An Arboretum of Decision Trees, or Varieties of Node in Decision Trees

Peter J. Hammond, Emeritus Professor of Economics
at Stanford and Warwick Universities
email: peter.hammond@stanford.edu
or p.j.hammond@warwick.ac.uk

2023 November 21st; typeset from arboretum.tex

Dynamic Inconsistency of Finite Horizon Planning

Jan de Villiers Graaff (1957)
Theoretical Welfare Economics (Cambridge)

There is a “horizon problem” in intertemporal welfare economics.

If one plans for only a finite horizon,
eventually the plan will run out, and will need revision.

Dynamic Consistency of Infinite Horizon Planning

Remedy for dynamic inconsistency:
face up to the complications of an infinite horizon.

PJH with James A. Mirrlees (1973) "Agreeable Plans"
in J.A. Mirrlees and N.H. Stern (eds.) *Models of Economic Growth*
(London: Macmillan, 1973), ch. 13, pp. 283–299.

PJH *Consistent Planning and Intertemporal Welfare Economics*
Ph.D. dissertation, Cambridge, 1973.

Facing up to some technical issues troubling Arrow in 1999:

PJH with Graciela Chichilnisky and Nicholas H. Stern (2020)
"Fundamental Utilitarianism and Intergenerational Equity
with Extinction Discounting"
Social Choice and Welfare 54: 397–427
(Special Issue: In Memory of Kenneth J. Arrow).

Systematic Study of Dynamic Consistency

Graham Pyatt recommended that I should read this key paper:

Robert H. Strotz (1956)
"Myopia and Inconsistency in Dynamic Utility Maximization"
Review of Economic Studies 23: 165–180.

Later paper:

Edmund S. Phelps and Robert A. Pollak (1968)
"On Second-Best National Saving and Game-Equilibrium Growth"
Review of Economic Studies 35: 185–199.

Their paper contrasts naïve versus sophisticated behaviour
in an economic growth model
with "hyperbolic" discounting of future generations' utilities.

That is, where each generation t's objective
is a weighted sum of its own and future generations' utilities
that takes the form $u_t + \alpha \sum_{\tau=1}^{\infty} \delta^{\tau-1} u_{t+\tau}$, where $\alpha, \delta \in (0, 1)$.

Raiffa's Preceding Work on Decision Trees

Howard Raiffa (1968) *Decision Analysis: Introductory Lectures on Choices under Uncertainty* (Addison-Wesley).

Focuses on decision trees (or "decision-flow diagrams", p. 10) with:

1. risk associated with roulette lotteries,
   involving specified hypothetical probabilities,
   rather than uncertainty associated with horse lotteries;
2. outcomes that are monetary amounts measured in dollars.

My Preceding Work on Decision Trees

Characterizations of expected utility, both objective and subjective:

(1983) "Ex-Post Optimality as a Dynamically Consistent Objective for Collective Choice Under Uncertainty"
in P.K. Pattanaik and M. Salles (eds.) *Social Choice and Welfare* (North Holland) ch. 10, pp. 175–205.

(1988) "Consequentialist Foundations for Expected Utility" *Theory and Decision* 25: 25–78.

Several (too many?) subsequent papers, including (1996) "Consequentialism, Structural Rationality and Game Theory"
in K.J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt (eds.) *The Rational Foundations of Economic Behaviour* (IEA Conference Volume No. 114) ch. 2, pp. 25–42.

(2022)
"Prerationality as Avoiding Predictably Regrettable Consequences"
Revue économique 73 (6): 943–976 (honouring Philippe Mongin).

Consequence Domains

A basic consequence domain is an abstract non-empty set Y.

Given a finite non-empty set S of states of the world s,
this can be extended to:

1. the domain $\Delta(Y)$ of finite roulette lotteries in the form of a probability mass function $Y \ni y \mapsto \lambda(y) \in [0,1]$ for which there exists a non-empty finite support $\operatorname{supp} \lambda \subseteq Y$ with: (i) $\lambda(y) > 0 \iff y \in \operatorname{supp} \lambda$; (ii) $\sum_{y \in \operatorname{supp} \lambda} \lambda(y) = 1$;
2. the product domain $Y^S = \prod_{s \in S} Y_s$ of finite horse lotteries (or Savage acts) in the form of mappings $S \ni s \mapsto y_s \in Y_s$;
3. the product domain $\Delta^S(Y^S) = \prod_{s \in S} \Delta(Y_s)$ of finite AA lotteries (or Anscombe/Aumann acts) in the form of mappings $S \ni s \mapsto \lambda_s \in \Delta(Y_s)$.

Thus, an AA lottery is a horse lottery
whose "prize" in each state $s \in S$ is a roulette lottery.

Consequence Choice Functions

Arrow, Kenneth J. (1959)
"Rational Choice Functions and Orderings"
Economica 26: 121–127.

Consider a consequence domain Y, which could be a lottery consequence domain like $\Delta(Y)$ or $\Delta^S(Y^S)$.

Let 2^Y denote the power set of all subsets of Y.

Let $\mathcal{F}(Y) \subset 2^Y$ denote the domain of feasible sets, which are defined as non-empty finite subsets of Y.

A choice function on a consequence domain Y is a mapping $\mathcal{F}(Y) \ni F \mapsto C(F) \in \mathcal{F}(Y)$ that satisfies $C(F) \subseteq F$ throughout the domain $\mathcal{F}(Y)$.

Arrow and Decisiveness

Arrow, Kenneth J. (1961) *Social Choice and Individual Values* (2nd. edition).

> *Abstention from a decision cannot exist;*
> *some [option] will prevail.*

When, despite Arrow, abstention really is possible,
it should be treated as an additional option n
signifying "none of the above".

Thus, instead of the feasible set being the pair $F = \{a, b\}$,
it becomes the triple $F' = \{a, b, n\}$.

Complete Base Preference Relations

The base preference relation $\succsim$ induced by the choice function C is the binary relation $\succsim$ on Y defined for all pairs $a, b \in Y$ by $a \succsim b \iff a \in C(\{a, b\})$.

