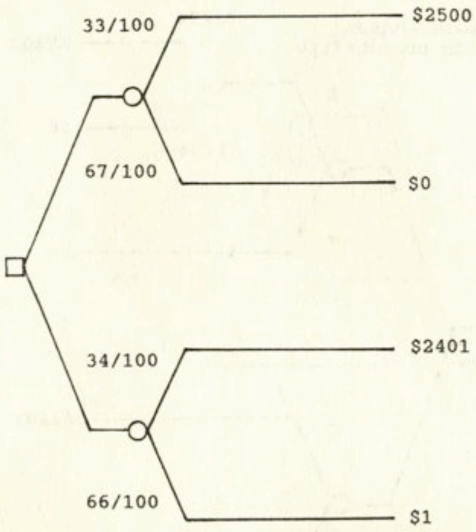

Figure 1.5

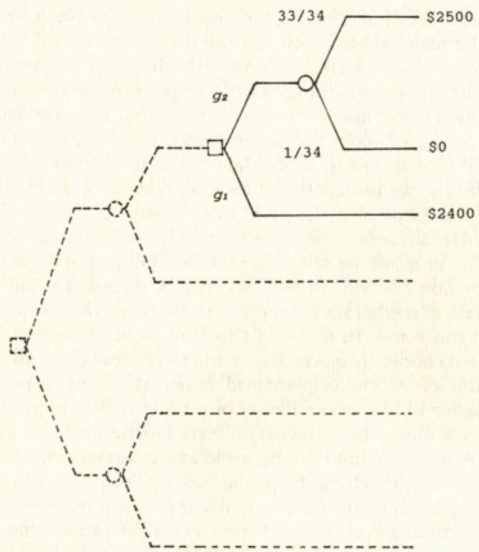

Figure 1.6

Three Kinds of Choice

McClennen's example:

- Myopic choice goes up at the first choice point and down at the second (should it be reached), and thus ends up worse than in any other strategy.
- Sophisticated choice, foreseeing the trap, goes down at the first choice point.
- Resolute choice goes up at the first choice point and also at the second (should it be reached) and thus withstands the seduction to go down at the second.
- The issue of feasibility of resolute choice

Three Kinds of Choice

The Party Example:

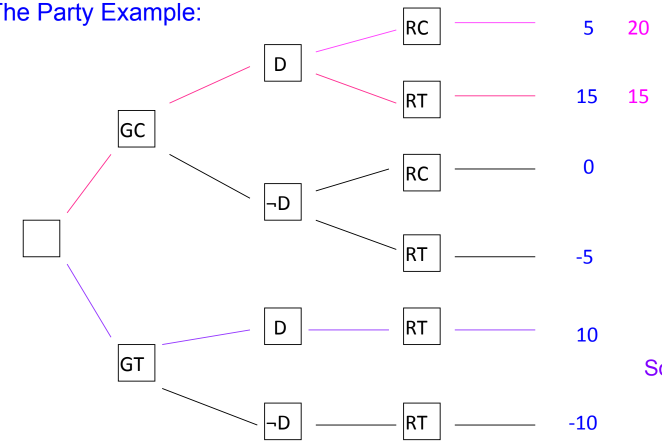

G: go
R: return
C: by car
T: by taxi
D: drink

Myopic
Sophisticated
Resolute
Choice

Three Kinds of Choice

Note that there is a double recursive deliberation involved in that picture (contrary to standard decision theory where just one roll-back analysis is made for a decision tree):

(1) There is a (degenerate) recursive roll-back analysis for the decision situation after drinking.

(2) And there is a full roll-back analysis for the entire graph – or rather for the modified graph which replaces the graph for the changed decision situation (after drinking) by the action considered optimal in that situation (as determined in the first step).

In other words: what is done in the changed decision situation is reevaluated in the original decision situation.

This is better represented by something like the following graph:

Three Kinds of Choice

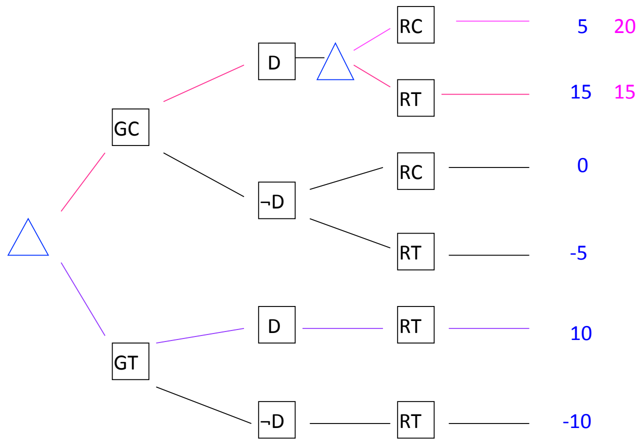

Myopic
Sophisticated
Resolute
Choice

How to Choose Between Sophisticated and Resolute Choice

- One problem is the alleged unfeasibility of resolute choice, because of which McClennen is usually confonted with the incredolous stare.
- McClennen himself (ch. 11) gives a complicated argument why sophisticated choice may be pragmatically deficient in a way in which resolute choice is not.
- The issue is difficult. I think the point is not to show that one method is superior to the other once for all. Rather, sometimes one method is feasible and recommendable, and sometimes the other. The problem is not to decide between the methods, but to find a criterion telling when one should favor one method over the other.
- My next goal is to argue that the present conceptual means don't provide the resources for stating such a criterion.

Global Decision Models, or the Agent Normal Form

If we want to generally capture such situations in which changes of one's decision situation are envisaged, then, it seems, we have to insert a decision node before each action node, summarizing the (possibly changed) decision situation one might be in before that action – where each such decision situation is of the local kind initially sketched. I call this a global decision model.

This corresponds to what game theorists call the agent normal form of a game that represents each action node by just one player.

The picture then is that each local decider or agent tries to do the best from his point of view. And so does the first local decider or agent expecting that the later ones will do their best.

Global Decision Models, or the Agent Normal Form

As such this does not yet say what it means to "try to do the best", what the global decision rule for such global decision situations is to be. Certainly, sophisticated choice as sketched above may be rigorously stated as such a global rule (and it basically agrees what game theorists say about the agent normal form.) But one might as well consider resolute choice that does not insert a decision node before each action node and decides about an entire course of actions at the decision nodes inserted.

One should also note that local deciders or agents are not persons. In fact, in the present decision-, not game-theoretic context all such local deciders are possible stages of one and the same person. And intuitively one should expect that this person tries to somehow integrate all the local agents into one coherent picture.

This very obscure remark indicates that there is something wrong with, or missing in, those agent normal forms.

This is what I now want to argue more rigorously:

Ambiguous Global Decision Models

Interpretation 1: At chance node C you receive either of two pieces of information and thus move to the belief states in Δ_2 or Δ_3.

Interpretation 2: At chance node C you forget something in either of two ways and thus move to the belief states in Δ_2 or Δ_3.

