Advanced Topics Data Science and AI
Automated Planning and Acting

Simple Decision Making

Tanya Braun

UNIVERSITÄT ZU LÜBECK
INSTITUT FÜR INFORMATIONSSYSTEME

Content

1. Planning and Acting with **Deterministic** Models
2. Planning and Acting with **Refinement** Methods
3. Planning and Acting with **Temporal** Models
4. Planning and Acting with **Nondeterministic** Models
5. Making Simple Decisions
    a. Utility Theory
    b. Decision Theory
    c. Relational Domains
6. Making Complex Decisions
7. Planning and Acting with **Probabilistic** Models
8. Provably Beneficial AI

- Other: open world, perceiving, learning
    - If time permits

Literature

- We now switch from
  - Automated Planning and Acting
    - Malik Ghallab, Dana Nau, Paolo Traverso
    - Main source
- to
  - Artificial Intelligence: A Modern Approach (3rd ed.)
    - Stuart Russell, Peter Norvig
    - Decision theory
      - Ch. 16 + 17
  - Plus recent research papers mentioned in footnotes

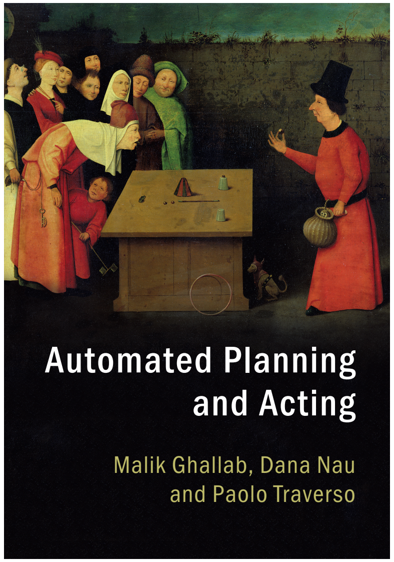

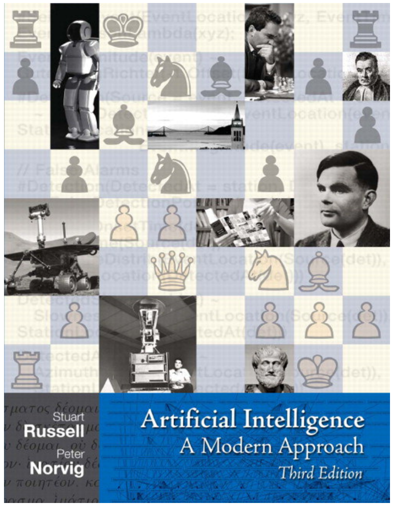

Acknowledgements

- Material from Lise Getoor, Jean-Claude Latombe, Daphne Koller, and Stuart Russell
- Compiled by Ralf Möller

Decision Making under Uncertainty

- Many environments have multiple possible outcomes
- Some of these outcomes may be good; others may be bad
- Some may be very likely; others unlikely
- What's a poor agent going to do??

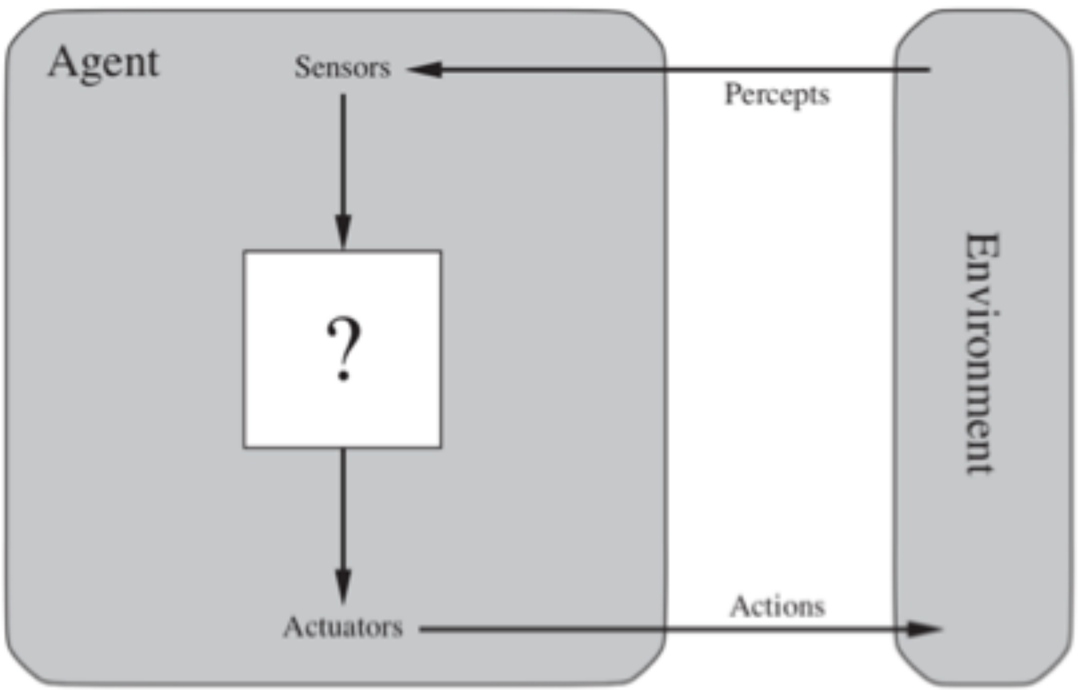

Nondeterministic vs. Probabilistic Uncertainty

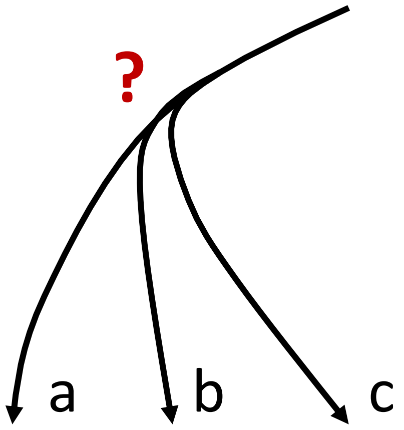

Nondeterministic model

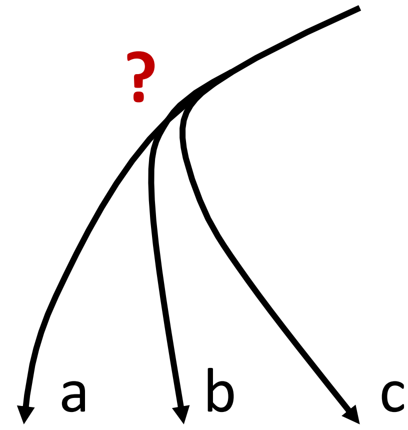

Probabilistic model

- $\{a, b, c\}$
- Decision that is
  best for worst case

- $\{a(p_a), b(p_b), c(p_c)\}$
- Decision that
  maximises expected
  utility value

Expected Utility

- Random variable X with n range values $x_1, \dots, x_n$ and distribution $(p_1, \dots, p_n)$
  - E.g.: X is the state reached after doing an action $A = a$ under uncertainty
- Function U of X
  - E.g., U is the utility of a state
- The expected utility of $A = a$ is

$$EU[A = a] = \sum_{i=1}^{n} P(X = x_i | A = a) \cdot U(X = x_i)$$

One State/One Action Example

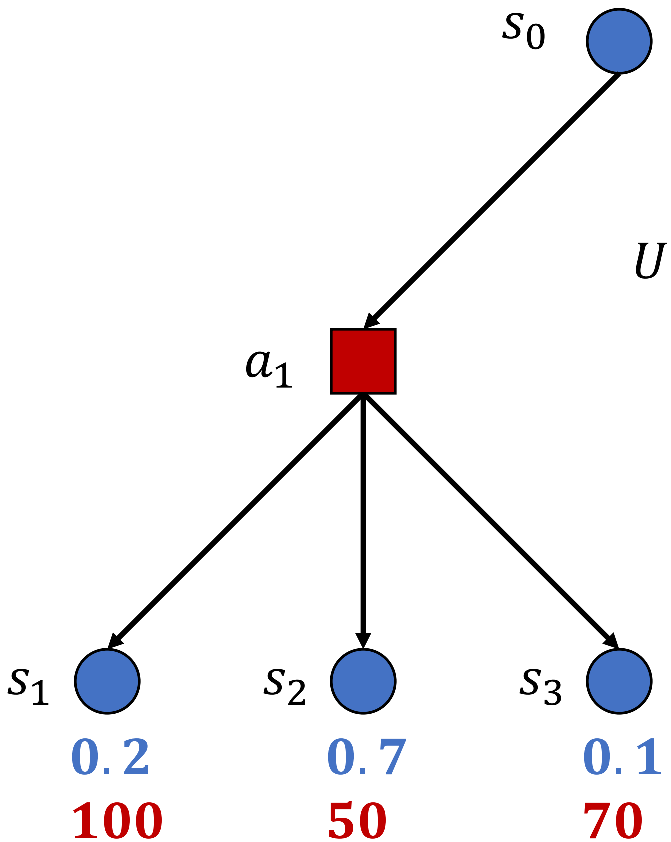

$$U(s_0) = 100 \cdot 0.2 + 50 \cdot 0.7 + 70 \cdot 0.1$$
$$= 20 + 35 + 7$$
$$= 62$$