That is the weak base relation; the corresponding

- strict base relation $\succ$ satisfies $a \succ b \iff b \notin C(\{a, b\})$;
- indifference relation $\sim$ satisfies $a \sim b \iff C(\{a, b\}) = \{a, b\}$.

Assumption

Choice is always decisive, with $C(F) \neq \emptyset$ for every finite set F.

Implication: the base relation defined by $a \succsim b \iff a \in C(\{a, b\})$ is always complete,
in the sense that either $a \succsim b$, or $b \succsim a$, or both.

In particular, $a \succ b \iff b \notin C(\{a, b\}) \iff C(\{a, b\}) = \{a\}$

Arrow's Two Audacious Questions

Arrow's *Social Choice and Individual Values*
raised the audacious question:
is there a satisfactory procedure for aggregating "individual values"
into a (complete and transitive) social preference ordering?

His dictatorship or "impossibility" theorem converted politics
from Bismarck's "art of the possible"
into the "science of the impossible".

Arrow also asked the less audacious question:
(why) should society have a social preference ordering?

Outline

The Potential Addict Example

Peter J. Hammond (1976) "Changing Tastes and Coherent Dynamic Choice" *Review of Economic Studies* 43: 159–173.

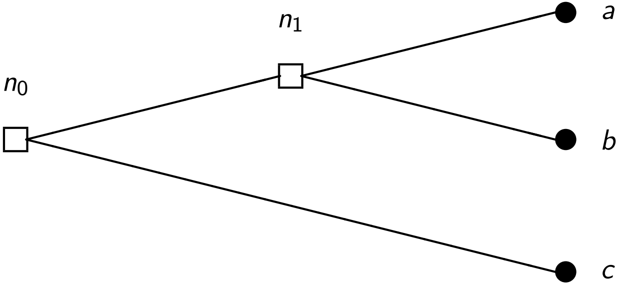

Figure: The potential addict's decision tree

Essential Inconsistency

Interpret:

- a as addiction or *akrasia* (weakness of will);
- b as bliss, or the best outcome;
- c as caution,
  or commitment to avoid any possibility of addiction.

Because the activity is addictive, typically one has $a \succ_1 b$
according to the (strict) preference relation $\succ_1$
that applies at node n_1.

But at node n_0, option b seems best and a seems worst,
so $b \succ_0 c \succ_0 a$ — in particular, $b \succ_0 a$.

This contradiction between the srict preferences $b \succ_0 a$
and $a \succ_1 b$ is an essential inconsistency.

The Case of a Preference Cycle

Example

Consider a preference cycle involving:

1. the triple $\{a, b, c\}$;
2. cyclic base preferences satisfying $a \succ b$, $b \succ c$, and $c \succ a$.

Given the feasible set $F = \{a, b, c\}$, suppose that $C(F) = \{b\}$.

This brings us back to the potential addict example.

Back to the Potential Addict Example

With $F = \{a, b, c\}$, we have $C(F) = \{b\}$
and $a \succ b$, $b \succ c$, and $c \succ a$.

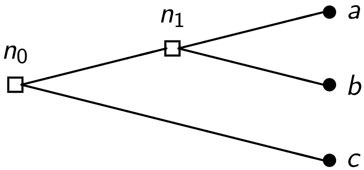

Any naïve decision maker may think they face
the reduced form decision problem,
with a single choice of $C(F)$ at node n_0, leading to $C(F) = \{b\}$.

Any sophisticated decision maker, however,
will realize that $a \succ b$ makes b infeasible,
so they choose $C(\{a, c\}) = \{c\}$. □

Naïve versus Sophisticated Behaviour

With $F = \{a, b, c\}$, we have $C(F) = \{b\}$
and $a \succ b$, $b \succ c$, and $c \succ a$.

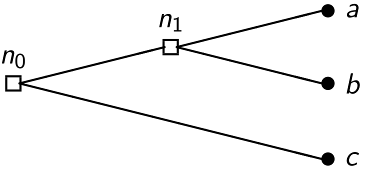

1. Naïve behaviour — at n_0 choose n_1, planning to go to the best possible outcome b later.

   But at node n_1, faced with the feasible set $\{a, b\}$, the revised base preference $a \succ_1 b$ says that the agent will choose the worst possible outcome a (addiction).

2. Sophisticated behaviour — at n_0 choose c (caution), avoiding the possibility of a altogether.

Outline

The Independence Axiom

Definition

The base preference relation $\succsim$ on the space $\Delta(Y)$ of letteries satisfies the independence axiom just in case, for every scalar $\alpha \in (0, 1] \subset \mathbb{R}$ and every triple $\lambda, \mu, \nu \in \Delta(Y)$, one has $\alpha\lambda + (1 - \alpha)\nu \succsim \alpha\mu + (1 - \alpha)\nu \iff \lambda \succsim \mu$.

Lemma

The base preference relation $\succsim$ on the space $\Delta(Y)$ of letteries satisfies the independence axiom just in case, for every scalar $\alpha \in (0, 1] \subset \mathbb{R}$ and every triple $\lambda, \mu, \nu \in \Delta(Y)$, one has

$$\alpha\lambda + (1 - \alpha)\nu \left\{ \begin{array}{c} \succ \\ \sim \\ \prec \end{array} \right\} \alpha\mu + (1 - \alpha)\nu \text{ according as } \lambda \left\{ \begin{array}{c} \succ \\ \sim \\ \prec \end{array} \right\} \mu$$

Essential Inconsistency Again

Suppose we have a strict violation of independence,
in the sense that there exist four lotteries $\lambda, \mu, \nu, \rho \in \Delta(Y)$
such that $\lambda \succ \mu$ and yet

$$\alpha\mu + (1-\alpha)\nu \succ \rho \succ \alpha\lambda + (1-\alpha)\nu$$

Then we have an essential inconsistency,
as in the first example with a strict preference cycle.

To demonstrate this, consider the decision tree on the next slide
where we distinguish between:

- decision nodes like n_0 and n_2, indicated by hollow squares,
  where the agent makes a choice;
- chance nodes like n_1, indicated by a hollow circle,
  where the next move is determined by a lottery
  with specified hypothetical ("objective") probabilities.