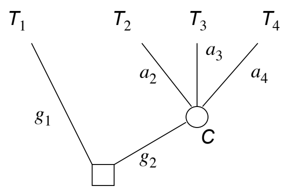

where each T_i has the form:

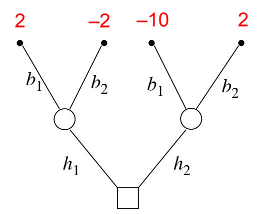

Ambiguous Global Decision Models

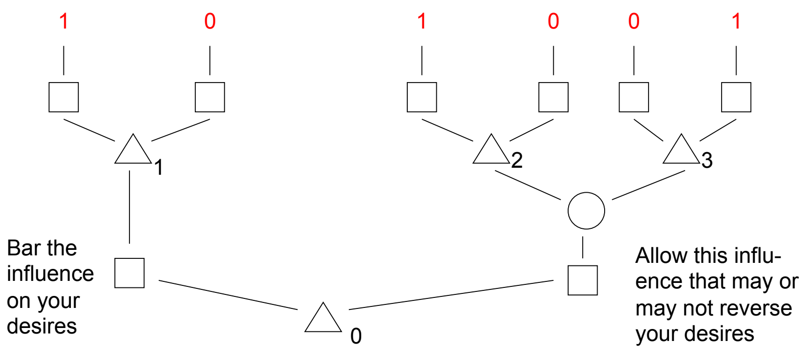

Ambiguous Global Decision Models

The crucial point of these examples that in each of their two interpretations we have exactly the same probabilities and utilities in all local decision situations and hence exactly the same global decision model.

Still, it is intuitively very clear that different choices are rational in each of the two interpretations. More precisely, the recursive deliberations are done in a different way.

The conclusion must be that global decision models do not provide sufficient resources for adequately dealing with decision situations in which the change of local decision situations is envisaged. But what is missing?

Reflexive Rationality

Wolfgang Spohn

Frege Lectures

University of Tartu, June 25 – 27, 2013

Session 3: Reflexive Decision Theory

Towards Reflexive Decision Theory

- In the first lecture I have introduced decision graphs, i.e., causal graphs with action nodes.
- More specifically, I have (not quite explicitly) introduced basic decision models $\langle U, H, \longrightarrow, p, u \rangle$, where
  - U is the set of all variables considered,
  - H is the set of action variables,
  - $\longrightarrow$ represents direct causal dependence between variables,
  - p provides conditional probabilities $p(x \mid v)$ for each value x of a chance variable X conditional on values v of the parents of X (and all probabilities entailed, which amount to probability measures for $\times(U - H)$ conditional on action sequences h in $\times H$).
  - u provides utilities for all worlds in $W = \times U$.

Towards Reflexive Decision Theory

- Action nodes were represented by □ and chance nodes by ○. (Note the ambiguity: in decision trees nodes represent events, in decision graphs they represent variables.)
- Occasionally, I have added decision nodes △ to the graphic representation. Again note the ambiguity. Standardly, decision and action nodes are not distinguished; in fact, in decision trees action nodes are usually called decision nodes – an original sin.
- The decision nodes introduce a reflexive perspective; according to such a representation, the subject reflects on the decision situations it might get into. We might call decision graphs thus enriched reflexive decision graphs.
- However, so far I have not provided any theory for those reflexive decision graphs. This is what I want to do in the following.

Causal Graphs and Bayesian Nets

A causal graph is a directed acyclic graph $\langle U, \longrightarrow \rangle$ the nodes in U of which represent variables and the arrows of which represent direct causal dependence between variables ("direct" relative to the given frame = the set of nodes or variables = U).

Causal dependence entails temporal succession. Hence, the variables must be specific, temporally located variables, and the arrows (or vertices) have to agree with the temporal order.

A Bayesian net $\langle U, \longrightarrow, p \rangle$ is a directed acyclic graph $\langle U, \longrightarrow \rangle$ together with a probability measure p for U agreeing with the graph, i.e. such that for each node A in U the set of parents of A is the minimal set X of variables preceding A such that A is independent from all the other variables preceding A given X w.r.t. p.

(If arrows are interpreted causally this is tantamount to assuming the Markov condition for causal chains and the principle of the common cause for causal forks.)

Reductions of causal graphs and Bayesian nets

We may take the probabilities as reflecting the causal relations. Or we may take the former as defining the latter. Then, however, a frame-relative notion of causation results (relative to the frame *U*).Thus, the theoretical task arises to inquire the relation between the frame-relative causal dependencies in richer and coarser frames. Thereby we can see to which extent causal relations in the coarser frame are indicative of causal relations in the richer frame and thus, in the final analysis, of the real causal relations in the fictitious universal frame.

[The point of the exploration starting here will be fully clear only at slide 23.]

Hence, suppose $\langle U, \longrightarrow, p\rangle$ is a Bayesian net. Let $U^* = U - \{C\}$ and p^* be the marginalization of p to U^*. How, then, does the coarser Bayesian net $\langle U^*, \longrightarrow^*, p^*\rangle$, the reduction of $\langle U, \longrightarrow, p\rangle$ by C, look like? That is, what is $\longrightarrow^*$?

We have to distinguish three cases: the IC-, the CC-, and the N-case.

Reductions of causal graphs and Bayesian nets

The IC (indirect cause) case:

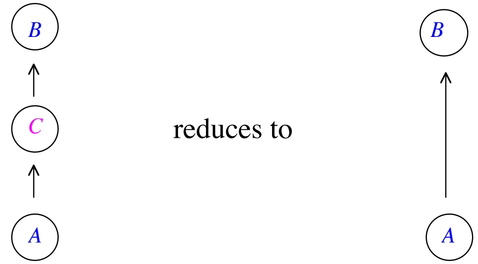

Reductions of causal graphs and Bayesian nets

The CC (common cause) case:

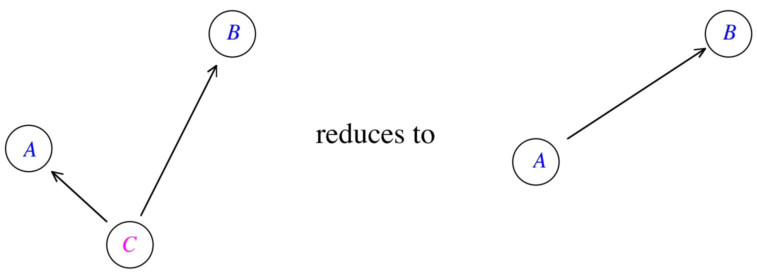

Reductions of causal graphs and Bayesian nets

The N (shaped or neighbor of a CC) case (but let's ignore it):

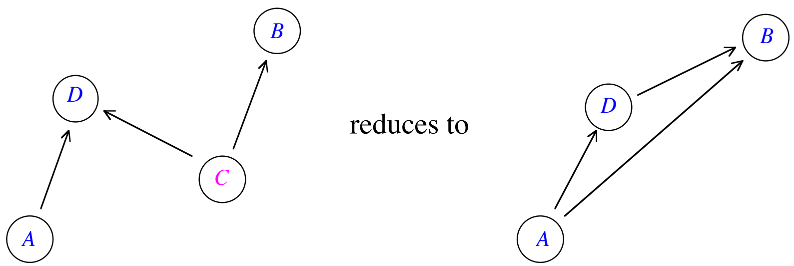

Reductions of causal graphs and Bayesian nets

Theorem: Let $\langle U, \longrightarrow, p\rangle$ be a Bayesian net and $\langle U^*, \longrightarrow^*, p^*\rangle$ be the reduction of $\langle U, \longrightarrow, p\rangle$ by C (as explained on the previous pages). Then we have: if p is faithful to $\langle U, \longrightarrow\rangle$, then $\langle U^*, \longrightarrow^*, p^*\rangle$ is indeed a Bayesian net.

