CES -161 - Modelos Probabilísticos em Grafos

[Rational] Decisions with Bayesian Networks

Prof. Paulo André Castro
pauloac@ita.br
www.comp.ita.br/~pauloac Sala 110, IEC-ITA

Decisões [Racionais] com Redes Bayesianas

- Preferências Racionais
- Utilidades x Dinheiro
- Redes de Decisão
- Classificação e Avaliação de classificadores

Lotteries

An agent chooses among prizes (A, B, etc.) and lotteries, i.e., situations with uncertain prizes

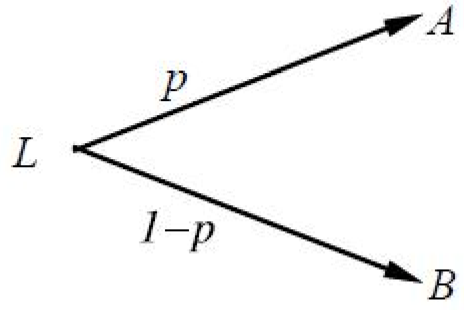

Lottery $L = [p, A;\ (1-p), B]$

Notation:

$A \succ B$	A preferred to B
$A \sim B$	indifference between A and B
$A \succsim B$	B not preferred to A

Preferências Racionais

Idea: preferences of a rational agent must obey constraints.
Rational preferences $\Rightarrow$
behavior describable as maximization of expected utility

Constraints:

Orderability

$$(A \succ B) \vee (B \succ A) \vee (A \sim B)$$

Transitivity

$$(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$$

Continuity

$$A \succ B \succ C \Rightarrow \exists p \; [p, A; \; 1-p, C] \sim B$$

Substitutability

$$A \sim B \Rightarrow [p, A; \; 1-p, C] \sim [p, B; 1-p, C]$$

Monotonicity

$$A \succ B \Rightarrow ([p, A; \; 1-p, B] \succsim [q, A; \; 1-q, B] \Leftrightarrow \; p \geq q)$$

Violação das Restrições leva a "Irracionalidade"

For example: an agent with intransitive preferences can be induced to give away all its money

If $B \succ C$, then an agent who has C would pay (say) 1 cent to get B

If $A \succ B$, then an agent who has B would pay (say) 1 cent to get A

If $C \succ A$, then an agent who has A would pay (say) 1 cent to get C

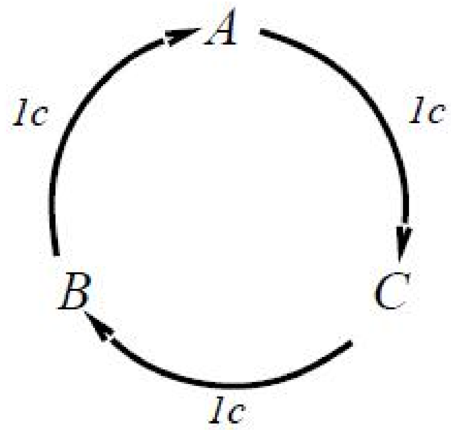

Maximizing Expected Utility (MEU)

Theorem (Ramsey, 1931; von Neumann and Morgenstern, 1944):
Given preferences satisfying the constraints
there exists a real-valued function U such that

$$U(A) \geq U(B) \quad \Leftrightarrow \quad A \succsim B$$

$$U([p_1, S_1; \ \ldots \ ; \ p_n, S_n]) = \Sigma_i \ p_i U(S_i)$$

MEU principle:

Choose the action that maximizes expected utility

Note: an agent can be entirely rational (consistent with MEU) without ever representing or manipulating utilities and probabilities

E.g., a lookup table for perfect tictactoe

Estimando Utilidades

Utilities map states to real numbers. Which numbers?

Standard approach to assessment of human utilities:
compare a given state A to a standard lottery L_p that has
"best possible prize" $u_\top$ with probability p
"worst possible catastrophe" $u_\bot$ with probability $(1-p)$
adjust lottery probability p until $A \sim L_p$

pay $30 ~

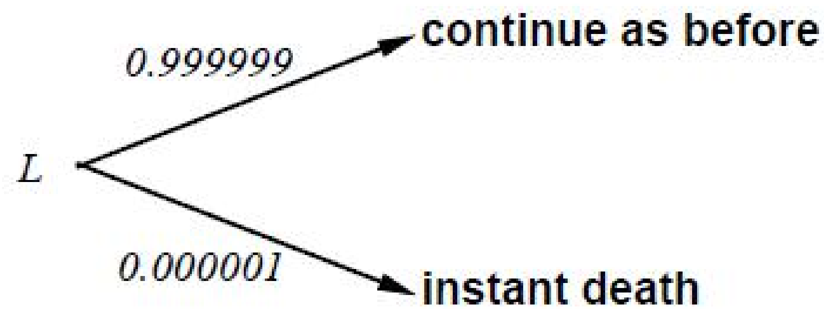

Definindo Funções de Utilidades através de loterias

- Dado o intervalo [0,1] entre a "pior catastrofe possível" e "o melhor prêmio possível", ao encontrar uma loteria [p,1;1-p,0] que seja indiferente a uma situação S o número p é a utilidade de S

- Em ambientes, com prêmios determinísticos pode-se apenas estabelecer a ordem de preferências, nesse caso usa-se o termo utilidades ordinais

- Funções de utilidades ordinais podem ser chamadas de funções de valor e são invariantes para qualquer transformação monotônica

Utilidade do dinheiro

- Preferências de um grupo sobre dinheiro certo (x) e loteria [p,M;1-p,0]

For each x, adjust p until half the class votes for lottery (M=10,000)