One State/Two Actions Example

$U_1(s_0) = 62$

$U_2(s_0) = 74$

$U(s_0) = \max\{U_1(s_0), U_2(s_0)\}$
$= 74$

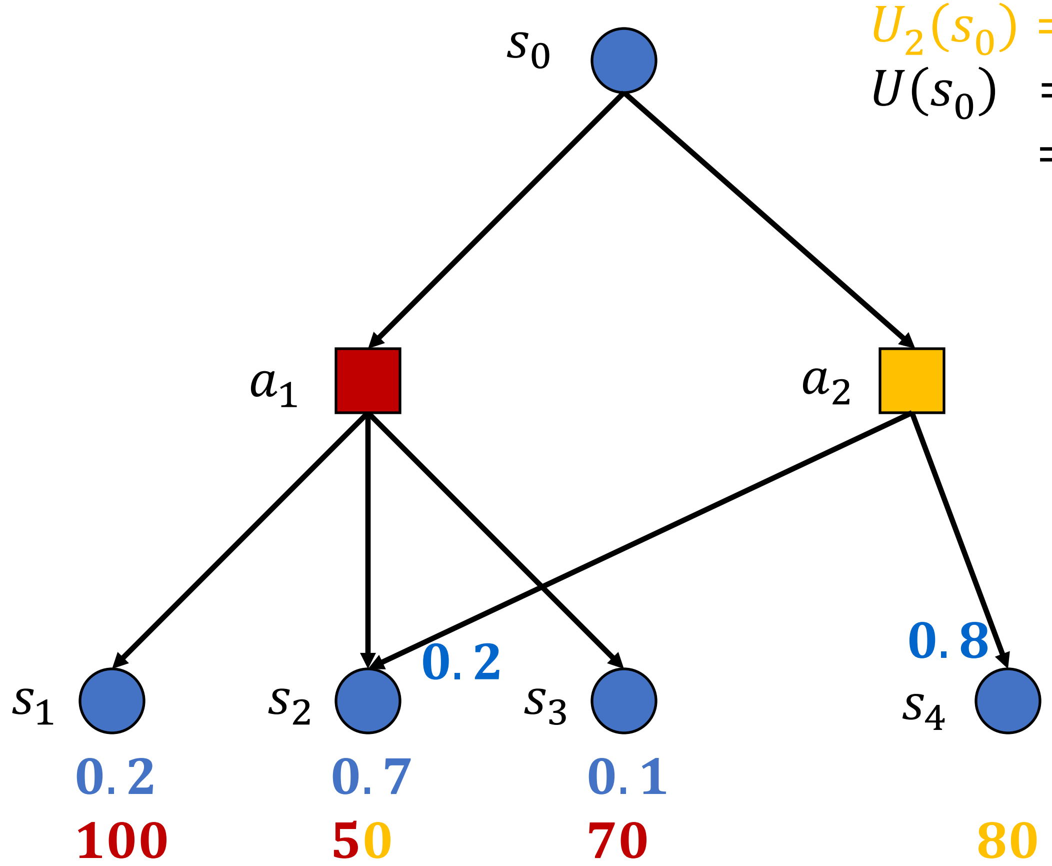

Introducing Action Costs

$U_1(s_0) = 62 - 5$
$U_2(s_0) = 74 - 25$
$U(s_0) \quad = \max\{U_1(s_0), U_2(s_0)\}$
$\quad = 57$

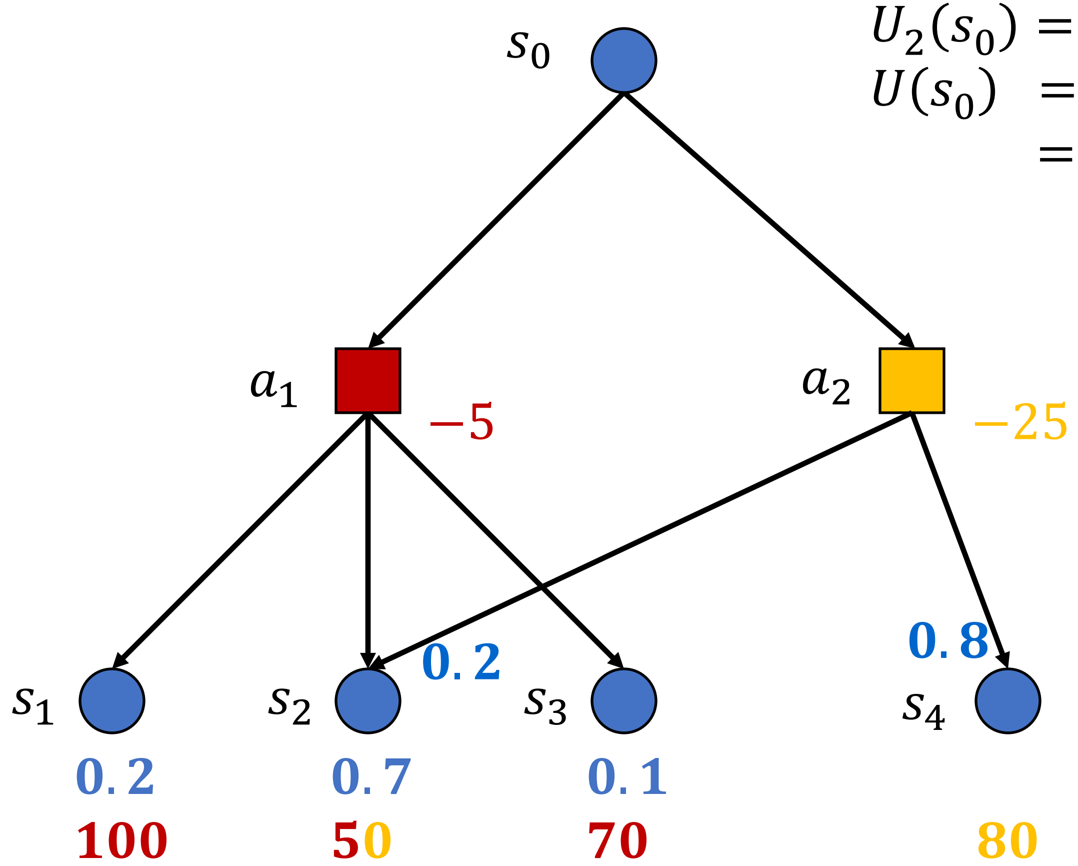

MEU Principle

- A rational agent should choose the action that maximizes agent's expected utility
- This is the basis of the field of decision theory
- The MEU principle provides a normative criterion for rational choice of action

AI is solved!!!

Not quite…

- Must have complete model of:
  - Actions
  - Utilities
  - States
- Even if you have a complete model, it might be computationally intractable
- In fact, a truly rational agent takes into account the utility of reasoning as well – bounded rationality
- Nevertheless, great progress has been made in this area, and we are able to solve much more complex decision-theoretic problems than ever before

Setting

- Agent can perform actions in an environment
  - Environment
    - Episodic, i.e., not sequential
      - Next episode does not depend on the previous episode
      - So called static models (vs. dynamic/temporal, next lecture)
    - Non-deterministic
      - Outcomes of actions not unique
      - Associated with probabilities (→ probabilistic model)
    - Partially observable
      - Latent, i.e., not observable, random variables
  - Agent has preferences over states/action outcomes
    - Encoded in utility or utility function → Utility theory
- "Decision theory = Utility theory + Probability theory"
  - Model the world with a probabilistic model
  - Model preferences with a utility (function)
  - Find action that leads to the maximum expected utility, also called decision making
    - Lecture title: "Simple decisions" because episodic

Outline (mainly Ch. 16)

Utility theory

- Preferences
- Utilities
- Dominance
- Preference structure

Decision theory

- Decision networks
- Value of information
- Relational domains

Preferences

- An agent chooses among prizes (A, B, etc.) and lotteries, i.e., situations with uncertain prizes
  - Outcome of a nondeterministic action is a lottery
- Lottery $L = [p, A; (1-p), B]$
  - A and B can be lotteries again
  - Prizes are special lotteries: $[1, X; 0, \text{not } X]$
  - More than two outcomes:
    - $L = [p_1, S_1; p_2, S_2; \cdots; p_n, S_n], \sum_{i=1}^{n} p_i = 1$
- Notation
  - $A \succ B$ $\quad$ A preferred to B
  - $A \sim B$ $\quad$ indifference between A and B
  - $A \succsim B$ $\quad$ B not preferred to A

Rational preferences

- Idea: preferences of a rational agent must obey constraints
- Rational preferences ⇒ behaviour describable as maximisation of expected utility

Rational preferences contd.

- Violating constraints leads to self-evident irrationality
- 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

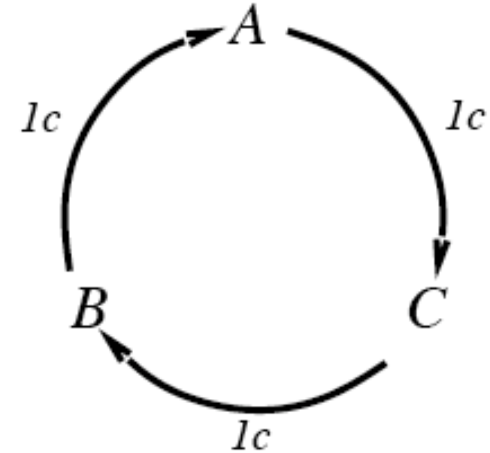

Axioms of Utility Theory

1. Orderability
   - $(A \succ B) \vee (A \prec B) \vee (A \sim B)$
   - $\{\prec, \succ, \sim\}$ jointly exhaustive, pairwise disjoint
2. Transitivity
   - $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
3. Continuity
   - $A \succ B \succ C \Rightarrow$
     $\exists p\ [p, A;\ 1-p, C] \sim B$
4. Substitutability
   - $A \sim B \Rightarrow$
     $[p, A; 1-p, C] \sim [p, B; 1-p, C]$
   - Also holds if replacing $\sim$ with $\succ$
5. Monotonicity
   - $A \succ B \Rightarrow$
     $(p \geq q \Leftrightarrow$
     $[p, A; 1-p, B]$
     $\succsim\ [q, A; 1-q, B])$
6. Decomposability
   - $[p, A;\ 1-p, [q, B;\ 1-q, C]] \sim$
     $[p, A;\ (1-p)q, B; (1-p)(1-q), C]$

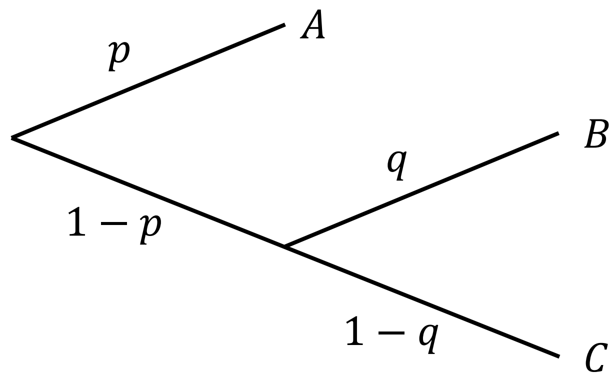

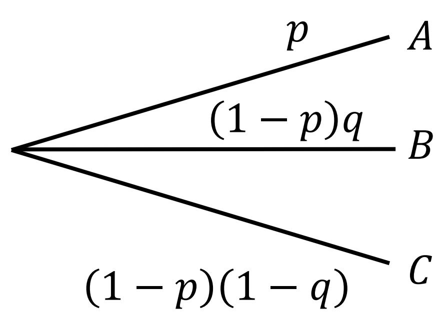

Decomposability: There is no fun in gambling.