Or conversely: Each arrow (direct causal dependence) in the coarser graph may unfold in the richer graph into a causal chain (IC case) or a causal fork (CC case), and each triangle may unfold into an N-shaped structure (N case).

Causal graphs with action nodes

We want to look now at causal graphs and Bayesian nets form the point of view of an agent who wants to place herself in the graph.

Hence, we distinguish a set $H \subseteq U$ of action variables (symbolized by boxes) and call the remaining variables in $W = U - H$ chance variables.

Now, let $\langle U, H, \longrightarrow, p\rangle$ be a Bayesian net with action nodes. Is this suited for characterizing the agent? No, it is rather suited to characterize what other persons believe about the agent and how his actions are embedded into the world.

Truncations of Bayesian nets

Pearl has proposed (and I have made an equivalent proposal 20 years earlier) to characterize the agent's point of view by the truncation $\langle U, H, \longrightarrow^*, p^*\rangle$ of $\langle U, H, \longrightarrow, p\rangle$ w.r.t. H, where $\longrightarrow^*$ is obtained from $\longrightarrow$ by deleting all arrows in $\langle U, H, \longrightarrow\rangle$ ending at action nodes in H, and where p^* is the so-called truncated factorization of p.

Example:

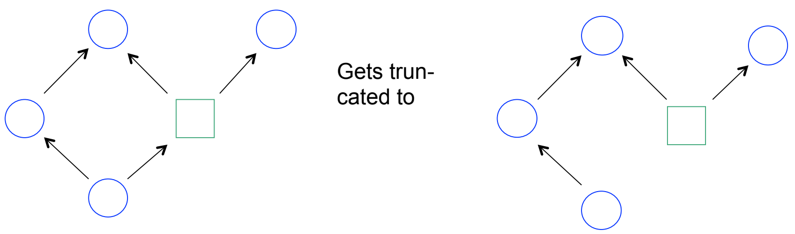

Truncations of Bayesian nets

It follows that action variables are exogenous in the truncated graph.

One may argue about whether or not the truncated factorization p^* of p should assume any probabilities for the action variables themselves. Long ago I have defended the principle: no probabilities for actions (so that p^* is in fact a family of probability measures for the occurrence variables in W, one for each course of action).

This principle entails the exogeneity of action variables.

Towards Reflexive Decision Models

But why should we accept the "no probabilities for actions" principle or the exogeneity of actions even from the agent's point of view? Is the agent unable to take a doxastic attitude towards her own actions? Is she unable to have a causal explanation for her own actions? Must she see her actions as uncaused?

Surely not. She can have beliefs about and explanations for her actions, just as external observers can have them. And if she takes herself to be rational, it is clear what the causes of her actions are. They just consist in the decision situations out of which she acts (where a decision situation is always her subjective view of her situation).

Hence, we have to enrich in turn the truncated Bayesian net by decision nodes containing possible decision situations, thus returning from decision graphs to very special influence diagrams. I call these enriched structures reflexive decision models.

Towards Reflexive Decision Models

$\delta = \langle U, H, D, \longrightarrow, p, u\rangle$ is a reflexive decision model iff the following conditions are satisfied:

(1) $\langle U, \longrightarrow\rangle$ is a causal graph.
$H \subseteq U$ is the set of action nodes (marked by boxes).
$D \subseteq U$ is the set of decision nodes (marked by triangles)
with $H \cap D = \emptyset$.
$W = U - (H \cup D)$ is the set of chance nodes (marked by circles).

Towards Reflexive Decision Models

(2) p is a probability measure for U agreeing with $\langle U, \longrightarrow\rangle$.

(3) u is a utility function (for $U - D$). (Hence, being in a certain decision situation does not receive any intrinsic utility.)

(4) p is to assigns probabilities to actions conditional to decision situations which decide about the actions such that irrational actions receive conditional probabilities 0; a precise statement of this condition requires, however, to be precise about the relevant decision rule, something we are still working up to.

Towards Reflexive Decision Models

(5) each action node has exactly one decision node as parent;
and each decision node has at least one action node as a child.

We cannot allow that an action node is governed by more than one decision node.

But we may allow:

(a) that a decision node governs more than one action node,

(b) that a decision node does not immediately temporally precede its action children,

(c) that a decision node has other children besides action nodes.

Towards Reflexive Decision Models

(6) Finally, we have to make assumptions about the possible decision situations which are the values of decision nodes. In particular, there is a temporally first decision node $\Delta_0 \in D$ which consists of decision situations the agent might have been in at that time and a specific decision situation $\delta_0 \in \Delta_0$ the agent is actually in, i.e., believes to be in, i.e., $p(\delta_0) = 1$. (This is Eells' so-called reflexivity condition.)

What is the relation between δ_0 and $\underline{\delta}$? Obviously, δ_0 represents the very same decision situation in an unreflected way which $\underline{\delta}$ represents in a reflexive way. The formal condition is this:

(7) δ_0 is the truncated reduction of $\underline{\delta}$ by $\{\Delta_0\}$, i.e., results from first reducing $\langle U, \longrightarrow, p\rangle$ by the first decision node Δ_0 and then truncating the resulting net w.r.t. to the action children of δ_0. Let me explain:

Truncated reductions of reflexive decision models

Example:

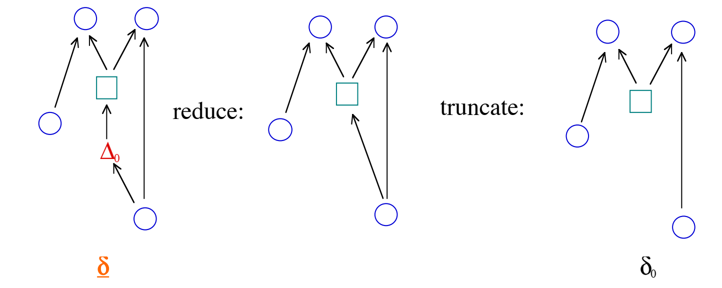

Truncated reductions of reflexive decision models

But what if Δ_0 has occurrence variables as children? Compare:

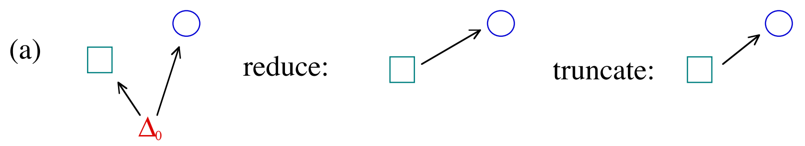

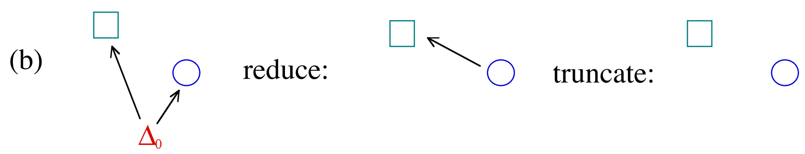

Truncated reductions of reflexive decision models

Thus, two structurally identical reflexive decision situations result in two structurally different unreflexive decision situations. This must not be! Hence, case (b) must rather be read thus:

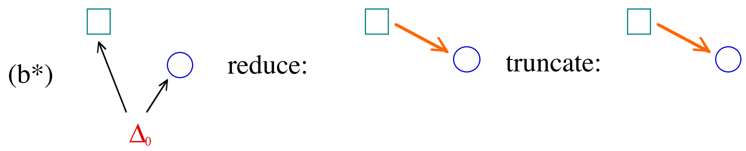

i.e., the arrows produced by common causes in reduction must always start at the action nodes irrespective of temporal order!

Truncated reductions of reflexive decision models

The N-case can't apply in truncated reductions (but, again, forget about it):

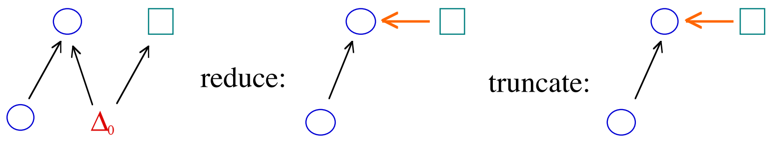

Truncated reductions of reflexive decision models

So, the conclusion is that truncated reductions must be allowed to produce backward arrows, which, however, do not represent mysterious backwards causation, but rather hide that the decision situation is just an ordinary common cause of the action and other occurrences.

Core example: How is the causal relation between smoking and lung cancer (so that we better shouldn't smoke)? Compare three stories:

- Smoking itself causes cancer ⇒ stop smoking!
- Some gene is a common cause of the desire to smoke and thus smoking and cancer ⇒ don't stop smoking!
- Only the desire to smoke (not the smoking itself) causes lung cancer ⇒ stop smoking!

Two standard puzzles are suited for exemplifying these considerations: the Toxin Puzzle and Newcomb's Problem.

The Toxin Puzzle: Version 1

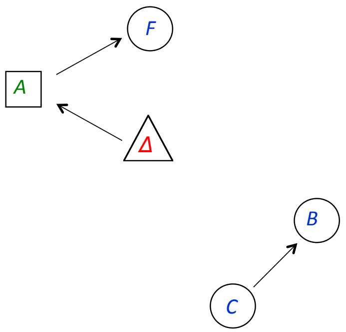

F: Feel sick (or not) after noon

A: Action: Drink the toxin (or not) at noon

Δ : Decision (Intention): Decide to drink (or not) the toxin short before noon

B: Money is on the bank (or not) after midnight

C: Cerebroscope is positive (or negative) at midnight

The Toxin Puzzle: Version 2

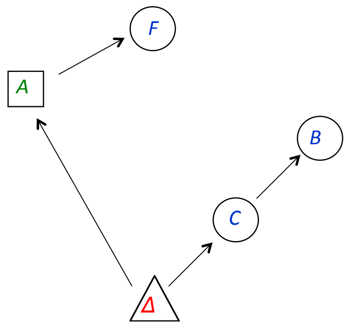

F: Feel sick (or not) after noon

A: Action: Drink the toxin (or not) at noon

B: Money is on the bank (or not) after midnight

C: Cerebroscope is positive (or negative) at midnight

Δ : Decision (Intention): Decide to drink (or not) the toxin short before midnight

The Toxin Puzzle: Version 2

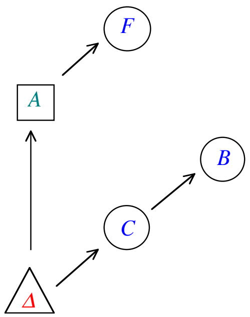

reduce &
truncate

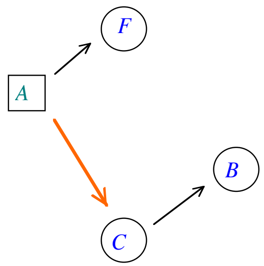

Newcomb‘s Problem: Eells‘ Version

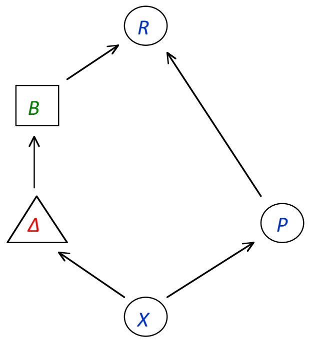

R: Reward between 0 and 1,001 M$

B: Action: Taking 1 or 2 boxes

Δ : Decision to take 1 or 2 boxes

P: Prediction that you will take 1 or 2 boxes

X: Some conjectured common cause of prediction and decision

Newcomb‘s Problem: My Version

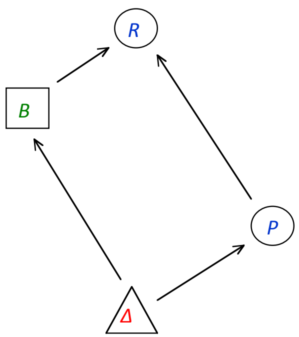

R: Reward between 0 and 1,001 M$

B: Action: Taking 1 or 2 boxes

P: Prediction that you will take 1 or 2 boxes

Δ : Decision to take 1 or 2 boxes

Newcomb’s Problem: My Version

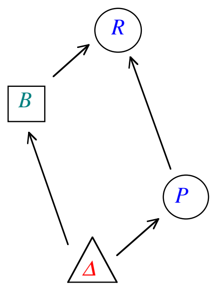

reduce &
truncate

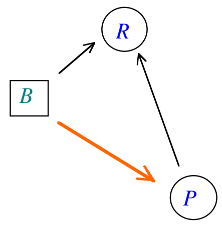

Adding Metastructure to Reflexive Decision Models

Our discussions so far suggest two places where to add further structure to reflexive decision models.

- First, reflexive decision models, as presented so far as decision graphs, in fact represent more explicitly the global decision models sketched in the last session as decision trees. However, we had observed that global decision models were insufficient. And reflexive decision mo-dels are still incomplete, since axiom 4 (slide 17) refers to a decision rule for those models which I have not yet specified.
- Second, the discussion of the Toxin Puzzle and Newcomb‘s Problem has suggested that it need not be fixed in advance where the decision nodes are to be placed. There is some freedom of choice.