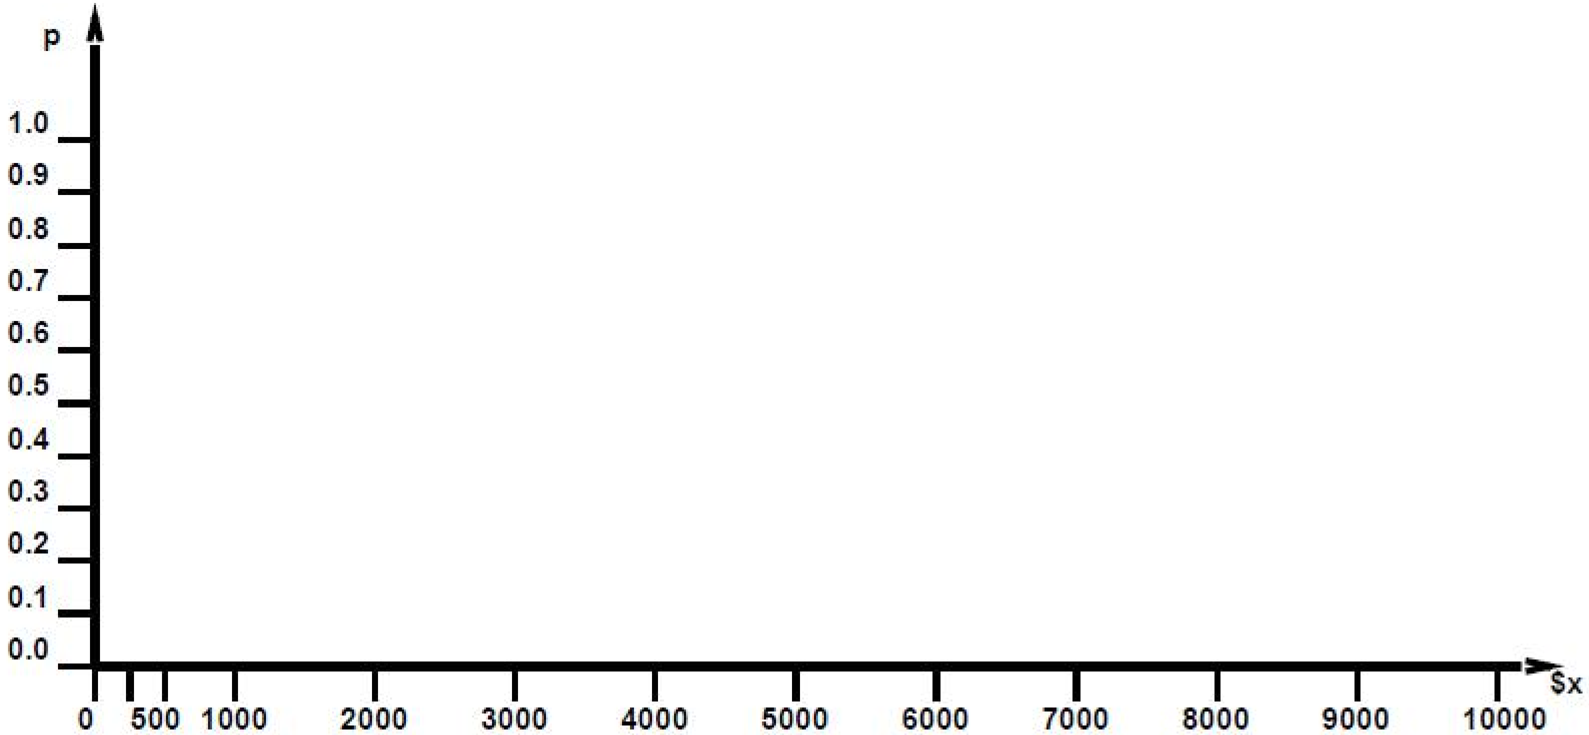

Dinheiro vs Utilidade

- Dinheiro não tem uma relação linear (ou simples) com utilidade!
- Ao estimar a utilidade em vários experimentos, observa-se que dada uma loteria L com valor esperado EMV(L) tem-se U(L) < U (EMV(L)), isto é as pessoas são aversas a risco
- Um gráfico típico de dinheiro ($) vs Utilidade (U):

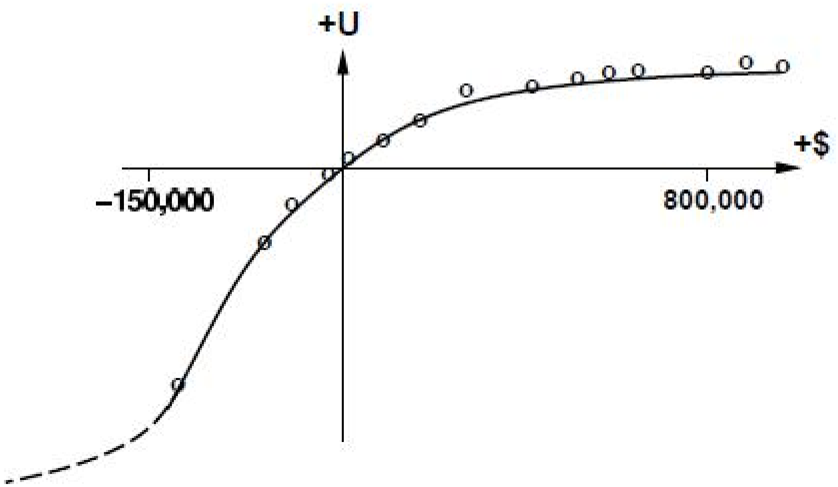

The Saint Petersburg Paradox

- The paradox is named from Daniel Bernoulli's presentation of the problem and his solution, published in 1738 in St. Petersburg

- A casino offers a game of chance for a single player in which a fair coin is tossed at each stage. The pot starts at 1 dollar and is doubled every time a head appears. The first time a tail appears, the game ends and the player wins whatever is in the pot.

- Thus the player wins 1 dollar if a tail appears on the first toss, 2 dollars if a head appears on the first toss and a tail on the second.
- Two questions:
  - How much would you accept to pay for playing this game?
  - What is the expected monetary value of the game?

The Saint Petersburg Paradox

- As Bernoulli stated:
  - The determination of the value of an item must not be based on the price, but rather on the utility it yields…. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount
- Bernoulli proposed that utility of money should be logarithmic. U(M)= a*log2(M)+b
- This makes EMV to be a finite value.