A Decision Tree with One Chance Node

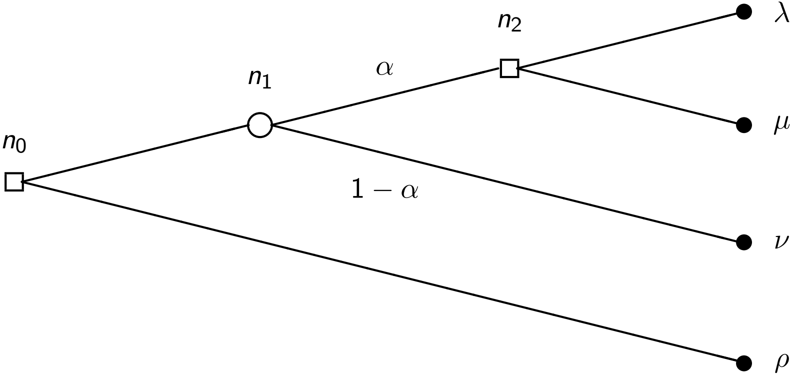

The hypothetical probabilities at the chance node n_1 are α and $1-\alpha$, where $0 < \alpha < 1$.

Three Possible Planned Lottery Consequences

Viewed from the initial node n_0,
the decision tree offers the feasible set

$$F = \{\alpha\lambda + (1-\alpha)\nu,\ \alpha\mu + (1-\alpha)\nu,\ \rho\}$$

consisting of the three consequence lotteries:

1. $\alpha\lambda + (1-\alpha)\nu$ in case the agent
   moves up at decision node n_0 and up at decision node n_2;
2. $\alpha\mu + (1-\alpha)\nu$ in case the agent
   moves up at decision node n_0 and down at decision node n_2;
3. ρ in case the agent moves down at decision node n_0.

What should the agent plan to choose
at the two decision nodes n_0 and n_2?

Naïve versus Sophisticated Choice

Recall the preferences $\lambda \succ \mu$ and yet

$$\alpha\mu + (1-\alpha)\nu \succ \rho \succ \alpha\lambda + (1-\alpha)\nu$$

1. A naïve agent considers only the feasible set

   $$F = \{\alpha\lambda + (1-\alpha)\nu,\ \alpha\mu + (1-\alpha)\nu,\ \rho\}$$

   and so plans to choose the optimal element $\alpha\mu + (1-\alpha)\nu$.

   But this naïve agent must choose up at n_0, and then face the continuation subtree with feasible set $\{\lambda, \mu\}$, from which λ is chosen, so the end result is $\alpha\lambda + (1-\alpha)\nu$.
2. A sophisticated agent, however, foresees that the consequence in the last continuation subtree of moving up at n_0 would be λ rather than μ, so the end result is $\alpha\lambda + (1-\alpha)\nu$.

   So the sophisticated agent chooses ρ instead of $\alpha\lambda + (1-\alpha)\nu$ the lottery consequence that results from moving down at n_0.

Outline

Source and Preliminaries

Ellsberg, Daniel (1961) “Risk, Ambiguity, and the Savage Axioms” *Quarterly Journal of Economics* 75 (4): 643–669.

> *Another example yields a direct test of one of the Savage postulates. Imagine an urn known to contain 30 red balls and 60 black and yellow balls, the latter in unknown proportion. One ball is to be drawn at random from the urn; the following actions are considered: (p. 653)*

Notation:

- R, B, Y are the colours red, black and yellow
- For each possible colour or state $\omega \in \{R, B, Y\}$, let:
  1. 1_ω indicates a prize of \$100;
  2. 0_ω indicates no prize.

Ellsberg's Four Bets and Their Results

	30	60				30	60	
	R	B	Y			R	B	Y
I	1_R	0_B	0_Y		III	1_R	0_B	1_Y
II	0_R	1_B	0_Y		IV	0_R	1_B	1_Y

A very frequent pattern of response is:
action I preferred to II, and IV preferred to III.

Less frequent is: II preferred to I, and III preferred to IV.

Both of these, of course, violate the Sure-thing Principle,
which requires the ordering of I to II
to be preserved in III and IV
(since the two pairs differ only in their third column,
constant for each pair). (p. 654)

Ambiguity Aversion

	30	60				30	60	
	R	B	Y	‖		R	B	Y
I	1_R	0_B	0_Y	‖	III	1_R	0_B	1_Y
II	0_R	1_B	0_Y	‖	IV	0_R	1_B	1_Y

Note that:

- I involves the known probability of $1/3$ of winning the prize;
- II involves an "ambiguous" probability of winning the prize;
- III involves an "ambiguous" probability of not winning;
- IV involves the known probability of $1/3$ of not winning.

So the preferences I $\succ$ II and IV $\succ$ III may be due to a phenomenon that has come to be known as "ambiguity aversion".

The Ellsberg Preferences Violate Independence

Raiffa, Howard (1961) "Risk, Ambiguity, and the Savage Axioms: Comment" *Quarterly Journal of Economics* 75 (4): 690–694.

Consider bet V $= \left(\frac{1}{2}(0_R + 1_R), \frac{1}{2}(0_B + 1_B), \frac{1}{2}(0_Y + 1_Y)\right)$
where, regardless of the colour ω:

- with probability $\frac{1}{2}$, one wins no prize;
- with probability $\frac{1}{2}$, one wins the prize.

The Ellsberg preferences for bet I over bet II,
and for bet IV over bet II, imply that

$$(1_R, 0_B, 0_Y) \succ (0_R, 1_B, 0_Y) \text{ and } (1_R, 0_B, 1_Y) \succ (0_R, 1_B, 1_Y)$$

If the independence axiom held, applying it twice would give

$$\begin{aligned} \text{bet V} &= \tfrac{1}{2}(1_R, 0_B, 0_Y) + \tfrac{1}{2}(0_R, 1_B, 1_Y) \\ &\succ \tfrac{1}{2}(0_R, 1_B, 0_Y) + \tfrac{1}{2}(0_R, 1_B, 1_Y) \\ &\succ \tfrac{1}{2}(1_R, 0_B, 1_Y) + \tfrac{1}{2}(0_R, 1_B, 1_Y) = \text{bet V} \end{aligned}$$

which is a contradiction.