And Then There Was Utility

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

$$U(A) \geq U(B) \Leftrightarrow A \gtrsim B$$
$$U([p_1, S_1; \ldots; p_n, S_n]) = \sum_i p_i U(S_i)$$

- MEU principle
  - Choose the action that maximises 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

Utilities

- 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 outcome" $\top$ with probability p
    - "worst possible catastrophe" $\bot$ with probability $(1-p)$
  - Adjust lottery probability p until $A \sim L_p$

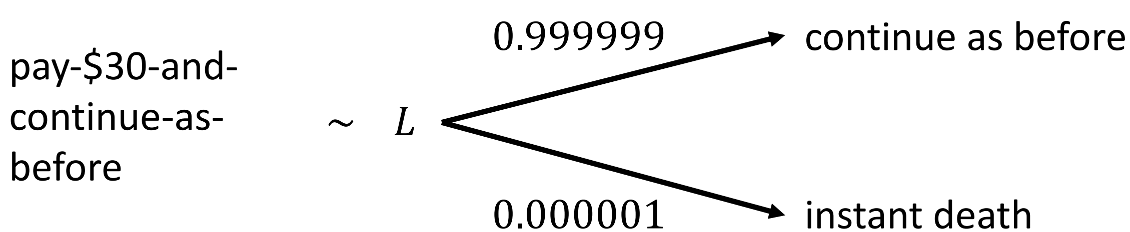

Utility scales

- Normalised utilities: $u_\top = 1.0, u_\perp = 0.0$
  - Utility of lottery $L \sim$ (pay-\$30-and-continue-as-before): $U(L) = u_\top \cdot 0.999999 + u_\perp \cdot 0.000001 = 0.999999$
- Micromorts: one-millionth chance of death
  - Useful for Russian roulette, paying to reduce product risks, etc.
- QALYs: quality-adjusted life years
  - Useful for medical decisions involving substantial risk
- Behaviour is invariant w.r.t. positive linear transformation

$$U'(x) = k_1 U(x) + k_2$$

  - No unique utility function; $U'(x)$ and $U(x)$ yield same behaviour

Ordinal Utility Functions

- With deterministic prizes only (no lottery choices), only ordinal utility can be determined, i.e., total order on prizes
  - Ordinal utility function also called value function
  - Provides a ranking of alternatives (states), but not a meaningful metric scale (numbers do not matter)

Money

- Money does not behave as a utility function
- Given a lottery L with expected monetary value $EMV(L)$, usually $U(L) < U(S_{EMV(L)})$, i.e., people are risk-averse
  - S_n: state of possessing total wealth $n
  - Utility curve
    - For what probability p am I indifferent between a prize x and a lottery $[p, \$M; (1-p), \$0]$ for large M?
    - Right: Typical empirical data, extrapolated with risk-prone behaviour for negative wealth

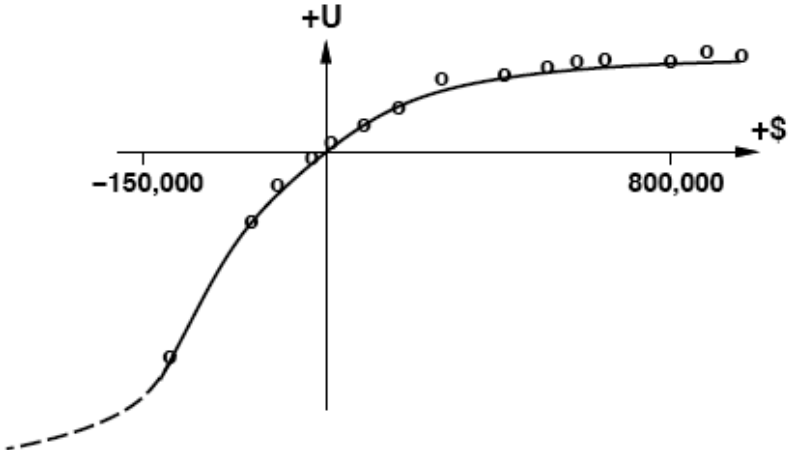

Money Versus Utility

- Money ≠ Utility
  - More money is better, but not always in a linear relationship to the amount of money
- Expected Monetary Value
  - Risk-averse
    - $U(L) < U(S_{EMV(L)})$
  - Risk-seeking
    - $U(L) > U(S_{EMV(L)})$
  - Risk-neutral
    - $U(L) = U(S_{EMV(L)})$
    - Linear curve
    - For small changes in wealth relative to current wealth

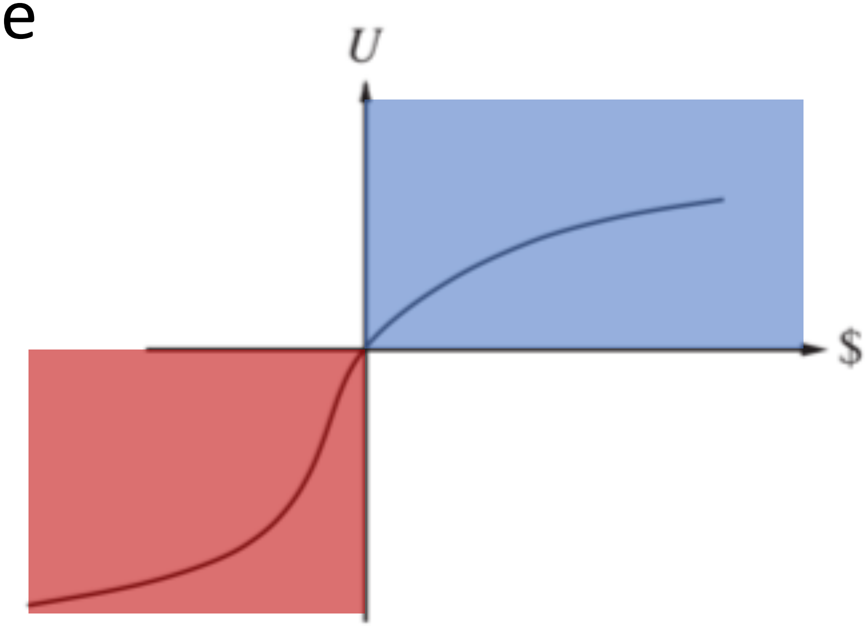

Multiattribute Utility Theory

- A given state may have multiple utilities
  - ...because of multiple evaluation criteria
  - ...because of multiple agents (interested parties) with different utility functions

- We will look at
  - Cases in which decisions can be made *without* combining the attribute values into a single utility value
    - Strict dominance
  - Cases in which the utilities of attribute combinations can be specified very concisely

Strict dominance

- Typically define attributes such that U is monotonic in each →
- Strict dominance
  - Choice B strictly dominates choice A iff

$$\forall\, i : X_i(B) \geq X_i(A) \text{ (and hence } U(B) \geq U(A))$$

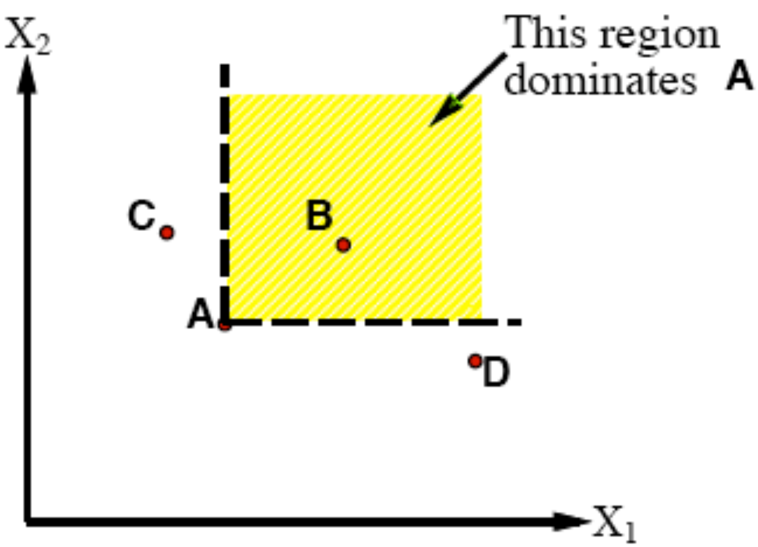

Deterministic attributes

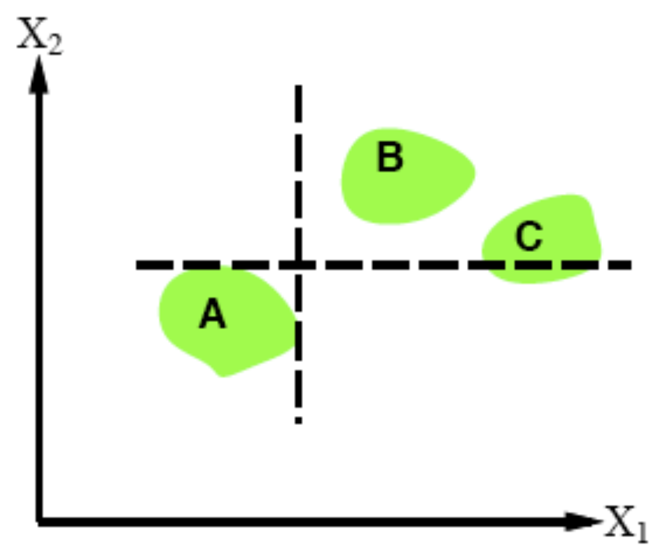

Uncertain attributes

Stochastic dominance

- Cumulative distribution p_1 first-order stochastically dominates distribution p_2 iff
$$\forall x : p_1(x) \leq p_2(x)$$
  - With a strict inequality for some interval
  - Then, $E_{p_1} > E_{p_2}$ (E referring to expected value)
    - The reverse is not necessarily true
  - Does not imply that every possible return of the superior distribution is larger than every possible return of the inferior distribution
- Example:
  - As we have *negative costs*, S2 dominates S1 with $\forall x : p_{S_2}(x) \leq p_{S_1}(x)$