- But it's always possible to recreate the paradox by changing the function!!!
  - Alternative theories may provide a better description model (*Prospect Theory*)

Problemas na Teoria da maximização da utilidade esperada

- A teoria da maximização da utilidade esperada é uma teoria normativa. Ela descreve como um agente deve reagir. Entretanto, não é uma teoria descritiva da tomada de decisões reais

- Há evidências experimentais que as pessoas violam os axiomas da teoria da utilidade

Escolha A ou B

- A: 80% de chance de ganhar $4000
- B: 100% de chance de ganhar $3.000

Escolha C ou D

- C: 20% de chance de ganhar $4000
- D: 25% de chance de ganhar $3. 000

Supondo U(0)=0

- Se maioria escolhe B em detrimento de A e C em detrimento de D,
  - De A e B, temos que $0,8*U(4000) < U(3000)$
  - De C e D temos que $0,8U(4000) > U(3000)$
- Contraditório!!!!

Teorias alternativas

- Em linhas gerais as pessoas divergem da teoria da maximização da utilidade esperada em situações de probabilidade muito alta e/ou muito baixa

- Há algumas teorias alternativas que se propõem a descrever o comportamento humano real. Uma das mais relevantes foi proposta por Kahneman e Tversky. Esta teoria propõe um modelo alternativo que descreve esse efeito "certeza" e outros

Decisões [Racionais] com Redes Bayesianas

- Preferências Racionais

- Utilidades x Dinheiro

- **Redes de Decisão**

- Classificação e Avaliação de classificadores

Decision Networks

- By now we know how to use Bayesian networks to represent uncertainty and do probabilistic inference.

- Now, we extend them to support decision making adding an explicit representation of both the **actions** under consideration and the value or **utility** of the resultant outcomes gives us **decision networks** (also called **influence diagrams** by Howard and Matheson, 1981).

- Bayesian decision networks combine probabilistic reasoning with utilities, helping us to make decisions that **maximize the expected utility**

Principle of Expected Utility

- The principle of maximum expected utility asserts that an essential part of the nature of rational agents is to choose that action which maximizes expected utility.

$$EU(A|E) = \sum_i P(O_i|E,A)\, U(O_i|A)$$

where

- E is the available evidence,
- A is a non-deterministic action with possible outcome states O_i,
- $U(O_i|A)$ is the utility of each of the outcome states, given that action A is taken,
- $P(O_i|E,A)$ is the conditional probability distribution over the possible outcome states, given that evidence E is observed and action A taken.

- In some cases $U(O_i \mid A) = U(O_i)$

Decision Network Node types

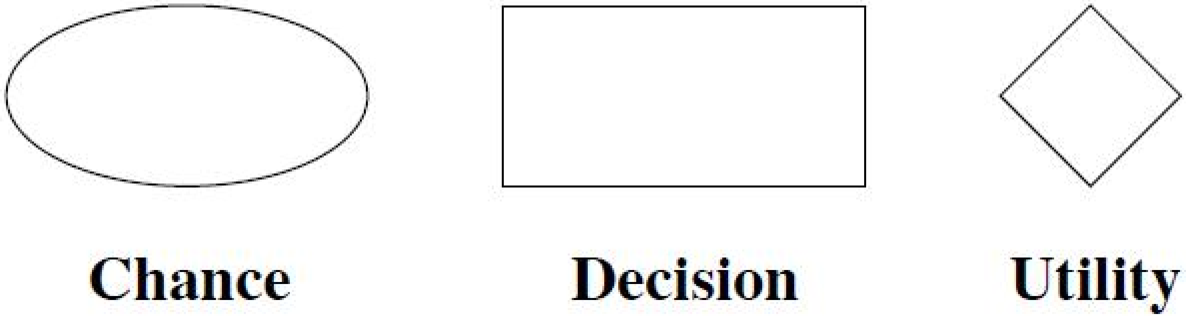

Decision network node types.

Rede de Decisão

- Nós de acaso: (elipses) representam variáveis aleatórias. Cada nó de acaso tem um distribuição de probabilidade condicional (dados os nós pais)

- Nós de Decisão : (retângulos) representam as possíveis ações

- Nós de utilidade: (losangos) representam as preferências do agente e podem ser usadas para definir as ações através da seleção da ação que maximiza a utilidade esperada

Football team example

Clare's football team, Melbourne, is going to play her friend John's team, Carlton. John offers Clare a friendly bet: whoever's team loses will buy the wine next time they go out for dinner. They never spend more than $15 on wine when they eat out. When deciding whether to accept this bet, Clare will have to assess her team's chances of winning (which will vary according to the weather on the day). She also knows that she will be happy if her team wins and miserable if her team loses, regardless of the bet.

- Assume that there is 30% of raining and in that case Melbourne has 60% of wining, but only 25% if it is not wet…
- Make your decision network!
  - Variables, conditional dependence, action and utility node

Football team example

Clare's football team, Melbourne, is going to play her friend John's team, Carlton. John offers Clare a friendly bet: whoever's team loses will buy the wine next time they go out for dinner. They never spend more than $15 on wine when they eat out. When deciding whether to accept this bet, Clare will have to assess her team's chances of winning (which will vary according to the weather on the day). She also knows that she will be happy if her team wins and miserable if her team loses, regardless of the bet.