Let me indicate these amendments of reflexive decision models.

Towards a Decision Rule for Reflexive Decision Models

- I argued that something is missing in global decision models (= reflex-ive decision models so far). Reflecting on the examples demonstrating this, the following suggestion appears plausible to me:
- We have to distinguish between moving to a superior, to an inferior, or to an equally valuable (actual or only possible) local decision situation, i.e., basic decision model.
- This superiority assessment is entirely subjective, there are so far no prescriptions. But we can presumably agree on the following instances:

Towards a Decision Rule for Reflexive Decision Models

(1) Learning or getting information moves one to a superior situation (in which one can make a better informed decision).
(2) Putting on eye glasses and thus seeing sharply and not dimly moves one to a superior situation.
(3) Forgetting (relevant facts) moves one to an inferior situation.
(4) Developing positive addictions (music, wine?) moves one to a superior (or equally valuable?) situation.
(5) Developing negative addictions (drugs) moves one to an inferior situation.
(6) Maturation and cultivation move one to a superior situation.
(7) Aging moves one to an equally valuable situation.
(8) Losing one's desire to eat through eating moves one to an equally valuable situation.
(9) Acquiring new desires by tasting, making new experiences, etc. moves one to an equally valuable (or superior?) situation.

Towards a Decision Rule for Reflexive Decision Models

The next task would be to precisely describe how the roll-back analysis or recursive deliberation depends on this superiority assessment. Three rules are suggested by the examples:

- If the change moves one to an inferior situation, then what is found optimal in that inferior situation needs to be reevaluated from the point of view of the previous superior situation.
- If the change moves one to a superior situation, then what is found optimal in that superior situation need not to be reevaluated from the point of view of the previous inferior situation.
- The same holds for moves to an equally valuable situation; I can't see so far that superior and equally valuable situations should make a difference to the roll-back analysis.

Towards a Decision Rule for Reflexive Decision Models

It's quite a different matter to generalize these three rules to a completely general decision rule for reflexive decision models. The recursive considerations may mix then in intricate ways. But I think I see how this can be done. The test for the adequacy of the very complicated resulting rule would then consist in simple situations which one can also intuitively grasp, such as the various examples I have considered.

Choosing when to choose

- The decision rule coming forward from that refined roll-back analysis refers to reflexive decision models with fixed decision nodes, but refines sophisticated and resolute choice by additionally taking that superiority assessment into account.
- However, it is neutral between sophisticated and resolute choice, i.e., it is able to accommodate both:
- In extreme form, sophisticated choice is characterized by having, for each action node, one immediately preceding decision node and performing the refined roll-back analysis for that reflexive decision model.
- In extreme form, resolute choice is characterized by having all action nodes decided in the initial decision node (which requires only a rudimentary roll-back analysis).

Choosing when to choose

- Our treatment of the Toxin Puzzle (and of Newcomb's Problem) suggests that there is indeed a choice where to place the decision nodes; we can opt for a late decision or an early decision (= commitment).
- This helps for a criterion deciding whether sophisticated or resolute choice is optimal in a given decision situation:
  - apply the decision rule for reflexive decision models to various distributions of decision nodes,
  - determine the maximal expected utility achievable according to each distribution,
  - and then choose such a distribution for which this value is maximal.
- The Toxin Puzzle exemplifies how the maximal expected utility can vary for different distributions of decision nodes and how one can maximize the maximal expected utility.

Choosing when to choose

- It should be clear that the procedures sketched do not only work for the extreme forms of sophisticated and resolute choice. They work for any distribution of decision nodes and hence for any mixtures of sophisticated and resolute choice, in which there are one or more decision nodes governing more than one action node.
- The resulting theory is obviously very complicated. I hope, however, that each complication was well motivated. And moreover I hope to have at least indicated that all the complications can be integrated into one coherent theory.

Reflexive Rationality

Wolfgang Spohn

Frege Lectures

University of Tartu, June 25 – 27, 2013

Session 4: Standard Game Theory

Game Theory

- Game theory is decision theory, or a theory of practical rationality, for several rational agents in social interaction.
- Hence, one may see decision theory as the trivial one-person specialization of game theory. This is the traditional perspective from 1944 where standard game theory started. (Thereby, some notion of rationality is appealed to that is still in need of explication and that reduces in the special case to the decision-theoretic notion.)
- Or one may see game theory as a special case of decision theory, as a decision theory for a special type of situations. One player may conceive of the actions of the other players simply as further chance nodes, and then decision theory as presented applies.

Game Theory

- However, if so, then the actions of the other players are very special chance nodes. The other players are not random devices, but also rational agents, to which presumably precisely the theory of rationality applies one intends to develop.
- Moreover, the theory is not for one, but for each player. And if our modeling assumes that each player sees the others as mere random devices, then it would assume gross error on behalf of all players. This must be avoided. Rather, the model assumptions about the players should be such that they impute the very same assumptions to the players.
- This is the more recent perspective of so-called epistemic game theory (starting ca. 1980). It suggests that the reduction of game-theoretic to decision-theoretic rationality may be complicated.

Games in Extensive Form

- Games can be represented in extensive form (just like decisions).
- A game tree is like a decision tree with chance and decision nodes.
  - One difference is that there are decision nodes for all of the n (≥ 2) players.
  - The probabilities at the chance nodes are fixed; this means that they are objective probabilities and commonly known or that they are intersubjectively shared for some other reason.
  - Another difference is that all branches of the tree receive utilities for all players, which may, of course, diverge.
  - A further difference is that the decision nodes of one player are partitioned into so-called information sets. The idea is that the player cannot distinguish at which node of an information set he actually is; to this extent he doesn't know the past evolution of the play.
  - This has the formal consequence that the possible actions a player can take at her decision node must be the same for all nodes in an information set.

Games in Extensive Form

- Although it is most venerable, I don't like the extensive form:
  - it is very explicit in some respects (that's good),
  - but at the cost of being clumsy and redundant (again, one and the same event or action is often multiply represented),
  - the order in the branches of the tree mixes temporal and epistemic aspects (as the notion of an information set displays),
  - and they represent the epistemic states of the players very poorly (everything besides the information sets is left implicit).
- As in decision theory, the extensive tree form can be reduced to the normal matrix form, which will do for us.
- However, one should observe that this reduction is not innocent; it creates artifacts and loses details – whence game theorists very often return to the extensive form.

Games in Normal Form

- As in decision theory, we can define strategies for each player, and given the probabilities at the chance nodes each strategy combination of all players receives an expected utility for each player. (Strategies themselves are so far not probabilistically assessed.)
- Let's simplify things and consider only two-person games with two players Ann and Bob. Many interesting phenomena already emerge in the two-person case, and the more-person cases may become most involved.
- So, let $A = \{a_1,\ldots,a_m\}$ be the set of (pure) strategies of Ann and $B = \{b_1,\ldots,b_n\}$ be the set of (pure) strategies of Bob.
- Ann has the expected utility function u and Bob the expected utility function v for $A \times B$, so that $u_{ij} = u(a_i, b_j)$ and $v_{ij} = v(a_i, b_j)$.