Definition of Directed Trees

In mathematics, a graph (N, E) is a set N of vertices or nodes
whose pairs may or may not be connected
by edges $e = (n, n')$ with $n \neq n'$
in a specified set $E \subseteq N \times N$.

The graph (N, E) is directed just in case
there is an antisymmetric and complete binary relation $>_{+1}$ on E
— i.e., for each edge $(n, n') \in E$, either $n >_{+1} n'$, or $n' >_{+1} n$.

The sequence $(n_1, n_2, \ldots, n_\ell)$ of ℓ nodes in N
is a path of length $\ell \in \mathbb{N}$ in the directed graph (N, E)
just in case $n_{k+1} >_{+1} n_k$ for $k = 1, 2, \ldots, \ell - 1$.

The graph (N, E) is a directed tree just in case there is
a unique initial node $n_0 \in N$ (or root, or seed, or entry point)
such that, for each other node $n \in N \setminus \{n_0\}$,
there is a unique path $(n_0, n_1, n_2, \ldots, n)$
which starts at node n_0 and ends at node n.

Botanical Correctness

Remark

Calling the unique initial node of a decision tree a "root" is botanically incorrect.

Real trees have entire root systems.

Succeeding Nodes and Continuation Subtrees

Let $T = (N, E)$ denote any finite directed tree.

Given any node $n \in N$, say that:

1. the node $n' \in N$ is an immediate successor of n just in case the ordered pair (n, n') is a directed edge of T, and then let

$$N_{+1}(n) := \{n' \in N \mid (n, n') \in E\} = \{n' \in N \mid n' >_{+1} n\}$$

   denote the set of all the immediate successors of n;
2. the node $n' \in N$ is a (ultimately) a successor of n
   just in case either $n' = n$,
   or there is a path $(n_1, n_2, \ldots, n_\ell)$ of nodes in N
   such that $n_1 = n$ and $n_\ell = n'$;
3. the continuation subtree $T_{\geq n} = (N_{\geq n}, E_{\geq n})$
   is the unique tree in which:
   - $N_{\geq n}$ is the set of nodes in T that succeed n;
   - $E_{\geq n}$ is the restriction $E \cap (N_{\geq n} \times N_{\geq n})$ of edges in E to directed pairs of nodes that both succeed n.

Old Definition of Decision Trees

A finite decision tree is a DRAG
with different kinds of node or vertex:

1. decision nodes n where the decision maker or agent must choose an edge emanating from n;
2. chance nodes where an edge emanating from n is determined by a roulette lottery;
3. event nodes where an edge emanating from n is determined by a horse lottery;
4. terminal nodes n from which no edge emanates, and which are mapped to a consequence $\gamma(n)$ in a specified consequence domain Y.

Remark

Raiffa focused on decision trees which lack event nodes, and whose consequence domain consists of scalars in $\mathbb{R}$ *that represent payoffs measured in dollars.*

Outline

Behaviour Rules and Continuation Subtrees

At any decision node n of any decision tree T with graph (N, E), let:

1. $N_{+1}(n)$ denote the set of nodes $n' \in N$
   that immediately succeed n
   along the unique corresponding directed edge $(n, n') \in E$
   that emanates from n and ends at n';
2. $\beta(T, n) \subseteq N_{+1}(n)$ denote the non-empty behaviour set
   of all immediately succeeding nodes that can be reached
   by deciding at n in accordance with a behaviour rule β;
3. $N_{\geq n}$ denote the set of all nodes (including n)
   that eventually succeed n in T;
4. $T_{\geq n}$ denote the continuation subtree with initial node n
   whose graph $(N_{\geq n}, E_{\geq n})$ results from the graph (N, E) of T
   by retaining only those nodes and associated edges
   that eventually succeed n in T.

Dynamic Consistency of Actual Behaviour, I

At any decision node n, the decision maker faces
the continuation subtree $T_{\geq n}$ whose initial node is n.

So possible actual behaviour at node n is, by definition,
specified by the non-empty choice set $\beta(T_{\geq n}, n) \subseteq N_{+1}(n)$.

This choice set overrules any previous plan or intention to:

1. include unacceptable nodes $\tilde{n} \in N_{+1}(n) \setminus \beta(T_{\geq n}, n)$;
2. exclude acceptable nodes $\tilde{n} \in \beta(T_{\geq n}, n)$.

Hence, in any preceding subtree $T_{\geq n'}$, which must satisfy $n \in N_{\geq n'}$,
the consistency condition $\beta(T_{\geq n'}, n) = \beta(T_{\geq n}, n)$ must hold.

Dynamic Consistency of Actual Behaviour, II

We have defined actual behaviour so that,
at any decision node n of T, it satisfies
the dynamic consistency condition $\beta(T_{\geq n'}, n) = \beta(T_{\geq n}, n)$
because both $\beta(T_{\geq n'}, n)$ and $\beta(T_{\geq n}, n)$ describe
the same non-empty subset of $N_{+1}(n)$.

Conclusion: whereas dynamic inconsistency:

1. is possible for preferences (Strotz), or plans, or intentions;
2. it is impossible for actual behaviour because,
   in any continuation decision tree that is so far unreached,
   plans are at best unrealized commitments.

Corollary: if plans or intentions are dynamically inconsistent,
then actual behaviour in practice cannot realize all these plans.

Outline

Facing up to Arrow's Less Audacious Question

William S. Vickrey (1963) *Metastatics and Macroeconomics*

Vickrey's book focuses on "metastatic" models
of perfectly competitive markets
in intertemporal general equilibrium.

These occur when each economic agent (consumer or producer)
chooses a single complete trading plan or strategy,
as in a game in normal or strategic form.

Answering Arrow's Less Audacious Question?

PJH (1977) "Dynamic Restrictions on Metastatic Choice" *Economica* 44: 337–350.

In finite decision trees for a single agent,
this "metastatic" reduction
to a game in strategic form works
if and only if the lone agent maximizes a preference ordering.