- Possible solution…

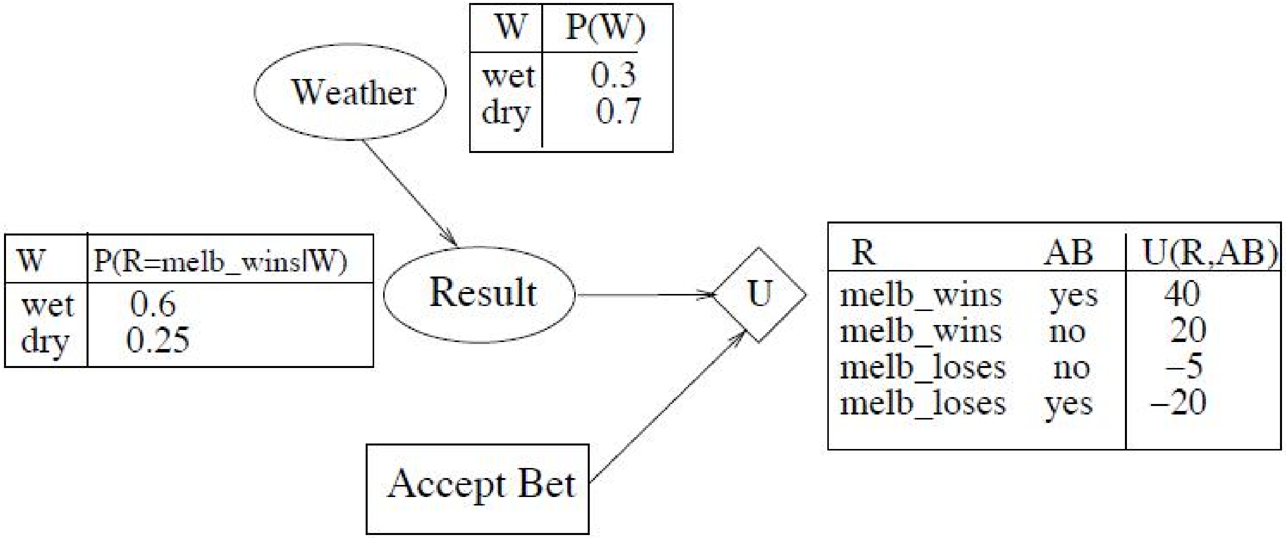

W	P(W)
wet	0.3
dry	0.7

W	P(R=melb_wins|W)
wet	0.6
dry	0.25

R	AB	U(R,AB)
melb_wins	yes	40
melb_wins	no	20
melb_loses	no	–5
melb_loses	yes	–20

Evaluating Decision Networks

To evaluate a decision network with a single decision node:

Decision Network Evaluation Algorithm (Single decision)

1. *Add any available evidence.*
2. *For each action value in the decision node:*
   (a) *Set the decision node to that value;*
   (b) *Calculate the posterior probabilities for the parent nodes of the utility node, as for Bayesian networks, using a standard inference algorithm;*
   (c) *Calculate the resulting expected utility for the action.*
3. *Return the action with the highest expected utility.*

Estimating Expected Utility

- Without evidence added, the probability of Melbourne winning is

$$P(R = melb\_wins) = P(W = w) \times P(R = melb\_wins|W = w) + P(W = d) \times P(R = melb\_wins|W = d)$$

- And:

$$P(R = melb\_loses|E) = 1 - P(R = melb\_wins|E).$$

$$\begin{aligned} EU(AB = yes) &= P(R = melb\_wins) \times U(R = melb\_wins|AB = yes) \\ &+ P(R = melb\_loses) \times U(R = melb\_loses|AB = yes) \\ &= (0.3 \times 0.6 + 0.7 \times 0.25) \times 40 + (0.3 \times 0.4 + 0.7 \times 0.75) \times -20 \\ &= 0.355 \times 40 + 0.645 \times -20 = 14.2 - 12.9 = 1.3 \end{aligned}$$

$$\begin{aligned} EU(AB = no) &= P(R = melb\_wins) \times U(R = melb\_wins|AB = no) \\ &+ P(R = melb\_loses) \times U(R = melb\_loses|AB = no) \\ &= (0.3 \times 0.6 + 0.7 \times 0.25) \times 20 + (0.3 \times 0.4 + 0.7 \times 0.75) \times -5 \\ &= 0.355 \times 20 + 0.645 \times -5 = 7.1 - 3.225 = 3.875 \end{aligned}$$

Outro Exemplo: Escolha da localização de um aeroporto

- Dependendo da posição pode-se alterar:
  - O risco de acidentes (logo, o número esperado de mortes.. **Deaths**
  - O incômodo causado pelo barulho dos aviões, quanto mais próximo de uma cidade pior….**Noise**
  - É fácil perceber que **Deaths** e **Noise** serão diretamente afetados pelo volume de **tráfego aéreo** no aeroporto.
  - Naturalmente, o custo também é alterado pela localização do aeroporto (Cost)
    - a desapropriação de um determinado terreno pode ser mais ou menos litigioso….e os custos de ligação de transportes do aeroporto a cidade podem ser maiores ou menores afetando a **construção**

Decision Network

Add action nodes and utility nodes to belief networks to enable rational decision making

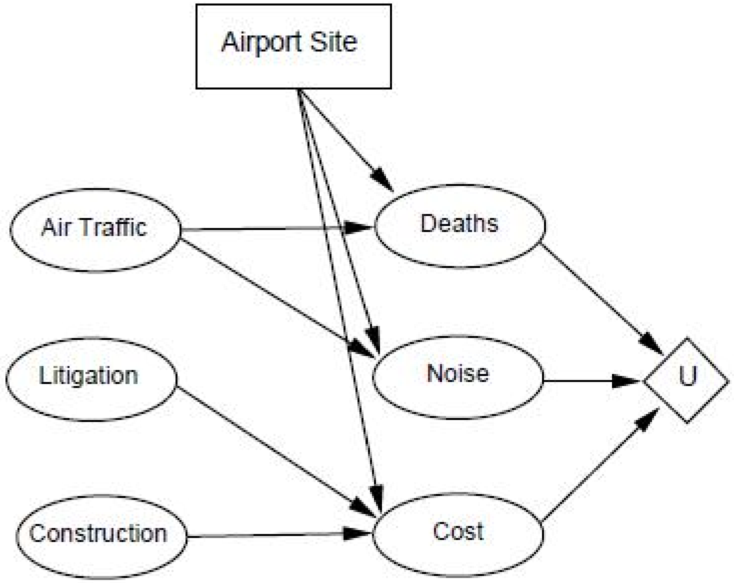

Algorithm:

For each value of action node
    compute expected value of utility node given action, evidence
Return MEU action

Podemos determinar distribuições de variáveis, mas como decidir?