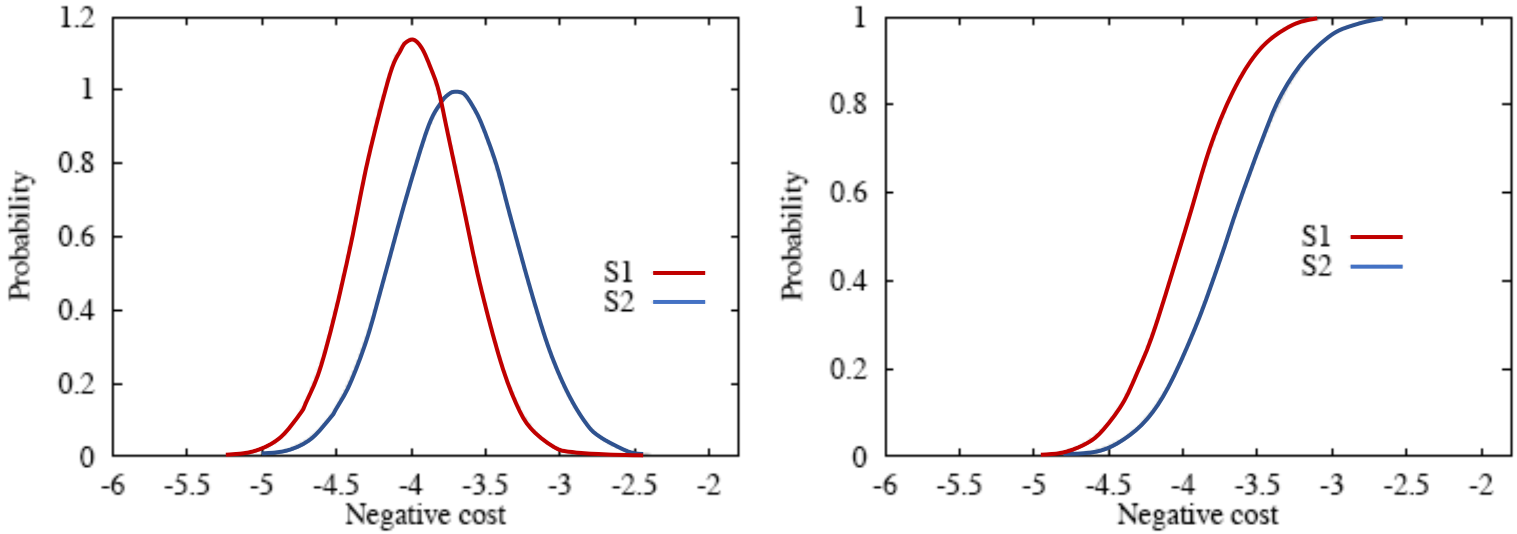

https://people.duke.edu/~dgraham/ECO_463/Handouts/StochasticDominance.pdf

Example

- Product P

Profit ($m)	Probability
0 to under 5	0.2
5 to under 10	0.3
10 to under 15	0.4
15 to under 20	0.1

- Product Q

Profit ($m)	Probability
0 to under 5	0.0
5 to under 10	0.1
10 to under 15	0.5
15 to under 20	0.3
20 to under 25	0.1

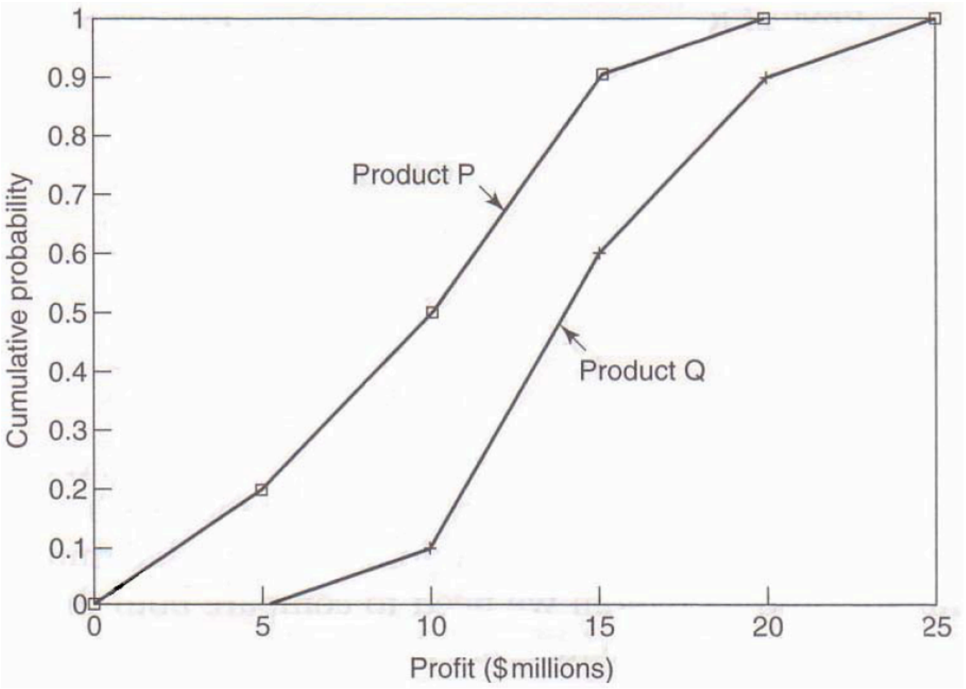

Stochastic dominance

- Cumulative distribution p_1 second-order stochastically dominates distribution p_2 iff

$$\forall\, t: \int_{-\infty}^{t} p_1(x)\,dx \le \int_{-\infty}^{t} p_2(x)\,dx$$

  - Or: $D(t) = \int_{-\infty}^{t} p_2(x) - p_1(x)\,dx \ge 0$
  - With a strict inequality for some interval
  - Then, $E_{p_1} \ge E_{p_2}$ (E referring to expected value)
- Examples with $t = z$:
  - Second-order stochastic dominance

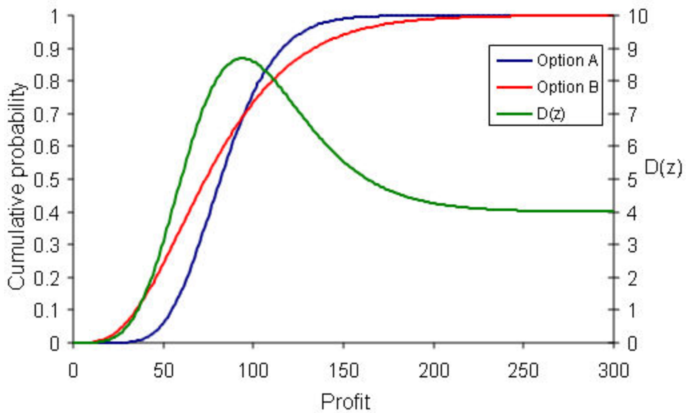

  - No dominance

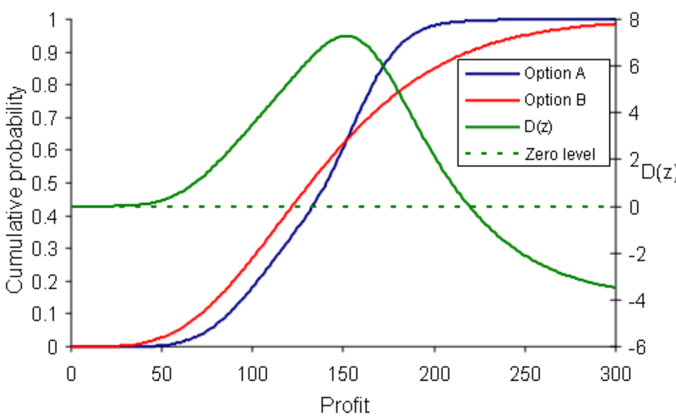

https://people.duke.edu/~dgraham/ECO_463/Handouts/StochasticDominance.pdf

Figures: https://www.vosesoftware.com/riskwiki/Stochasticdominancetests.php

Preference Structure

- To specify the complete utility function $U(x_1, \ldots, x_n)$, we need d^n values in the worst case
  - n attributes
  - each attribute with d distinct possible values
  - Worst case meaning: Agent's preferences have no regularity at all
- Supposition in multiattribute utility theory
  - Preferences of typical agents have much more structure
- Approach
  - Identify regularities in the preference behaviour
  - Use so-called representation theorems to show that an agent with a certain kind of preference structure has a utility function

$$U(x_1, \ldots, x_n) = F[f_1(x_1), \ldots, f_n(x_n)]$$

    - where F is hopefully a simple function such as addition

Preference structure: Deterministic

- X_1 and X_2 preferentially independent (PI) of X_3 iff
  - Preference between $\langle x_1, x_2, x_3 \rangle$ and $\langle x_1', x_2', x_3 \rangle$ does not depend on x_3
  - E.g., $\langle Noise, Cost, Safety \rangle$
    - $\langle 20{,}000\ suffer, \$4.6\ billion, 0.06\ deaths/month \rangle$
    - $\langle 70{,}000\ suffer, \$4.2\ billion, 0.06\ deaths/month \rangle$
- Theorem (Leontief, 1947)
  - If every pair of attributes is PI of its complement, then every subset of attributes is PI of its complement
    - Called mutual PI (MPI)
- Theorem (Debreu, 1960):
  - MPI ⇒ ∃ *additive* value function

$$V(x_1, \dots, x_n) = \sum_i V_i(x_i)$$

  - Hence assess n single-attribute functions
  - Often a good approximation

Preference structure: Stochastic

- Need to consider preferences over lotteries
- X is utility-independent (UI) of Y iff
  - Preferences over lotteries in X do not depend on y
- Mutual UI (Keeney, 1974): each subset is UI of its complement ⇒ ∃ *multiplicative* utility function
  - For $n = 3$:
    $$U = k_1U_1 + k_2U_2 + k_3U_3 \\ +k_1k_2U_1U_2 + k_2k_3U_2U_3 + k_3k_1U_3U_1 \\ +k_1k_2k_3U_1U_2U_3$$
  - I.e., requires only n single-attribute utility functions and n constants

Intermediate Summary

- Preferences
  - Preferences of a rational agent must obey constraints
- Utilities
  - Rational preferences = describable as maximisation of expected utility
  - Utility axioms
  - MEU principle
- Dominance
  - Strict dominance
  - First-order + second-order stochastic dominance
- Preference structure
  - (Mutual) preferential independence
  - (Mutual) utility independence

Outline

Utility theory

- Preferences
- Utilities
- Dominance
- Preference structure

Decision theory

- Decision networks
- Value of information
- Relational domains

Decision Networks

- Extend Bayesian networks (BNs) to handle actions and utilities
  - Or any other probabilistic (graphical) formalism
- Also called influence diagrams
- Use BN inference methods to solve MEU problems
- Perform Value of Information calculations

Decision Networks cont.