Games in Normal Form

- Then the game in normal form is represented as a double matrix:

	b_1	…	b_j	…	b_n
a_1	u_{11} v_{11}	…	u_{1j} v_{1j}	…	u_{1n} v_{1n}
…	…	…	…	…	…
a_i	u_{i1} v_{i1}	…	u_{ij} v_{ij}	…	u_{in} v_{in}
…	…	…	…	…	…
a_m	u_{m1} v_{m1}	…	u_{mj} v_{mj}	…	u_{mn} v_{mn}

- The standard assumption is that the whole set-up is mutual knowledge between Ann and Bob, i.e., Ann knows the sets A and B of possible strategies and the expected utlitiy functions u and v (which includes knowledge of the relevant probabilities). Bob knows all this as well, Ann knows that Bob does so, Bob knows that Ann does so, Ann knows that Bob knows that she knows all this, and so on.
- Otherwise, Ann and Bob would not be thinking about the same game.

Nash Equilibria

- What is the standard solution of such games? The basic idea is this:
- Given the unique recommendation of what is rational for Bob (which exists by assumption), Ann should choose a best reply against this recommendation, and likewise for Bob.
- That is, Ann should choose a_k and Bob b_l such that

  $u_{kl} \geq u_{il}$ for all $i = 1,\ldots,m$, and

  $v_{kl} \geq v_{kj}$ for all $j = 1,\ldots,n$.
- If and only if these two conditions are satisfied, then $\langle a_k , b_l \rangle$ form a Nash equilibrium in pure strategies.
- The problem is: Nash equilibria in pure strategies need not exist.
- The solution to this problem is: Ann and Bob need not choose pure strategies. They could also decide to play mixed strategies, and Nash equilibria in mixed strategies always exist. More precisely:

Nash Equilibria

- Let S be the set of all probability distributions over A (= mixed strategies of Ann) and T be the set of all probability distributions over B (= mixed strategies of Bob).
- Each mixed strategy combination $\langle s, t \rangle \in S \times T$ has an expected utility $u(s, t) = \sum_{ij} s(a_i)\, t(b_j)\, u_{ij}$ for Ann and an expected utility $v(s, t) = \sum_{ij} s(a_i)\, t(b_j)\, v_{ij}$ for Bob.
- Then, a strategy combination $\langle s^*, t^* \rangle$ is a Nash equilibrium (in mixed strategies) if and only if:

  $u(s^*, t^*) \geq u(s, t^*)$ for all $s \in S$, and

  $v(s^*, t^*) \geq v(s^*, t)$ for all $t \in T$.

Nash Equilibria

- In other words, none of the players has a reason to deviate from the equilibrium provided the other(s) conform(s) to it. In this sense, a Nash equilibrium is stable. [Note, however, in such an equilibrium the players need not necessarily have a reason to precisely stick to it; so, stability is not perfect.]
- Each game in normal form has at least one NE (Nash equilibrium in mixed strategies); so a rational solution in this sense always exists.
- However, a game may have many NE (Nash equilibria in mixed strategies); this generates a difficult selection problem.
- There is a huge discussion about the adequacy of this notion, pointing to various problems and many alternatives. Here, we need not engage in this discussion.

Matching Pennies

- This game is characterized by the following normal form:

	head	tail
head	1 -1	-1 1
tail	-1 1	1 -1

- This is the simplest example for a so-called zero-sum or constant-sum game; the gains of the one player are the losses of the other.
- It has no NE in pure strategies and exactly one in mixed strategies, where each player chooses his/her options with probability 0.5.
- In two-person zero-sum games the equilibrium strategies are at the same time their maximin strategies (the optimal defense against the opponent‘s wishing one‘s worst).
- Moreover, NE is essentially unique in those games.

Pure Coordination Game

- This game is characterized by the following normal form:

	a	*b*
a	1 1	0 0
b	0 0	1 1

- Here, the players have identical interests. Still, they have a problem.
- The pure coordination game has two NE in pure strategies with EU 1 for each, and one in mixed strategies (where each player plays (0,5 *a*, 0,5 *b*)) with EU 0,5 for each.
- It‘s not clear how they can reach coordination in this game in order to satisfy their interests.
- In cooperative game theory (where players are allowed to communicate and make agreements) this would be easy. However, we move in the context of non-cooperative game theory, which is always considered as basic (since cooperation should be treated as part of the game).

Bach or Strawinsky

- This game is characterized by the following normal form:

	B	*S*
B	2 1	0 0
S	0 0	1 2

- This is a coordination game with a slight conflict built into it.
- It has two NE in pure strategies ⟨*B, B*⟩ and ⟨*S, S*⟩ with obvious EU and one symmetric mixed NE ⟨ (1/3 *B*, 2/3 *S*), (1/3 *B*, 2/3 *S*) ⟩ with EU 2/3 for each.

Chicken, or: Hawk and Dove

- This game is characterized by the following normal form:

	H	*D*
H(awk)	0 0	7 2
D(ove)	2 7	6 6

- This is a very common social situation.
- This game has no pure NE; ⟨*D, D*⟩ looks reasonable, but is not stable. But it has exactly one symmetric mixed NE ⟨ (1/3 *H*, 2/3 *D*), (1/3 *H*, 2/3 *D*) ⟩ with EU 14/3 for each.

The Justification of Nash Equilibria

- I have mentioned one justification of NE: Given there is a unique rational solution of game situations, it must be a Nash equilibrium. (But is the presupposition satisfied? And what does „rational" mean here?)
- I have mentioned a second justification for the case of two-person zero-sum games.
- The strongest justification comes from the publicity requirement: whatever the theory recommends, its recommendations must allow publicity. And only NE allow that.
- However, this results in a reinterpretation of NE. We may conceive of a mixed strategy of Ann as something what Ann does (this was the former understanding) or as something what Bob believes about Ann. The latter results in an epistemic understanding of NE:

The Justification of Nash Equilibria

- That is, we may interpret a NE ⟨*s*, *t*⟩ as an *equilibrium of opinions*.
- For, why should Ann choose the mixed strategy *s*? Only when she does not care which of the pure strategies a_i with $s(a_i) > 0$ results from playing *s*. But how can she be indifferent? Only when all a_i with $s(a_i) > 0$ are equally good for her, i.e., have the same expected utility $\sum_j t(b_j)\, u_{ij}$ – where *t* now represents Ann's opinion about Bob's pure strategies. This indifference is guaranteed in the NE ⟨ *s*, *t* ⟩. The same hold for Bob's pure strategies b_j with $t(b_j) > 0$, when *s* represents his opinion about Ann's possible actions.
- Hence, only in such an equilibrium of opinions can the opinions of the players be mutual knowledge among the players, i.e., can the publicity requirement be satisfied.
- And then rationality means here the same as in decision theory, i.e., maximizing expected utility.