- Nós de ação e nós de utilidade na rede;

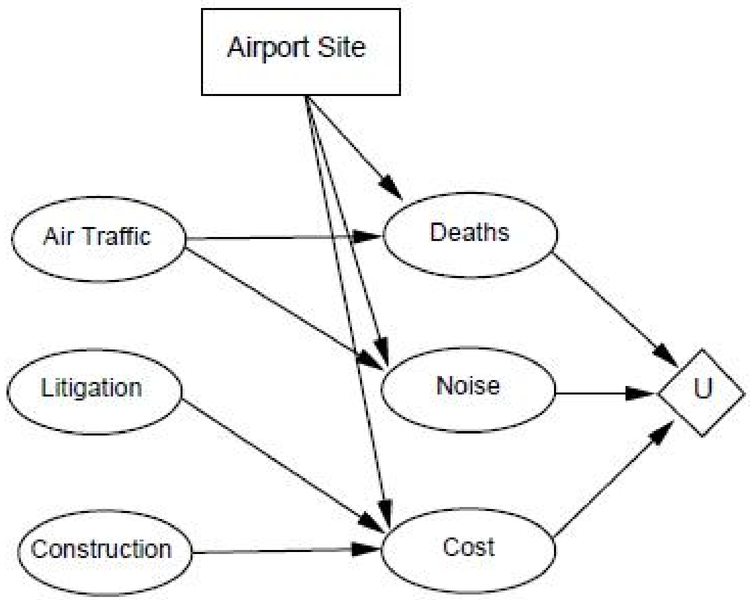

Optimal decisions: decision networks include utility information;
probabilistic inference required for $P(outcome|act\ on, ev\ dence)$

Processo de Decisão...

- pi = P(deaths=i| ASite=s, Noise=n) Ou
- P(outcome| action, evidence)

- Utilidade Esperada (action=a) = $\sum_i$ U($outcome_i$)*P($outcome_i$|action=a, evidence)
- Escolher ação que maximiza a utilidade esperada

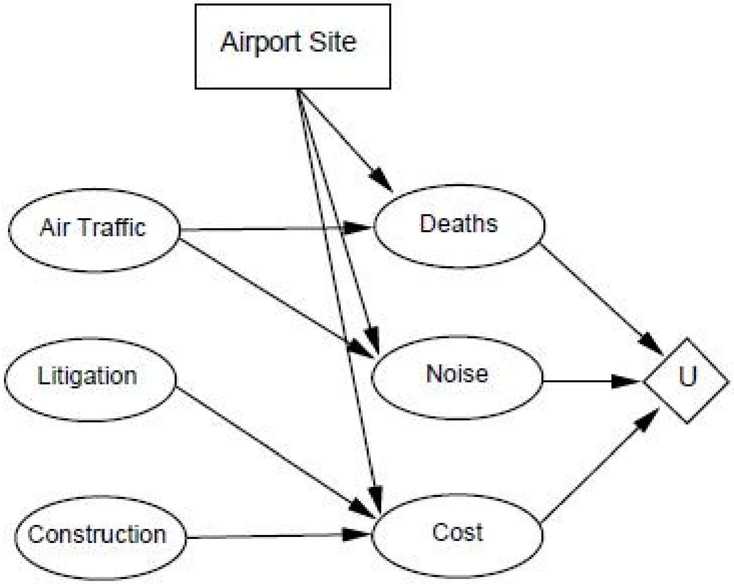

Decisões com Redes Bayesianas

- Preferências Racionais

- Utilidades x Dinheiro

- Redes de Decisão
  - Redes de Decisão e Decisão Sequenciada

- Modelo Decisório de Markov

Information Links: When a variable need to be observed

- There may be arcs from chance nodes to decision nodes — these are called information links (Jensen and Nielsen, 2007, p. 305).
- These links indicate when a chance node needs to be observed before the decision D is made — but after any decisions prior to D.
- With an information link in place, network evaluation can be extended to calculate explicitly what decision should be made, given the different values for that chance node.

Information Link: Example

- To illustrate the use of information links, let’ s use again the Melbourne football team example, but now Clare was only going to decide whether to accept the bet or not after she heard the weather forecast
- Previously :

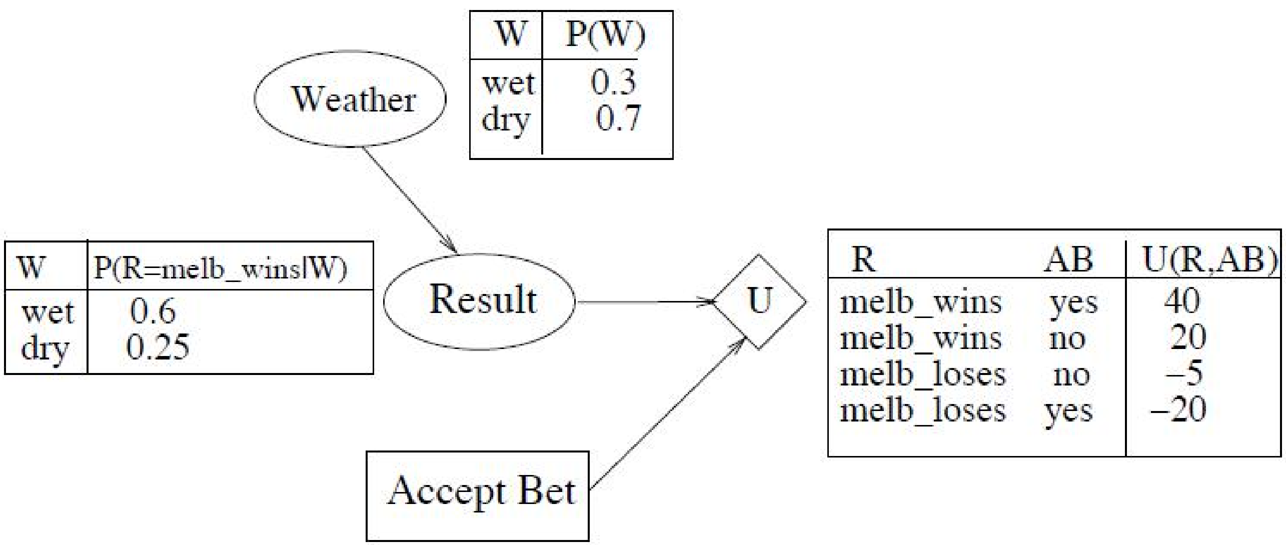

Previously Clare decided don't accept the bet

- What if there is a reasonable Forecast about the weather? How to change the network?