- Chance nodes:
  random variables
  - As in BNs

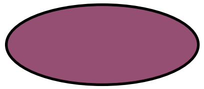

- Decision nodes:
  actions that decision maker can take

- Utility/value nodes:
  the utility of the outcome state

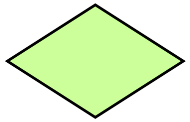

Umbrella Network

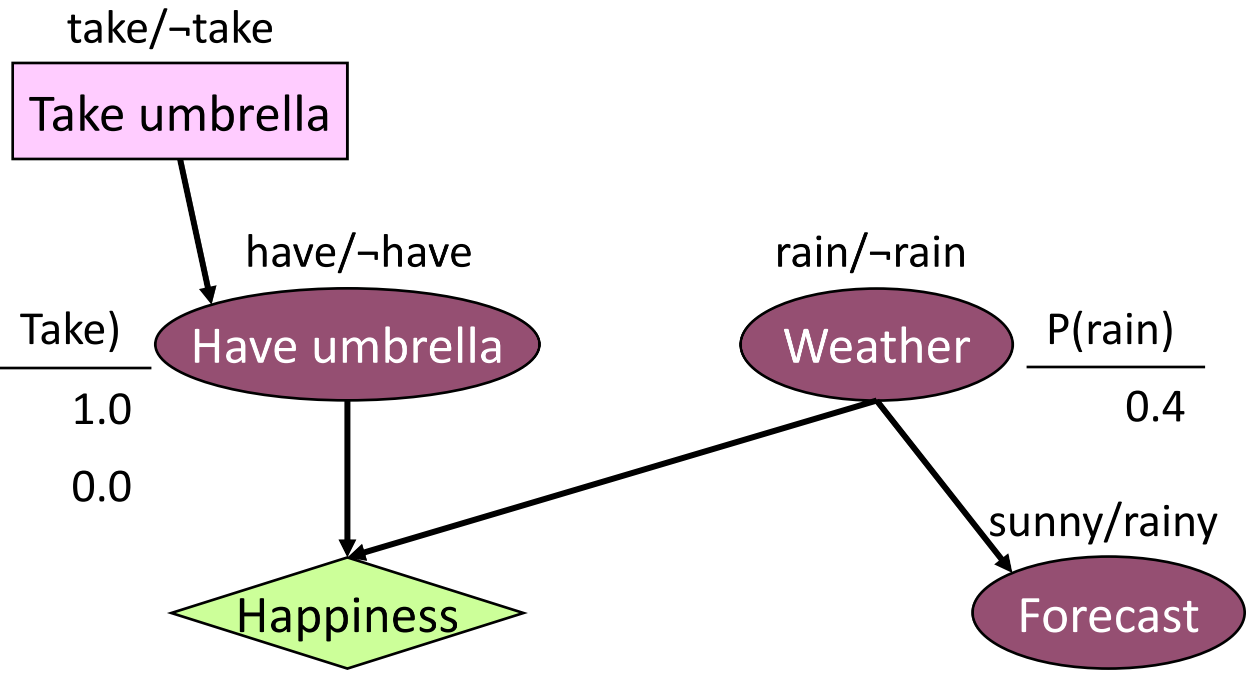

Take	P(have | Take)
take	1.0
¬take	0.0

P(rain)
0.4

Have	Weather	U(Have, Weather)
have	rain	-25
have	¬rain	0
¬have	rain	-100
¬have	¬rain	100

Weather	P(sunny | Weather)
rain	0.3
¬rain	0.8

Evaluating Decision Networks

- Set the evidence variables for current state
- For each possible value of the decision node:
  - Set decision node to that value
  - Calculate the posterior probability of the parent nodes of the utility node, using BN inference
  - Calculate the resulting utility for action
- Return the action with the highest utility

Umbrella Network

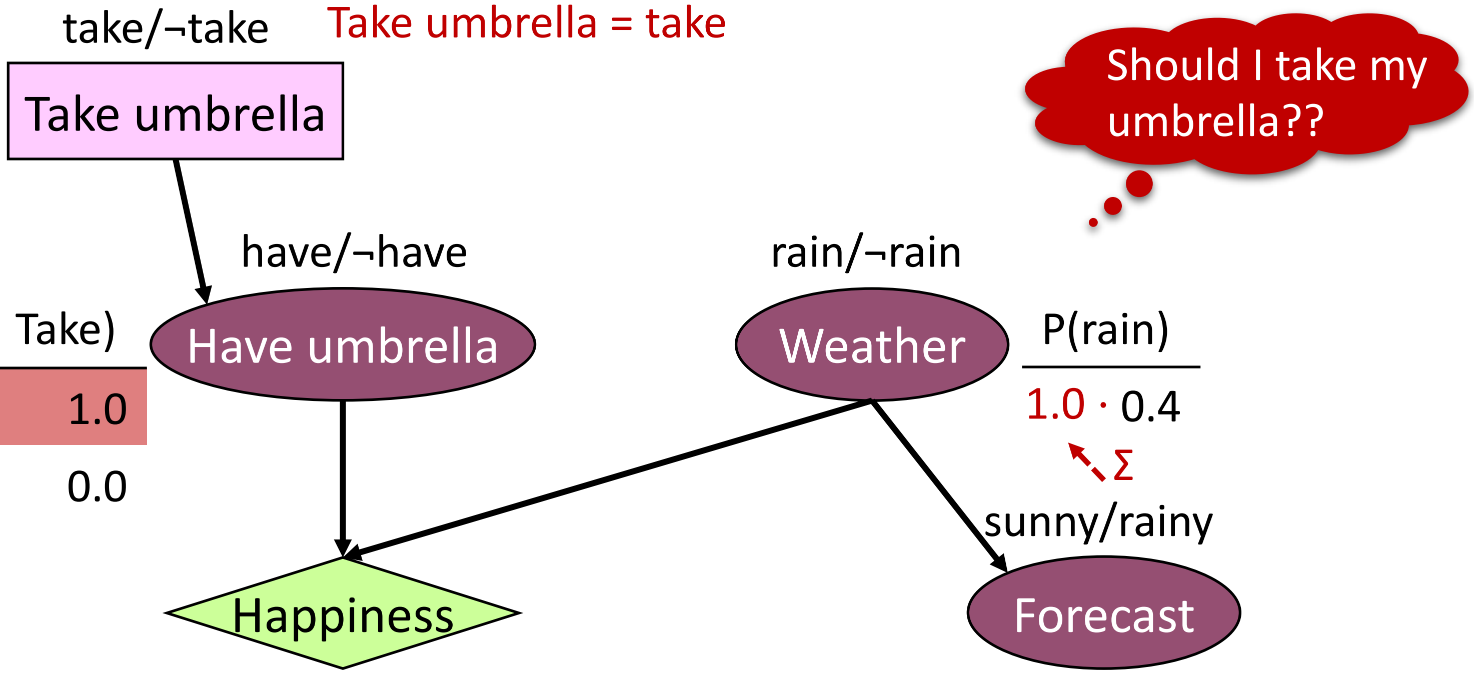

Take	P(have | Take)
take	1.0
¬take	0.0

P(rain)
1.0 · 0.4

↖ Σ

Have	Weather	U(Have, Weather)	
have	rain	1.0 · 0.4 · -25	= -10
have	¬rain	1.0 · 0.6 · 0	= 0
¬have	rain	0.0 · 0.4 · -100	= 0
¬have	¬rain	0.0 · 0.6 · 100	= 0

Weather	P(sunny | Weather)
rain	0.3
¬rain	0.8

EU(take) = -10 + 0 + 0 + 0 = -10

Umbrella Network

take/¬take Take umbrella = ¬take

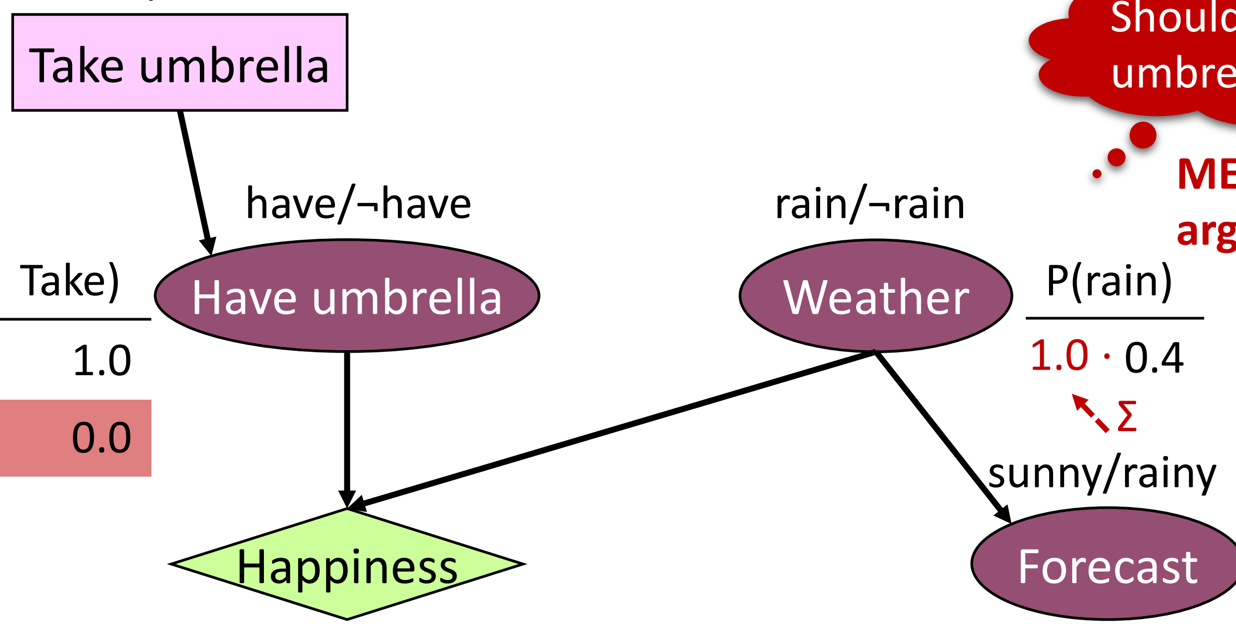

Should I take my umbrella??