The paper builds on Kenneth J. Arrow (1959) "Rational Choice Functions and Orderings" *Economica* 26: 121–127.

It also justifies the independence of irrelevant alternatives axiom in Arrow's *Social Choice and Individual Values*.

Extensive and Strategic Form Games

Extensive form games:
Ernst Zermelo (1913) "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels".

Games in strategic or normal form:
Von Neumann, J. (1928) "Zur Theorie der Gesellschaftsspiele" *Mathematische Annalen* 100: 295–320.

Von Neumann, J. and O. Morgenstern (1944; 3rd edn. 1953) *Theory of Games and Economic Behavior* (Princeton UP).

A Game in Normal or Strategic Form

Von Neumann (1928) and von Neumann and Morgenstern (1953) gave a (slightly inadequate) definition of a game in extensive form.

They explain the "normal" or "strategic" form as a reduction of the extensive form game in which each player makes only one move.

This one move consists of announcing to an "umpire" a "strategy" or complete plan of action for the whole game.

The umpire then executes each player's announced strategy, and determines the result of the game.

A modern example would be a game of chess played between two computer programs.

Von Neumann's Normal Form Invariance

Von Neumann's (1928) major claim of normal form invariance:
it is sufficient to analyse any game in extensive form
by considering only its normal form.

This is generally invalid, but is valid for:

1. Two-person "zero sum" or "strictly competitive" games.
   These are the only games with more than one player
   for which von Neumann's claim really works;
2. One-person games, whose extensive form is a decision tree.
   For these I claim that normal form invariance
   is an essential part of the definition of pre-rationality.
3. Zero-person games, or stochastic processes,
   in which a time series is randomly determined
   as an element in a suitable function space. (Kolmogorov).

Consequentialist Normal Form Invariance

A basic principle of "structural rationality":

a consequentialist version
of von Neumann's principle of normal form invariance
applied to finite decision trees with:

1. consequences rather than payoffs
   attached to terminal nodes (or "leaves");
2. moves by the agent at decision nodes;
3. risky moves by chance ("roulette lotteries") at chance nodes,
   according to an exogenously specified
   "hypothetical" or "objective" probability distribution;
4. uncertain moves by nature ("horse lotteries") at event nodes,
   where the agent's information partition of the state space
   becomes refined,
   as in Gérard Debreu's event trees in his *Theory of Value*.

Defining Prerational Preference Relations

PJH (2022) "Prerationality as Avoiding Predictably Regrettable Consequences" *Revue Économique* 73: 937–970.
(Special Issue: In Honour of Philippe Mongin)

A preference relation $\succsim$ on the consequence domain Y is prerational just in case there exists a consequentialist choice function $\mathcal{F}(Y) \ni F \mapsto C(F) \subseteq F$ defined on the family of finite non-empty subsets of Y which:

1. describes the possible consequences of a behaviour rule

$$\mathcal{D}(Y) \ni (T, n) \mapsto \beta(T, n) \subset N_{+1}(n)$$

   that is defined on the domain $\mathcal{D}(Y)$ of all pairs (T, n) where:
   - $T \in \mathcal{T}(Y)$, the unrestricted domain of all possible finite decision trees with consequences in Y;
   - n is a decision node in T;
2. has the property that $\succsim$ is the base relation corresponding to the restriction of $F \mapsto C(F) \subseteq F$ to all pair sets $F = \{y', y''\} \subset Y$ with $y' \neq y''$.

Bernoulli and Expected Utility Functions

A Bernoulli Utility Function on Y is a mapping $Y \ni y \mapsto u(y)$.

The von Neumann–Morgenstern expected utility function
generated by the Bernoulli Utility Function $Y \ni y \mapsto u(y) \in \mathbb{R}$
is the mapping $\Delta(Y) \ni \lambda \mapsto U(\lambda) \in \mathbb{R}$
having the property that, for all $\lambda \in \Delta(Y)$,
one has $U(\lambda) = \mathbb{E}_\lambda u = \sum_{y \in Y} \lambda(y)\, u(y)$.

Equivalence Classes of Bernoulli Utility Functions

The two Bernoulli utility functions

$$Y \ni y \mapsto u(y) \in \mathbb{R} \text{ and } Y \ni y \mapsto \tilde{u}(y) \in \mathbb{R}$$

as well as the two associated
von Neumann–Morgenstern expected utility functions

$$\Delta(Y) \ni \lambda \mapsto U(\lambda) \in \mathbb{R} \text{ and } \Delta(Y) \ni \lambda \mapsto \tilde{U}(\lambda) \in \mathbb{R}$$

are cardinally equivalent just in case,
for all $y, y', y'' \in Y$ with $u(y') \neq u(y``)$,
the ratios of utility differences satisfy

$$\frac{u(y) - u(y'')}{u(y') - u(y'')} = \frac{\tilde{u}(y) - \tilde{u}(y'')}{\tilde{u}(y') - \tilde{u}(y'')}$$

— or equivalently, just in case there exist
an additive constant $\alpha \in \mathbb{R}$,
and a strictly positive multiplicative constant $\rho \in \mathbb{R}$,
such that $\tilde{u}(y) = \alpha + \rho u(y)$ for all $y \in Y$.

Characterizing Prerational Preference Relations

A (weak) preference relation $\succsim$ on a given consequence domain Y is non-trivial just in case there exist three consequences $a, b, c \in Y$ such that $a \succ b$ and $b \succ c$.

Theorem

Suppose the domain $\mathcal{T}(Y)$ includes finite decision trees with both:

1. *chance nodes where a roulette lottery is resolved;*
2. *event nodes where a horse lottery is resolved.*

Then a non-trivial base preference relation $\succsim$ is both prerational and continuous w.r.t. strictly positive objective probabilities $\pi(n'|n)$ of different "prizes" $n' \in N_{+1}(n)$ at chance nodes n of T if and only if it can be represented by the expected value, with unique positive subjective probabilities, of each Bernoulli utility function $Y \ni y \mapsto u(y) \in \mathbb{R}$ in a unique cardinal equivalence class.