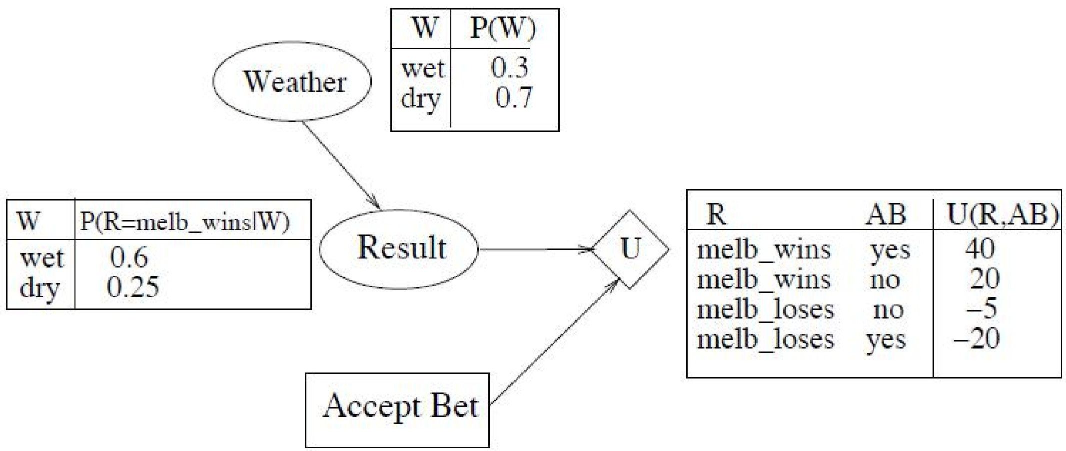

W	P(W)
wet	0.3
dry	0.7

W	P(R=melb_wins|W)
wet	0.6
dry	0.25

R	AB	U(R,AB)
melb_wins	yes	40
melb_wins	no	20
melb_loses	no	–5
melb_loses	yes	–20

$$EU(AB=yes) = P(R=melb\_wins) \times U(R=melb\_wins|AB=yes) + P(R=melb\_loses) \times U(R=melb\_loses|AB=yes) = 1.3$$

$$EU(AB=no) = P(R=melb\_wins) \times U(R=melb\_wins|AB=no) + P(R=melb\_loses) \times U(R=melb\_loses|AB=no) = 3.875$$

Introducing the New Node, Information Link and Decision Table

W	F	P(F|W)
wet	rainy	0.60
	cloudy	0.25
	sunny	0.15
dry	rainy	0.10
	cloudy	0.40
	sunny	0.50

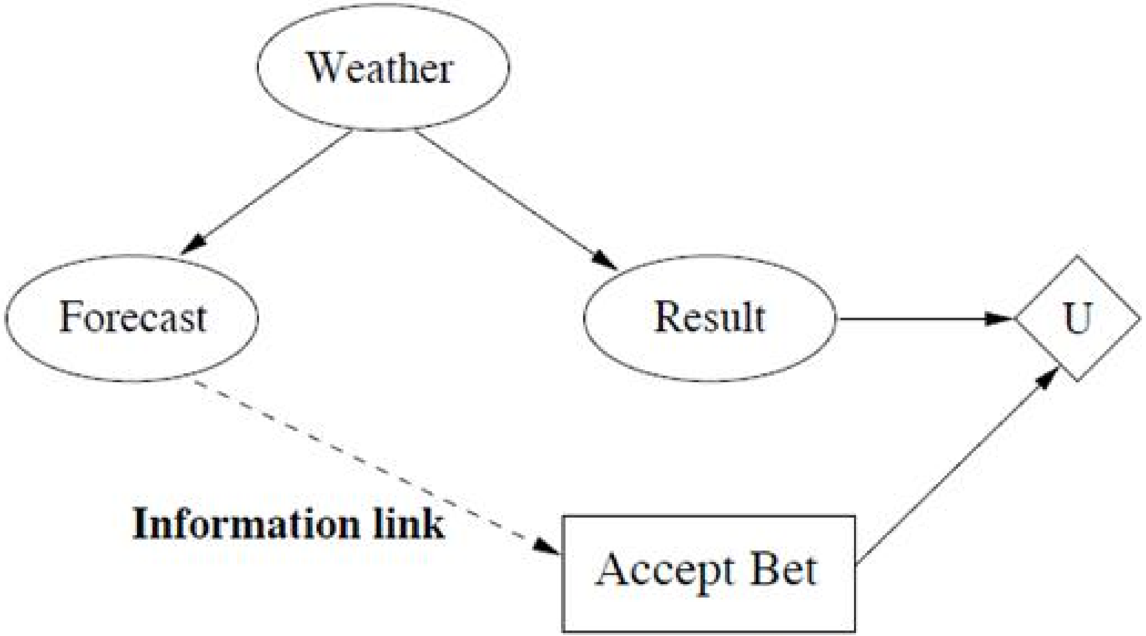

Decision Table

F	Accept Bet
rainy	**?**
cloudy	**?**
sunny	**?**

Decision Table Algorithm (single decision node, with information Link)

1. *Add any available evidence.*
2. *For each combination of values of the parents of the decision node:*
   (a) *For each action value in the decision node:*
      i. *Set the decision node to that value;*
      ii. *Calculate the posterior probabilities for the parent nodes of the utility node, as for Bayesian networks, using a standard inference algorithm;*
      iii. *Calculate the resulting expected utility for the action.*
   (b) *Record the action with the highest expected utility in the decision table.*
3. *Return the decision table.*

Decision Table obtained by following the algorithm

- **EU(AB=yes | F=rainy)** = P(R=melb_wins| F=rainy) x U(R=melb_wins| AB=yes,F=rainy)+ P(R=melb_loses| F=rainy) x U(R=melb_loses| AB=yes,F=rainy)
- **EU(AB=no | F=rainy)** = P(R=melb_wins| F=rainy) x U(R=melb_wins| AB=no, F=rainy)+ P(R=melb_loses| F=rainy) x U(R=melb_loses| AB=no, F=rainy)
- Analogus to other values of F…..

Decisions calculated for football team, given the new evidence node *Forecast*.