MEU(Take) = 20
argmax = ¬take

Take	P(have | Take)
take	1.0
¬take	0.0

P(rain)
1.0 · 0.4

↖ Σ

Have	Weather	U(Have, Weather)	
have	rain	0.0 · 0.4 · -25	= 0
have	¬rain	0.0 · 0.6 · 0	= 0
¬have	rain	1.0 · 0.4 · -100	= -40
¬have	¬rain	1.0 · 0.6 · 100	= 60

Weather	P(sunny | Weather)
rain	0.3
¬rain	0.8

EU(¬take) = 0 + 0 - 40 + 60 = 20

EU(take) = -10 + 0 + 0 + 0 = -10

Decision Making in Decision Nets

- Assumes that all available information provided to agent before it makes its decision
  - Hardly ever the case
  - Know what questions to ask!
- Information value theory
  - Choose what information to acquire
  - Assume that prior to selecting an action represented by a decision node, the agent can acquire the value of any of the potentially observable chance nodes
    - Simplified version of sequential decision making (next lecture)
      - Observation actions affect only agent's belief state, not the external physical state

Value of information

- Idea: Compute value of acquiring each possible piece of evidence
  - Can be done directly from decision network
- Example: Buying oil drilling rights
  - Two blocks A and B, exactly one has oil, worth k
  - Prior probabilities 0.5 each, mutually exclusive
  - Current price of each block is $k/2$
  - "Consultant" offers accurate survey of A
  - Fair price for survey?
  - Solution: Compute expected value of information
    - = expected value of best action given the information minus expected value of best action without information
    - Survey may say "oil in A" or "no oil in A", probability 0.5 each (given!)
      $= [0.5 \cdot \text{value of "buy A" given "oil in A"} + 0.5 \cdot \text{value of "buy B" given "no oil in A"}] - 0$
      $= (0.5 \cdot k/2) + (0.5 \cdot k/2) - 0 = k/2$

General formula

- Current evidence E, current best action α, possible action outcomes S_i, potential new evidence E_j

$$EU(\alpha|E) = \max_{\mathrm{a}} \sum_i U(S_i)P(S_i \mid E, a)$$

- Suppose we knew $E_j = e_{jk}$, then we would choose a_{jk} such that

$$EU\left(\alpha_{e_{jk}} \mid E, E_j = e_{jk}\right) = \max_{\mathrm{a}} \sum_i U(S_i)P\left(S_i \mid E, a, E_j = e_{jk}\right)$$

- E_j is a random variable whose value is currently unknown ⇒ must compute expected gain over all possible values:

$$VPI_E(E_j) = \left(\sum_k P(E_j \mid E)\, EU\left(\alpha_{e_{jk}} \mid E, E_j = e_{jk}\right)\right) - EU(\alpha, E)$$

  - VPI = value of perfect information

Properties of VPI

- Non-negative – in expectation
$$\forall j, E : VPI_E(E_j) \geq 0$$
- Non-additive – consider, e.g., obtaining E_j twice
$$VPI_E(E_j, E_k) \neq VPI_E(E_j) + VPI_E(E_k)$$
- Order-independent
$$VPI_E(E_j, E_k) = VPI_E(E_j) + VPI_{E,E_j}(E_k)$$
$$= VPI_E(E_k) + VPI_{E,E_k}(E_j)$$
- Note: When more than one piece of evidence can be gathered, maximising VPI for each to select one is not always optimal

  ⇒ Evidence-gathering becomes a sequential decision problem

Qualitative behaviors

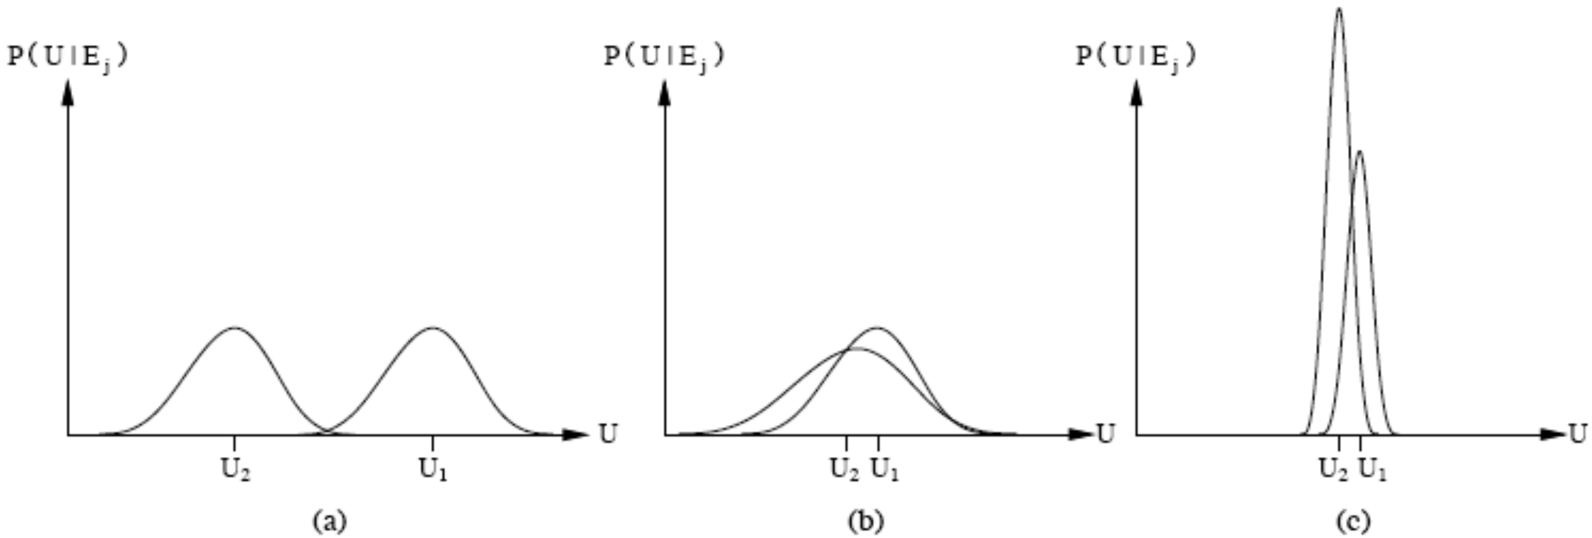

a) Choice is obvious, information worth little
b) Choice is non-obvious, information worth a lot
c) Choice is non-obvious, information worth little

- *Information has value to the extent that it is likely to cause a change of plan and to the extent that the new plan will be significantly better than the old plan*

Information Gathering Agent

```
function INFORMATION-GATHERING-AGENT(percept)
returns: an action
persistent: D, a decision network

   integrate percept into D
   j ← the value that maximises VPI(E_j)/Cost(E_j)
   if VPI(E_j) > Cost(E_j) then
      return Request(E_j)
   else
      return the best action from D
```

- Ask questions $Request(E_j)$ in a reasonable order
- Avoid irrelevant questions
- Take into account importance of piece of information j in relation to $Cost(E_j)$

Decision Making in Decision Nets II

- Solving MEU/query answering problems intractable in general
  - Query answering: Computing probability distributions (given evidence)
  - Exponential in tree width of the graphical model
    - Tree width ≈ Largest number of arguments in a table/factor to occur during calculations
- Regularities in graphical model may allow to reduce the tree width by explicitly encoding them and using them during calculations

Relational Domains

- Relations between objects/individuals/constants
  - Regularities/symmetries

- Constructs of first-order logic to parameterise a propositional formalism
  - Symmetries encoded compactly using logical variables
  - Parameterised random variables (PRVs) to denote sets of random variables behaving identically