Outline

Enlivenment Nodes

Tractability allows only bounded decision trees
of rather small size to be considered.

For example, a complete solution to the game of chess
is generally impossible unless one starts from an original position
with no more than 7 pieces (including both kings) on the board.

Then unmodelled events or decisions could emerge later
as new branches of an enlivened decision tree.

Tomorrow, seminar at the Sorbonne, 12:00–13:00:
"Prerationality in Enlivened Decision Trees".

The work in progress discussed there aims to extend
the principle of consequentialist pre-rationality
to a relevant domain of enlivened trees.

Such trees have a new class of enlivening nodes,
each of which is the initial node
of a previously unmodelled continuation subtree.

Outline

Quantum Nodes I

In physics and philosophy, there is:

1. an extensive older literature claiming (I think falsely) that quantum uncertainty cannot be described by a Kolmogorov probability space.

   Instead it is claimed that one needs a quantum logic of events (often in the form of an "orthocomplemented" lattice) which cannot be described by a single σ-algebra;

2. a recent literature on "quantum Bayesianism" and Ramsey/Savage subjective probability theory.

   This, *inter alia*, pays no attention to the Anscombe/Aumann distinction between roulette and horse lotteries.

Quantum Nodes II

These two literatures suggest the need
to consider quantum decision theory
based on decision trees that include quantum nodes
where uncertainty is resolved by a quantum lottery of some kind.

Presentation on "Quantum Observables, Contextual Boolean Algebras, and Bayesian Rationality in Decision Trees"
to the Workshop on the Applications of Topology
to Quantum Theory and Behavioral Economics
at the Fields Institute for Research in Mathematics,
Toronto, March 2023.

Quantum Nodes III

Main result: decomposition of a single quantum node
into a three-fold combination of:

1. first, a decision or event node that determines:
   - a quantum experiment;
   - an associated "contextual" σ-algebra of measurable events;
2. second, an event node that determines
   a suitable density operator or prior beliefs
   corresponding to a subjective probability distribution;
3. third, a chance node where the probability distribution
   of different possible observed experimental results
   in the relevant quantum context can be calculated
   by using the laws of quantum mechanics
   — especially Born's rule.

Outline

The History of a Consistent Interest in Dynamic Inconsistency

Choice Functions, Base Preferences, and Arrow's Question

Three Motivating Examples

Finite Decision Trees with Four Kinds of Node

Three New Kinds of Node in Decision Trees

- Enlivenment Nodes in Enlivened Decision Trees
- Quantum Nodes for Describing Quantum Uncertainty
- Consequence Nodes and Menu Consequences

Decision Trees with Intermediate Consequence Nodes

Beyond Deathbed Consequentialism

The old definition of decision tree allowed consequences to be attached only to terminal nodes.

This is a terminal or "deathbed" form of consequentialism, in which, for example, an individual's entire life history is evaluated only at a time when life is about to end.

We might want to go beyond terminal consequentialism.

For example, we might want to accommodate more directly consumers' decision trees in macroeconomics, where consequences are intertemporal consumption streams.

So we shall allow intermediate consequence nodes which have:

1. a unique immediate successor (so no branching occurs);
2. a specified consequence, which could even be a continuation decision "subtree" whose initial node is the consequence node;
3. a specified time at which the consequence occurs.

Outline

Consequence Nodes

So far the only consequence nodes in a decision tree T are terminal nodes N for which $N_{+1}(n) = \emptyset$.

We now allow a decision tree T to include additional non-terminal (or "intermediate") consequence nodes.

Each intermediate consequence node n has:

1. a unique immediate successor $N_{+1}(n)$ (so at n there is no branching due to a non-trivial decision, or because either a roulette lottery or a horse lottery has been resolved);
2. a specified consequence $\gamma(n) \in Y$, which could include even the continuation decision subtree $T_{\geq n}$ whose initial node is the consequence node;
3. a specified time $t(n) \in \mathbb{R}$ at which the consequence occurs.

A Canonical Example from Macro and Finance

Example

In the canonical discrete-time macroeconomic model
of consumption and portfolio choice, at each successive time $t \in \mathbb{N}$,
several things happen simultaneously:

1. the consumer's asset portfolio $\mathbf{a}_{t-1} \in \mathbb{R}^m$ just before time t
   takes the form of a multi-dimensional lottery
   whose random "prize" determines a new wealth level w_t;
2. the consumer chooses to allocate $w_t \in \mathbb{R}$
   between current consumption expenditure $c_t \in \mathbb{R}$
   and an updated portfolio of assets $\mathbf{a}_t \in \mathbb{R}^m$;
3. so there is:
   - first, a chance node that determines $w_t \in \mathbb{R}$;
   - second, a decision node that determines
     the pair $(\mathbf{a}_t, c_t) \in \mathbb{R}^{m+1}$;
   - third and last, a timed consequence node (t, c_t)
     — or possibly $(t, c_t, \mathbf{a}_t)$. □

Streams of Timed Consequences

In a decision tree $T = (N, E)$, recall the definition of a path of length $\ell \in \mathbb{N}$ as a sequence $(n_1, n_2, \ldots, n_\ell)$ of ℓ nodes in the directed graph (N, E) that satisfy $(n_k, n_{k+1}) \in E$ or $n_{k+1} >_{+1} n_k$ for $k = 1, 2, \ldots, \ell - 1$.

In case the decision tree T has consequence nodes, any path $(n_1, n_2, \ldots, n_\ell)$ in the tree includes a subsequence of $m \leq \ell$ consequence nodes, and so a finite sequence or stream $\mathbf{y} = (t_j, y_j)_{j=1}^m$ of timed consequences $(t_j, y_j) \in \mathbb{R} \times Y$ whose successive times satisfy $t_{j+1} > t_j$ for $j = 1, 2, \ldots, m - 1$.

We assume that every terminal node is a timed consequence node.