Prisoners‘ Dilemma (PD)

- This game is characterized by the following normal form:

	C	*D*
C(ooperate)	2 2	0 3
D(efect)	3 0	1 1

- This is also a very common social situation.
- This game has exactly one NE, a pure one: ⟨*D, D*⟩.
- This is unsatisfactory. Both players fare better, if they cooperate.
- This might be called a paradox of (the impossibility of) cooperation, or a conflict between individual and collective rationality, and it has provoked a huge discussion.
- The single-shot game looks tragic. But look at the finite repetition of the game. Then, again "always defect, always defect" is the only NE! This is no longer tragic, it is perfectly silly, close to a normative refutation of standard game theory.

The Ultimatum Game

- This game is characterized by the following normal form:

	Take 1	…	*Take 5*	…	*Take 9*
Accept	9 1	…	5 5	…	1 9
Reject	0 0	…	0 0	…	0 0

- This game has exactly one NE, namely: ⟨*Take 9* , *Accept* ⟩.
- Actually, we observe that the proponents often offer a fair division and that the respondents often reject, when offered a massively unfair division. And people don‘t think this is irrational.
- Another apparent failure of standard game theory!

Finally: the Centipede Game

- This game is more graphically represented in extensive form:

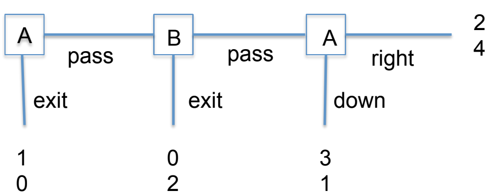

- This game has exactly one NE, namely: *exit* at the first left node.
- Obviously, you can indefinitely increase the number of legs and the final pay-offs.
- Still another apparent failure of standard game theory!

Reflexive Rationality

Wolfgang Spohn

Frege Lectures

University of Tartu, June 25 – 27, 2013

Session 5: Extension of Reflexive Decision Theory to Game Theory

A Hidden Assumption in Game Theory

- Does the reflexive perspective which I have introduced into decision theory have any consequences for game theory? It should have. But how?
- Look at the causal structure of two-person games in normal form. And take the options of the players to be simple actions (as they have been in all the examples we have considered) and not complex strategies (which in turn have a complex causal structure). Then the causal structure looks thus:

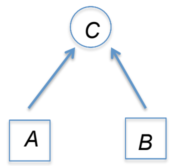

A: Ann's action variable (set of options)
B: Bob's action variable (set of options)
C: set of possible consequences

A Hidden Assumption in Game Theory

- Thus, the action variables *A* and *B* of Ann and Bob are causally independent. This is a basic assumption of non-cooperative games in normal form.
- (Of course, we could model also causal influence, when, e.g., Ann chooses first and Bob observes her choice. But then the game would have to be modelled in a different way.)
- According to how causal structure is reflected in Bayesian nets, this causal independence entails the probabilistic independence of the action variables.
- Nash equilibria respect this probabilistic independence, in both interpretations, as an equilibrium of mixed strategies or as an equilibrium of opinions.
- This assumption went undoubted for 60 years. Indeed, how could it be reasonably doubted?

A Hidden Assumption in Game Theory

- In discussing the Toxin Puzzle and generally the reduction of reflexive to simple decision models, we have seen how the inference from causal to probabilistic independence may be undermined, namely by the very decision nodes in the reflexive models.
- The causal graph for two-person game in normal form with decision nodes looks thus:

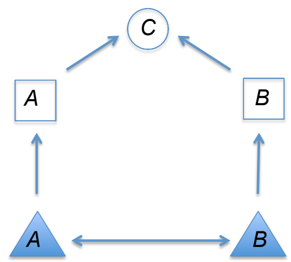

The triangles represent *A*'s and *B*'s decision node, the double arrow represents causal interaction.

A Hidden Assumption in Game Theory

- In this causal graph the action nodes are still causally independent, but the decision nodes are represented as causally entangled.
- This may well be so, since the decision nodes have complex ingredients and have no clear temporal extensions; so, all sorts of interpersonal exchange may be going on here.
- Of course, the decision nodes need not be so entangled. Each player may take a fresh decision after the interaction that no longer influences the other player(s). But they may be entangled.
- And for our purposes, the possibility of entanglement is sufficient. Because it entails that the common assumption of the probabilistic independence is unfounded! Let's give up this assumption!

Dependency Equilibria

- So, we want to allow now that Ann's conditional probabilities for Bob's actions or pure strategies vary. Thus, her opinions now take the form $q(b_j \mid a_i)$, where $q(. \mid a_i)$ is a distribution over B, for each $a_i \in A$. Reversely, Bob's opinions now take the form $r(a_i \mid b_j)$, where for each $b_j \in B$ $r(. \mid b_j)$ is a distribution over A.
- What does it mean under these assumptions for Ann to be rational? It means to maximize conditional expected utility, i.e., to choose on a_i for which $\sum_j q(b_j \mid a_i)\, u(a_i, b_j)$ is maximal. (Recall that this is where we arrived at in Lecture 1.)
- Likewise for Bob.

Dependency Equilibria

- However, it should be possible that these conditional probabilities q and r should be mutual knowledge, just like the utilities u and v. This has two consequences.
- The first consequence is that Ann's and Bob's conditional probabilities must combine into a single joint distribution p over $A \times B$, i.e., there must be a joint distribution p such that for all i and j $p(b_j \mid a_i) = q(b_j \mid a_i)$ and $p(a_i \mid b_j) = r(a_i \mid b_j)$. (Unlike in the case of NE, where s and t were independent, this condition is not automatically satisfied.)
- The second consequence is that all $a_i \in A$ with $p(a_i) > 0$ must have equal and maximal conditional expected utility for Ann. Otherwise, the probability distribution would contradict the rationality of Ann.
- Likewise for Bob.

Dependency Equilibria

So, let us define: The probability distribution p over $A \times B$ is a dependency equilibrium (DE) iff for all i with $p(a_i) > 0$ and all $k = 1, \ldots, m$

$$\sum_j p(b_j \mid a_i)\, u(a_i, b_j) \geq \sum_j p(b_j \mid a_k)\, u(a_k, b_j)$$

and reversely, for all j with $p(b_j) > 0$ and all $l = 1, \ldots, n$

$$\sum_i p(a_i \mid b_j)\, v(a_i, b_j) \geq \sum_i p(a_i \mid b_l)\, v(a_i, b_l),$$

i.e., if all of Ann's and Bob's actions that are not excluded, i.e., have positive probability according to p, have, respectively, the same maximal expected utility for Ann and Bob.

(This entails that dependency equilibria, DE, are awkward to calculate; we have to solve polynomial equations for that purpose.)

Dependency Equilibria

(1) For some $a_k \in A$ or $b_l \in B$ we may have $p(a_k) = 0$ or $p(b_l) = 0$ so that no conditional probabilities are defined for them and the definition just given makes no sense. This defect may, however, be removed in a precise and adequate way.

(2) Dependency equilibria are not to be confused with the correlated equilibria of Aumann (another interesting notion I am not going to explain here).