Compact Representation

- Factor graph to parfactor graph

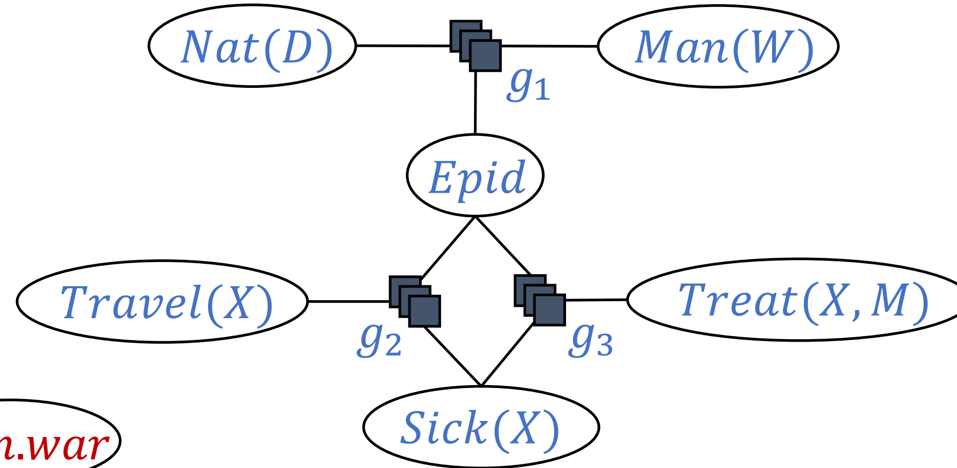

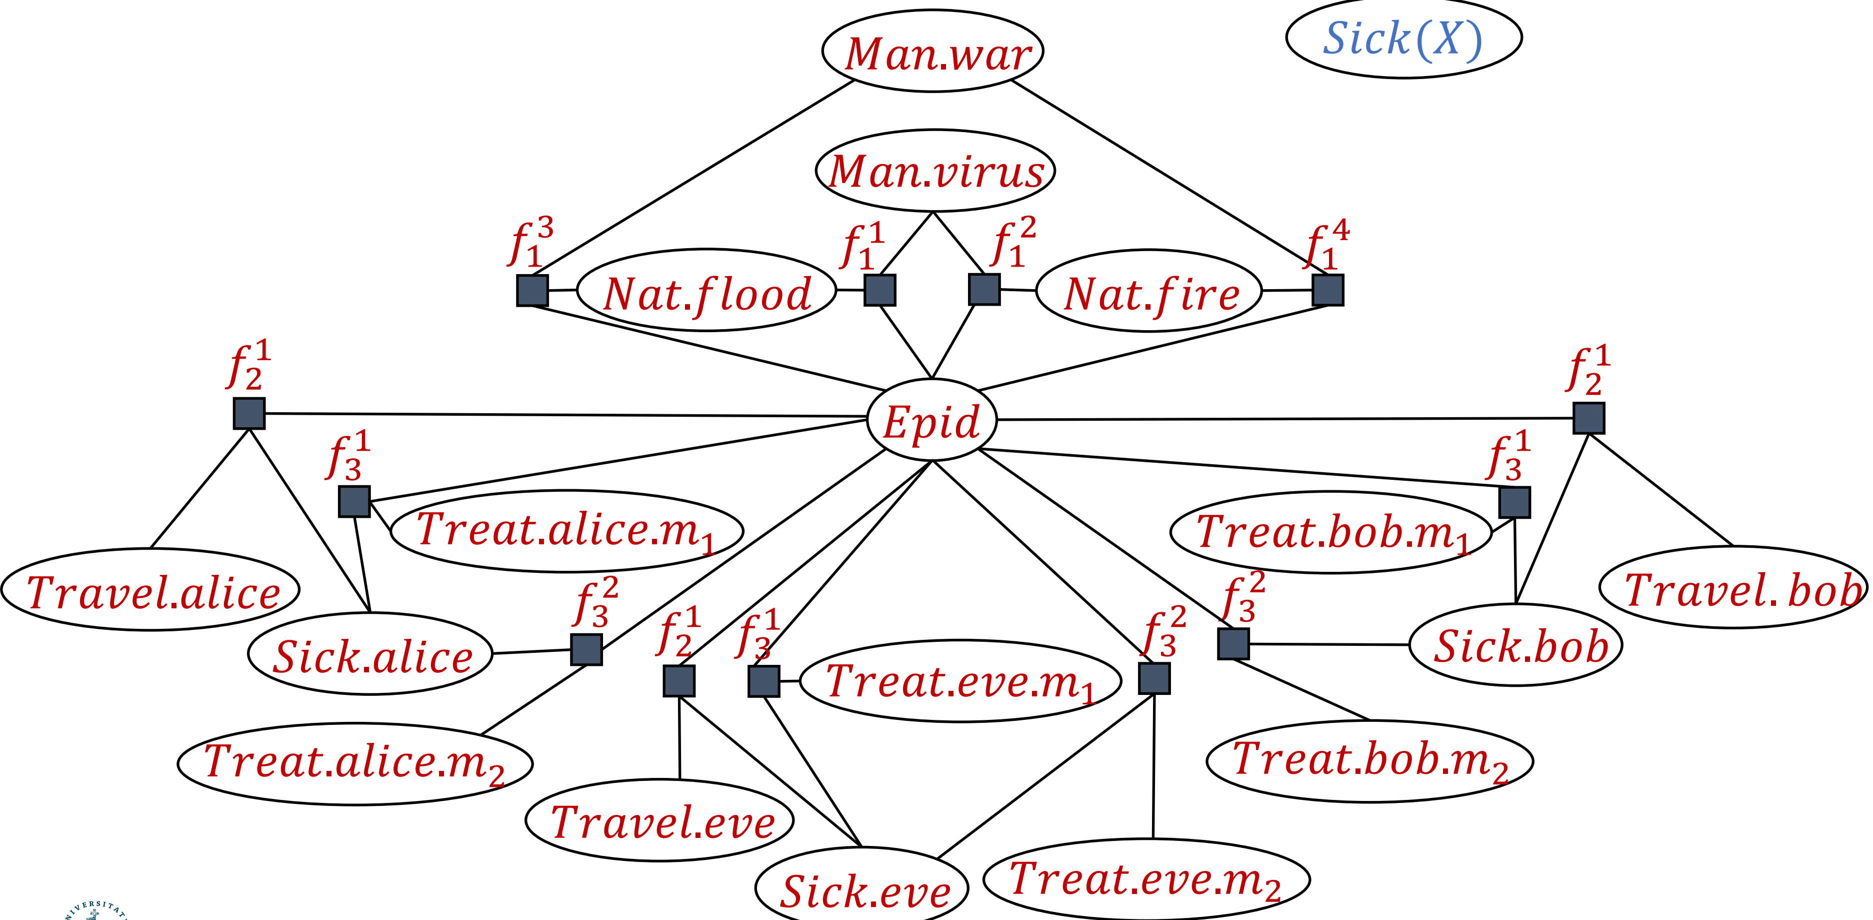

Compression: Pass the colours around

- If you have a (propositional) model available*
- **Colour nodes according to the evidence you have**
  - No evidence, say **red**
  - State „one“, say **brown**
  - State „two“, say **orange**
  - …
- **Colour factors distinctively according to their equivalences**
  For instance, assuming f_1 and f_2 to be identical and B appears at the second position within both, say **blue**

*can also be done at the „lifted“, i.e., relational level

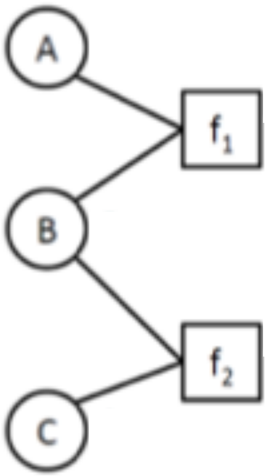

Singla and Domingos (2008), Kersting et al. (2009), Ahmadi et al. (2013)

Compression: Pass the colours around

1. Each factor collects the colours of its neighbouring nodes

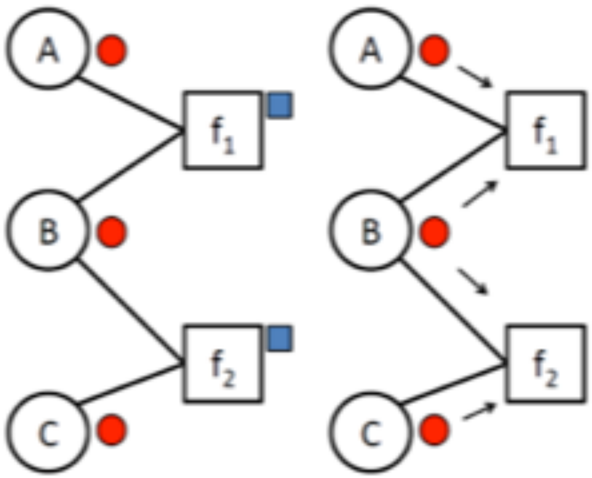

Compression: Pass the colours around

1. Each factor collects the colours of its neighbouring nodes
2. Each factor „signs“ its colour signature with its own colour

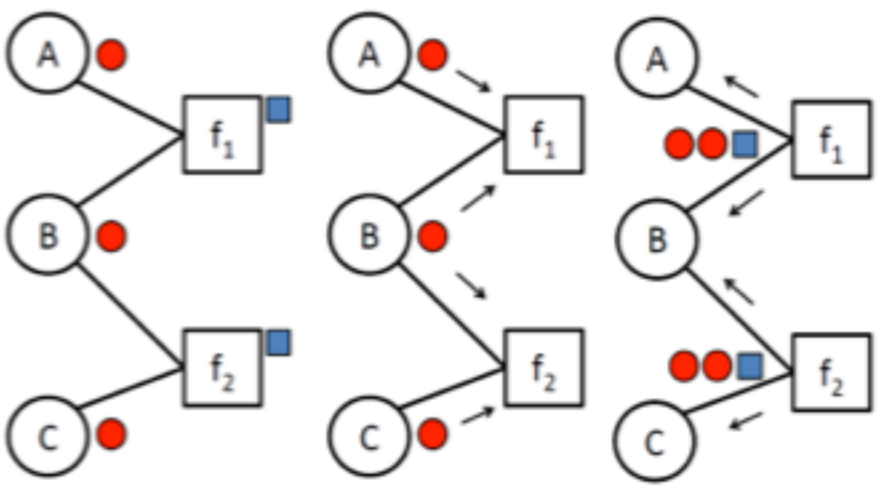

Compression: Pass the colours around

1. Each factor collects the colours of its neighbouring nodes
2. Each factor „signs“ its colour signature with its own colour
3. Each node collects the signatures of its neighbouring factors

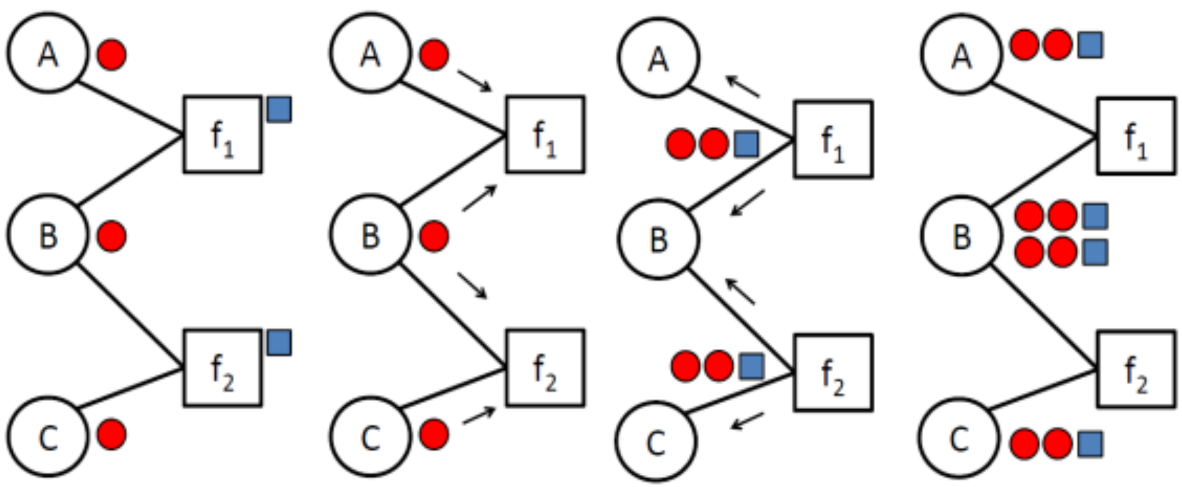

Compression: Pass the colours around

1. Each factor collects the colours of its neighbouring nodes
2. Each factor „signs“ its colour signature with its own colour
3. Each node collects the signatures of its neighbouring factors
4. Nodes are recoloured according to the collected signatures