Domains of Lotteries and of Finite Decision Trees

Given the consequence domain Y, let:

1. $\mathbf{Y}(Y)$ denote the domain of all possible timed consequence streams $\mathbf{y} = (t_j, y_j)_{j=1}^m$ that meet the previous definition;
2. $\mathbf{L}(Y)$ denote the domain of all possible Anscombe/Aumann roulette lotteries $\boldsymbol{\lambda}$ whose prizes are consequence streams in $\mathbf{Y}(Y)$;
3. $\mathfrak{T}(Y) = \mathcal{T}(\mathbf{Y}(Y))$ denote the domain of all finite decision trees with:
    - chance nodes where a roulette lottery is resolved;
    - event nodes where a horse lottery is resolved;
    - timed consequence nodes (t, y) (including all terminal nodes).

Evidently the extended consequence of following any decision strategy $\mathbf{d}$ in a decision tree $T \in \mathfrak{T}(Y)$ is an Anscombe/Aumann roulette lottery $\boldsymbol{\lambda}(\mathbf{d}) \in \mathbf{L}(Y)$.

Characterizing Prerational Preference Relations

The following result is immediately implied by our earlier theorem applied to the new "stream domain" $\mathfrak{T}(Y)$ of decision trees that have an Anscombe/Aumann roulette lottery $\mathbf{l}(n) \in \mathbf{L}(Y)$ over consequence streams attached to each terminal node of T.

Theorem

Given the domain $\mathfrak{T}(Y)$ of finite decision trees T,
a non-trivial base preference relation $\succsim$
over the associated domain $\mathbf{L}(Y)$
of Anscombe/Aumann roulette lotteries
is both prerational and continuous
w.r.t. strictly positive probabilities $\pi(n'|n)$
of different roulette lottery outcomes
if and only if it can be represented by the expected value,
with unique positive probabilities of horse lottery outcomes,
of each Bernoulli utility function $\mathbf{Y}(Y) \ni \mathbf{y} \mapsto u(\mathbf{y}) \in \mathbb{R}$
in a unique cardinal equivalence class.

Histories and Continuations of Consequence Streams

Given any finite decision tree $T \in \mathfrak{T}(Y)$ and any time $t \in \mathbb{R}$,
consider any consequence stream $\mathbf{y} = (t_j, y_j)_{j=1}^m$
that passes through the stream $\mathbf{n} = (n_j)_{j=1}^m$
of timed consequence nodes.

Given any decision node n of T,
the consequence stream $\mathbf{y}$ can be split into:

1. a history $\mathbf{h}_{<n}$ of timed consequences that precede n;
2. a continuation $\mathbf{y}_{>n}$ of timed consequences that succeed n.

History Dependent Preferences and Utility Functions

Any base preference relation $\succsim$ over the domain $\mathbf{L}(\mathbf{Y})$
of consequence lottery streams $\boldsymbol{\lambda}$ induces,
conditional on each history $\mathbf{h}_{<n}$,
a history-dependent conditional base preference relation $\succsim_{\mathbf{h}_{<n}}$
on the subdomain $\mathbf{L}_{>n}(\mathbf{Y})$
of continuation consequence lottery streams
that can arise after passing through n.

In case the base preference relation $\succsim$ over the domain $\mathbf{L}(\mathbf{Y})$
can be represented by the expected value
of the Bernoulli utility function $\mathbf{Y}(Y) \ni \mathbf{y} \mapsto u(\mathbf{y}) \in \mathbb{R}$,
then the conditional base preference relation $\succsim_{\mathbf{h}_{<n}}$
on the continuation subdomain $\mathbf{L}_{>n}(\mathbf{Y})$
can be represented by the expected value
of the parameterized conditional Bernoulli utility function

$$\mathbf{Y}_{>n}(Y) \ni \mathbf{y}_{>n} \mapsto u(\mathbf{h}_{<n}; \mathbf{y}_{>n}) \in \mathbb{R}$$

From Dynamic Programming to Behaviour

Given any finite decision tree $T \in \mathfrak{T}(Y)$,
to find the appropriate non-empty behaviour set $\beta(T, n) \subseteq N_{+1}$
use the idea of backward recursion in dynamic programming.

This requires starting at terminal nodes and their consequences,
and then determining the history valuation function
and associated optimal behaviour set
by solving the backward recurrence relation

$$\left.\begin{array}{rcl} V_n^T(\mathbf{h}_{<n}) & = & \\ \beta(T, n) & = & \arg \end{array}\right\} \max_{n' \in N_{+1}(n)} \left\{ \mathbb{E}_{n'} \left[V_{n'}^T(\mathbf{h}_{<n}, n') \right] \right\}$$

Please forgive the rather serious abuse of notation which fails,
for instance, to distinguish different kinds of node ...

History Independent Recursive Preferences

The base preference relation $\succsim$ over the domain $\mathbf{L}(\mathbf{Y})$
is history independent just in case
each conditional base preference relation $\succsim_{\mathbf{h}_{<n}}$ takes the form $\succsim_n$,
independent of what timed consequences precede n.

Similarly for the Bernoulli utility functions
whose expected values represent base preferences.

In this case both the base preference relations
and the associated Bernoulli utility functions
must be recursively separable, by definition.

Outline

Work in Progress, with Agustín Troccoli Moretti

Prospective title of joint paper:
Prerationality in Decision Trees that Include Menu Consequences

Precursors: Static Menus as Consequences

From the literature on social choice with rights:

PJH (1995) "Social Choice of Individual and Group Rights"
in W.A. Barnett, H. Moulin, M. Salles, and N. Schofield (eds.)
Social Choice, Welfare, and Ethics
(Cambridge University Press), ch. 3, pp. 55–77.

PJH (1997) "Game Forms versus Social Choice Rules
as Models of Rights" in Arrow, Sen, and Suzumura (eds.)
Social Choice Re-examined, Vol. II (London: Macmillan, 1997)
ch. 11, pp. 82–95.

Precursors: Decision Subtrees as Consequences

More relevant literature on decision trees
with menus included as parts of consequences:

Tjalling C. Koopmans (1964) "On Flexibility of Future Preference"
pp. 243–254 of M. W. Shelly, II, and G. W. Bryan (eds.)
Human Judgments and Optimality (New York: John Wiley).

David M. Kreps and Evan L. Porteus (1978)
"Temporal Resolution of Uncertainty and Dynamic Choice Theory"
Econometrica 46 (1): 185–200.