F	*Bel*(*W* = *wet*)	*Bel*(*R* = *melb_wins*)	EU(AB=yes)	EU(AB=no)
rainy	0.720	0.502	**10.12**	7.55
cloudy	0.211	0.324	-0.56	**3.10**
sunny	0.114	0.290	-2.61	**2.25**

Decision Table with recorded actions (highest expected utility)

W	F	P(F|W)
wet	rainy	0.60
	cloudy	0.25
	sunny	0.15
dry	rainy	0.10
	cloudy	0.40
	sunny	0.50

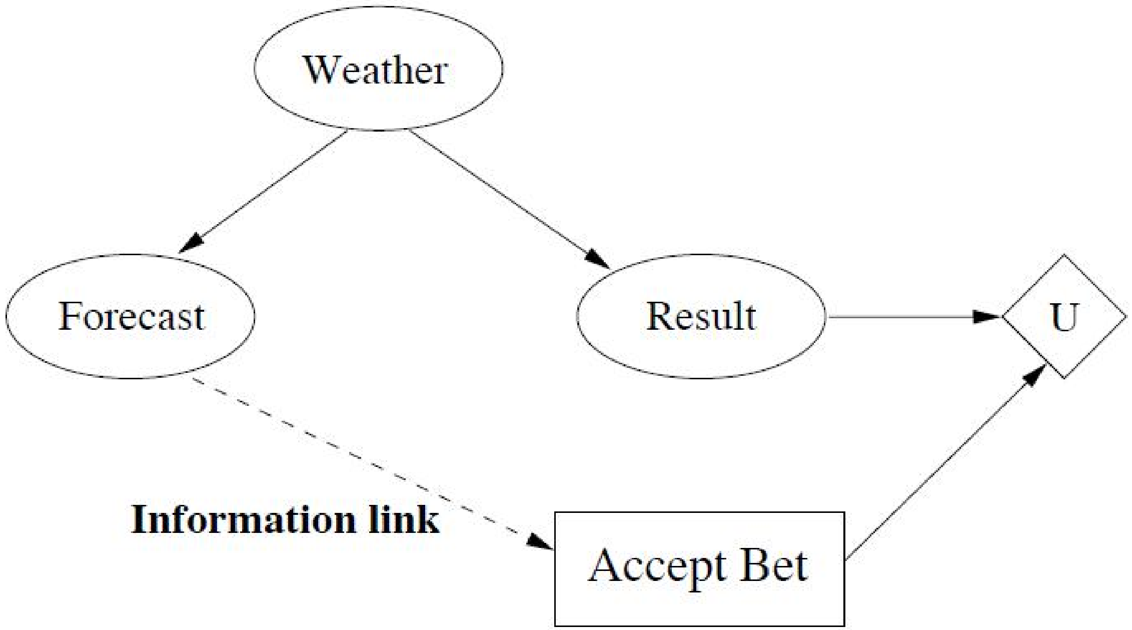

Decision Table

F	Accept Bet
rainy	yes
cloudy	no
sunny	no

Another Example: Fever

Suppose that you know that a fever can be caused by the flu. You can use a thermometer, which is fairly reliable, to test whether or not you have a fever. Suppose you also know that if you take aspirin it will almost certainly lower a fever to normal. Some people (about 5% of the population) have a negative reaction to aspirin. You'll be happy to get rid of your fever, as long as you don't suffer an adverse reaction if you take aspirin.

- Decision network?
- Variables: Fever, Flu, Therm, Reaction, FeverLater?
- Decision: Take aspirin
- Utility function: What is relevant to the utility function?