(3) DE form a wider class than NE. Those distributions p over $A \times B$ that factorize into independent s over A and t over B – NE apply only to such p – are obviously NE if and only if they are (degenerate) DE.

(4) Theorem: each pure strategy combination ⟨a_i, b_j⟩ or each p with $p(a_i, b_j) = 1$ that weakly Pareto dominates a NE is a DE.

(5) Conjecture: exactly those pure strategy combinations are DE that weakly Pareto dominate the maximin strategies of the players.

Matching Pennies

The game:

u \ v	b_1	b_2
a_1	1, 0	0, 1
a_2	0, 1	1, 0

NE:

p	b_1	b_2
a_1	$1/4$	$1/4$
a_2	$1/4$	$1/4$

DE:

p	b_1	b_2
a_1	x	$1/2 - x$
a_2	$1/2 - x$	x

where $0 \leq x \leq 1/2$

Hawk and Dove

The game:

u \ v	b_1	b_2
a_1	6, 6	2, 7
a_2	7, 2	0, 0

3 NE:

pure

p	b_1	b_2
a_1	0	1
a_2	0	0

pure

p	b_1	b_2
a_1	0	0
a_2	1	0

mixed

p	b_1	b_2
a_1	$4/9$	$2/9$
a_2	$2/9$	$1/9$

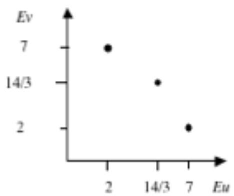

Hawk and Dove

The game:

u \ v	b_1	b_2
a_1	6, 6	2, 7
a_2	7, 2	0, 0

3 kinds of DE:

pure

p	b_1	b_2
a_1	0	1
a_2	0	0

pure

p	b_1	b_2
a_1	0	0
a_2	1	0

mixed

p	b_1	b_2
a_1	x	y
a_2	y	$1 - x - 2y$

where $y = 1/18\,(2 - 15x + \sqrt{4 + 156x + 9x^2})$

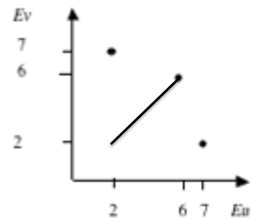

Prisoners‘ Dilemma (PD)

The game:

$u \backslash v$	b_1	b_2
a_1	2, 2	0, 3
a_2	3, 0	1, 1

NE:

p	b_1	b_2
a_1	0	0
a_2	0	1

Two kinds of DE:

p	b_1	b_2
a_1	$\frac{1}{2}\, x(1+x)$	$\frac{1}{2}\, x(1-x)$
a_2	$\frac{1}{2}\, x(1-x)$	$\frac{1}{2}\,(1-x)\,(2-x)$

where $0 \le x \le 1$

p	b_1	b_2
a_1	$\frac{3}{8}\,(1-x)\,(1+x)$	$\frac{1}{8}\,(1-x)\,(1-3x)$
a_2	$\frac{1}{8}\,(1+x)\,(1+3x)$	$\frac{3}{8}\,(1-x)\,(1+x)$

where $-1/3 \le x \le 1/3$

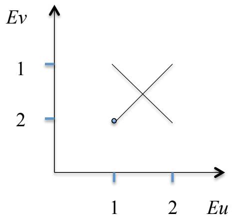

• : NE
— : DE

A General Remark

Facing recalcitrant or counter-intuitive examples:

- You may wiggle with the formal representation. There are lots of ways regarding, e.g., PD; for instance, we may always be inmidst an infinite repetition of PD; or always have an exit option; etc.
- Or you may wiggle with the utilities; e.g., assume other-regarding preferences, giving fairness an extra weight, etc. This mitigates the problems, but there is no guarantee that they vanish.
- Or you may – this seems to be a novel idea, though – wiggle with the probabilities, as exemplified already for the single-shot PD. (However, this works only if you accept my DE.)

Short Bibliography

Some Textbooks:

Raiffa, Howard (1968), *Introduction into Decision Analysis*, Reading, Mass.: Addison-Wesley[still my favorite textbook of standard decision theory as taught in economics]

Peterson, Martin (2009), *An Introduction to Decision Theory*, Cambridge: Cambridge University Press [an excellent and up-to-date introduction from the philosophical side]

Luce, R. Duncan, H. Raiffa (1957), *Games and Decisions*, New York: Wiley [somehow still very recommendable, although all modern developments are missing]

Hargreaves Heap, S.P., Y. Varoufakis (2004), *GameTheory: A Critical Text*, London: Routledge, 2nd ed. [perhaps the most recommendable introduction into game theory for philosophers, from experts with a foundational vein, little mathematics, many examples]

Myerson, Roger B. (1991), *Game Theory. Analysis of Conflict*, Cambridge, Mass.: Harvard University Press [probably still the most comprehensive textbook on game theory]

Relevant Literature:

Dependency equilibria are entirely of my own making (see my writings). It might be possible to relate them to:

Tuomela, Raimo (2000), *Cooperation*, Dordrecht: Kluwer [where similar ideas are sketched]

Concerning endogenous preference change, sophisticated choice and related phenomena there still do not seem to exist any textbooks or representative monographs. But there is:

McClennen, E.F. (1990), *Rationality and Dynamic Choice*, Cambridge: Cambridge University Press [as far as I know, still the only systematic and formal treatment from the philosophical side]

Elster, Jon (1979), *Ulysses and the Sirens. Studies in Rationality and Irrationality*, Cambridge: Cambridge University Press

Elster, Jon (1983), *Sour Grapes. Studies in the Subversion of Rationality*, Cambridge: Cambridge University Press

Elster, Jon (2000), *Ulysses Unbound: Studies in Rationality, Precommitment, and Constraints*, Cambridge: Cambridge University Press [all Elster books are rich and most instructive and inspiring, drawing on abounding material und displaying the deep relevance of those topics; however, no formal modeling]

Frank, Robert H. (1988), *Passions Within Reason. The Strategic Role of the Emotions*, New York: W. W. Norton & Company [another most inspiring book by a (non-standard) economist]

Own Writings:

„Strategic Rationality", Forschungsberichte der DFG-Forschergruppe *Logik in der Philosophie* Nr. 24, 1999, 55 S.

"Dependency Equilibria and the Causal Structure of Decision and Game Situations", *Homo Oeconomicus* 20 (2003) 195-255

"Dependency Equilibria", *Philosophy of Science* 74 (2007) 775-789

"From Nash to Dependency Equilibria", in: G. Bonnano, B. Loewe, W. van der Hoek (eds.), *Proceedings of LOFT 2008*, Texts in Logic and Games, Amsterdam: Amsterdam University Press, 2009

„Why the Received Models of Considering Preference Change Must Fail", in: T. Grüne-Yanoff, S. O. Hansson (eds.), *Preference Change: Approaches from Philosophy, Economics and Psychology*, Dordrecht: Springer, 2009

"Reversing 30 Years of Discussion: Why Causal Decision Theorists Should One-Box", *Synthese* 187 (2012) 95-122