Marciano Siniscalchi (2011) "Dynamic Choice under Ambiguity"
Theoretical Economics 6: 379–421.

Mark Machina and Marciano Siniscalchi (2014)
"Chapter 13 — Ambiguity and Ambiguity Aversion" pp. 729–807
of *Handbook of the Economics of Risk and Uncertainty, Volume 1*.

Another Precursor

PJH (1989)
"Consistent Plans, Consequentialism, and Expected Utility"
Econometrica 57: 1445–1449.

Uses an argument like that involved in menu consequences
in order to dispute some contentious claims that were made in:

Johnsen, T. H., and J. B. Donaldson (1985)
"The Structure of Intertemporal Preferences Under Uncertainty and Time Consistent Plans" *Econometrica* 53: 1451–1458.

These claims were that maximizing a non-expected utility function will be time consistent even outside
a restricted domain of "macroeconomic" decision trees.

My response has also helped me to realize much later
the need to formalize the notion of a timed consequence node.

The Potential Addict Example Reconsidered

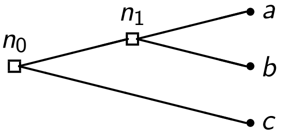

Figure: The potential addict's decision tree

Recall the "essentially inconsistent" base preferences,
with $b \succ_0 c \succ_0 a$ at node n_0, and $a \succ_1 b$ at node n_1.

Add a consequence node $n_{1/2}$ on the edge from n_0 to n_1,
with the timed menu consequence $(\frac{1}{2}, \{a, b\})$.

Then the "myopically addictive" behaviour norm
with $\beta(T, n_0) = \{n_1\}$ and $\beta(T, n_1) = \{a\}$ maximizes
any preference ordering over consequence streams that satisfies

$$\big((\tfrac{1}{2}, \{a, b\}), a\big) \succ \big((\tfrac{1}{2}, \{a, b\}), b\big) \text{ and } \big((\tfrac{1}{2}, \{a, b\}), a\big) \succ c$$

Rationalizing Allais's Dependent Preferences

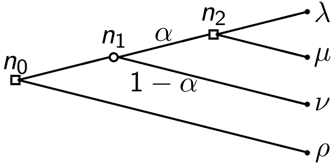

Suppose $\lambda \succ \mu$ and yet $\alpha\mu + (1-\alpha)\nu \succ \rho \succ \alpha\lambda + (1-\alpha)\nu$.

Add a consequence node $n_{1/2}$ on the edge from n_0 to n_1,
with the menu consequence $M := \{\alpha\lambda + (1-\alpha)\nu, \alpha\mu + (1-\alpha)\nu\}$.

Then the "myopically addictive" behaviour norm
with $\beta(T, n_0) = \{n_1\}$ and $\beta(T, n_2) = \{\lambda\}$
maximizes any preference ordering $\succsim$
over consequence lottery streams that satisfies

$$\big((\tfrac{1}{2}, M), \alpha\lambda + (1-\alpha)\nu\big) \succ \big((\tfrac{1}{2}, M), \alpha\mu + (1-\alpha)\nu\big)$$

and $\big((\tfrac{1}{2}, M), \alpha\lambda + (1-\alpha)\nu\big) \succ \rho$

Expected Utility Implies Non-Expected Utility!

Consider the space D
of cumulative distribution functions $\mathbb{R} \ni y \mapsto F(y) \in [0, 1]$
representing monetary lotteries.

Following Allais and Machina in particular,
consider general non-expected utility preferences $\succsim$
represented by a utility function $D \ni F \mapsto U(F) \in \mathbb{R}$.

Then $\succsim$ is also represented by the expected value $\mathbb{E}V(F, y)$
of the Bernoulli utility function $D \times \mathbb{R} \ni (F, y) \mapsto V(F, y)$
of two variables, provided we focus on the "menu consequence" F,
and define $V(F, y) := U(F)$, independent of y, for all $y \in \mathbb{R}$.

Characterizing Prerational Preference Relations, I

Appropriate streams $\mathbf{M}$ of non-empty finite subsets M
of the extended consequence domain $\mathbf{Y}(Y)$
can be regarded as menu consequence streams.

Such streams $\mathbf{M}$ will be additional arguments
of a Bergson utility function $u(\mathbf{y}, \mathbf{M})$
defined on consequence streams,
including menu consequence streams.

Characterizing Prerational Preference Relations, II

Conjecture

Consider any finite decision tree $T \in \mathcal{T}(Y)$
with consequences in the domain Y
and any arbitrary behaviour norm $N^d \ni n \twoheadrightarrow \beta(T, n)$ in T
whose possible consequences in the particular tree T
are given by $\Phi_\beta(T)$.

Then there exists an extension T^* of T with menu consequences
and a Bernoulli utility function u^*
over consequence streams with menu consequences
such that the expected utility maximizing behaviour norm β^* in T^*
has $\Phi_\beta(T)$ as the possible set of restrictions
to non-menu consequences
of the set $\Phi_{\beta^*}(T^*)$ of possible extended consequences.

Conclusions: New Kinds of Node in Decision Trees

The finite decision trees that have generally been considered up to now have four kinds of node:
(i) decision nodes; (ii) roulette lottery nodes;
(iii) horse lottery nodes; (iv) terminal consequence nodes.

Now we have three new kinds of node:

(v) enlivenment nodes,
to be considered in tomorrow's talk at the Sorbonne;
(vi) quantum nodes, which I claim can be reduced,
at least in many of the cases considered in physics textbooks,
to combinations of the first four kinds of node;
(vii) (timed) consequence nodes,
which were discussed toward the end of this talk.

Intermediate Consequence Nodes

In the absence of menu consequences,
any feasible sequence of timed consequence nodes
can be replaced by a single consequence stream
which is attached to a terminal node.

Menu consequences, however, seem to have the potential
to allow major modifications to standard results
concerning expected utility maximizing behaviour.

Conclusion: when we consider finite decision trees
which have intermediate consequence nodes
that include menu consequences,
there remains much work to do.

Envoi

Many thanks for the kind invitation and for your attention.