A Decision Network for the Fever example

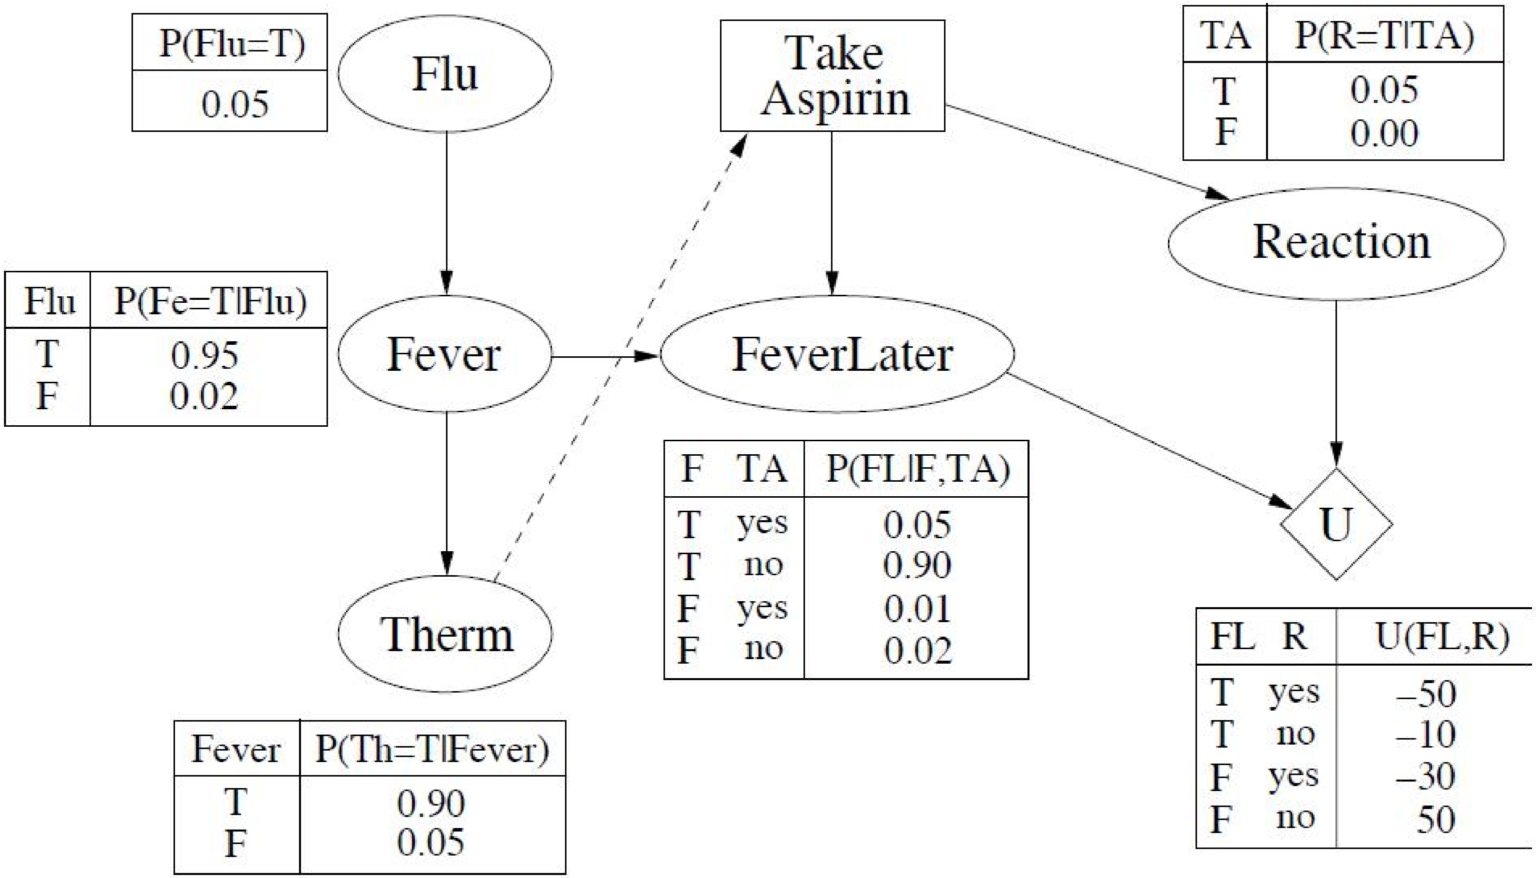

P(Flu=T)
0.05

Flu	P(Fe=T|Flu)
T	0.95
F	0.02

Fever	P(Th=T|Fever)
T	0.90
F	0.05

TA	P(R=T|TA)
T	0.05
F	0.00

F	TA	P(FL|F,TA)
T	yes	0.05
T	no	0.90
F	yes	0.01
F	no	0.02

FL	R	U(FL,R)
T	yes	–50
T	no	–10
F	yes	–30
F	no	50

- The action influences chance nodes, instead of the utility node !!!

Decision Table obtained by following the algorithm

Decisions calculated for the fever problem given different values for *Therm* and *Reaction*.

Evidence	*Bel(FeverLater=T)*	EU(TA=yes)	EU(TA=no)	Decision
None	0.046	45.27	**45.29**	no
Therm=F	0.018	45.40	**48.40**	no
Therm=T	0.273	**44.12**	19.13	yes
Therm=T & *Reaction=T*	0.033	-30.32	**0**	no

Types of Actions

- There are two main types of actions in decision problems, **intervening** and **nonintervening**.
- **Non-intervening actions** do not have a direct effect on the chance variables being modeled, as in the Football team example.
- **Intervening actions** do have direct effects on the world, in the fever example, deciding to take aspirin will affect the later fever situation.

Types of actions

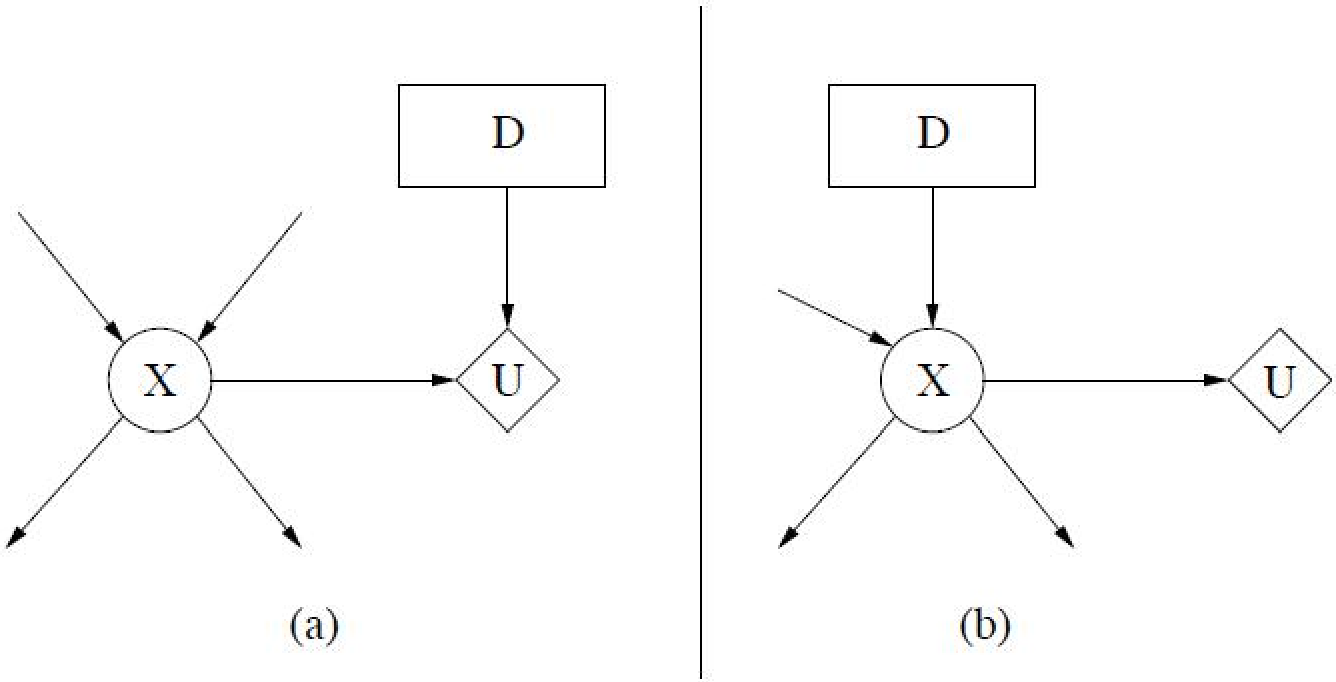

Generic decision networks for (a) non-intervening and (b) intervening actions.