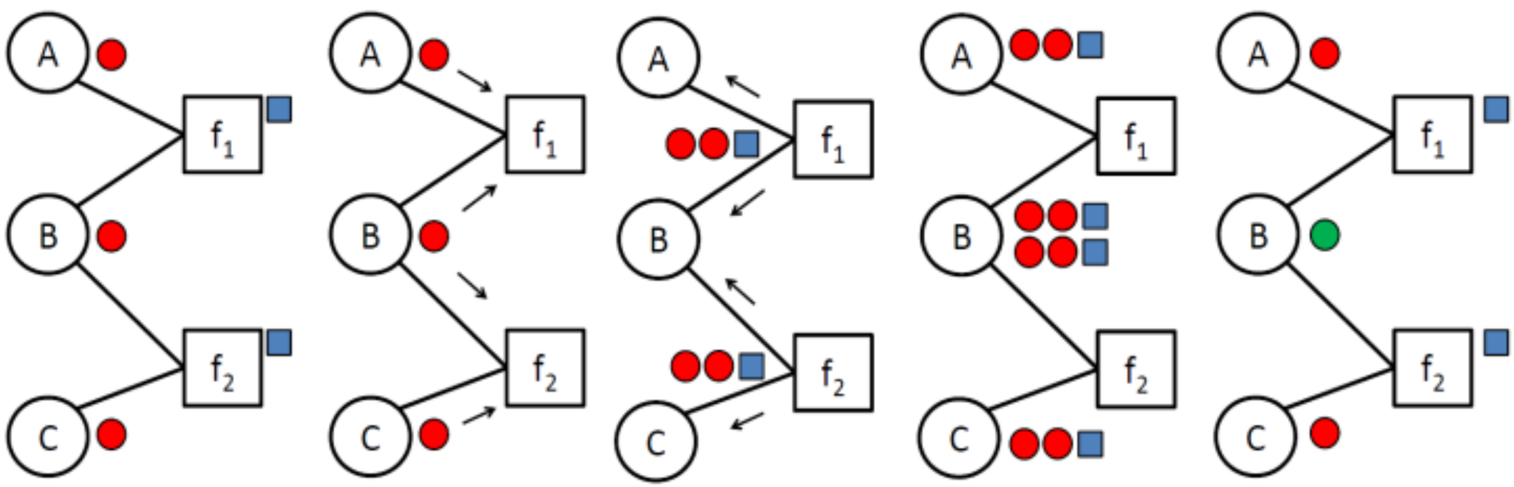

Compression: Pass the colours around

1. Each factor collects the colours of its neighbouring nodes
2. Each factor „signs“ its colour signature with its own colour
3. Each node collects the signatures of its neighbouring factors
4. Nodes are recoloured according to the collected signatures
5. If no new colour is created stop, otherwise go back to 1

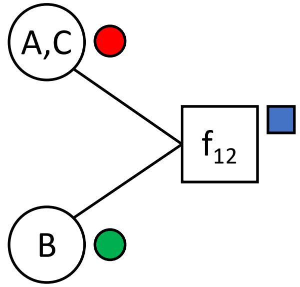

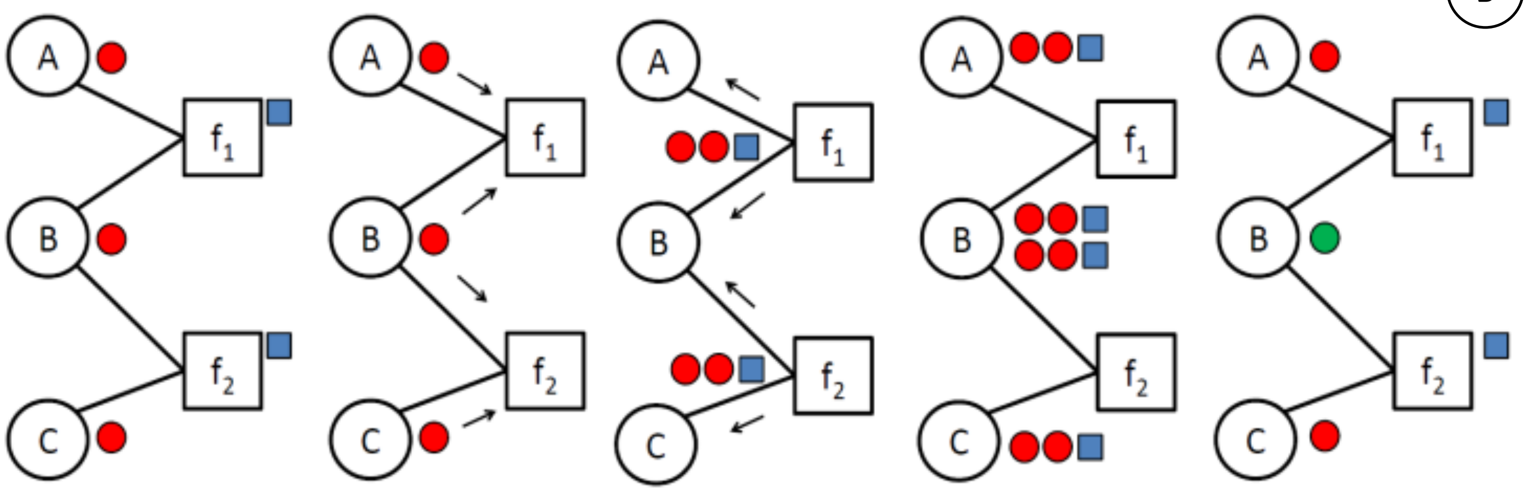

Compression

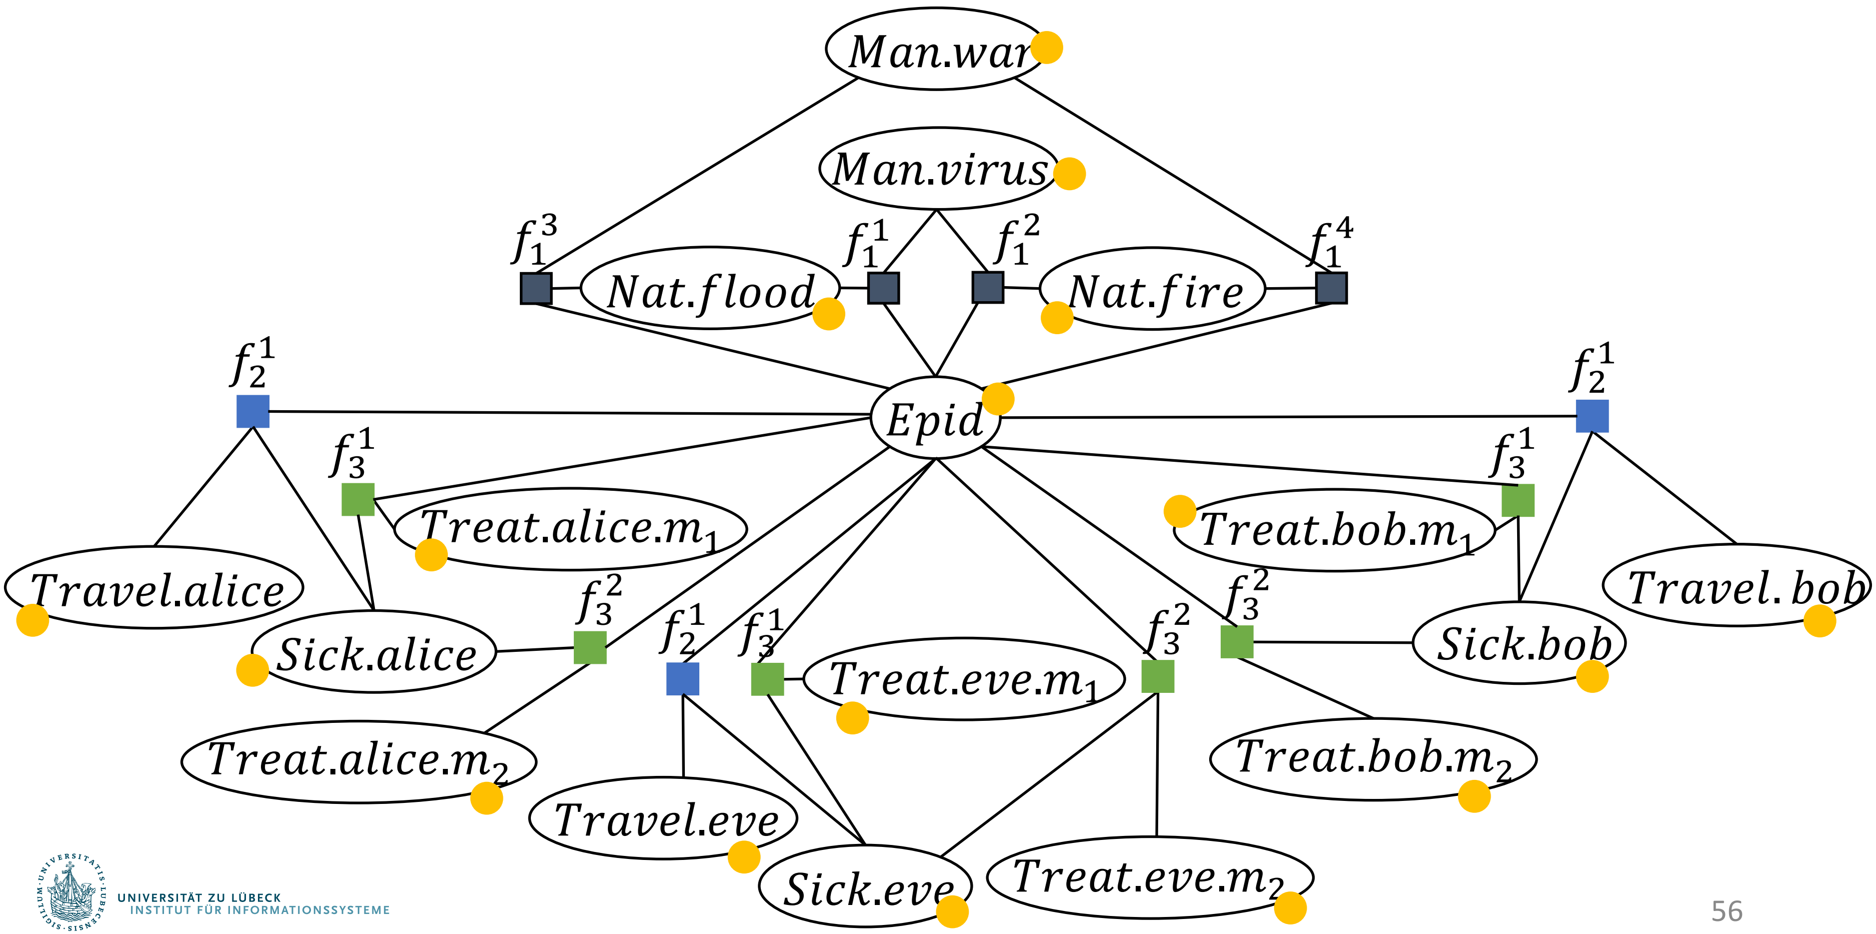

Compression

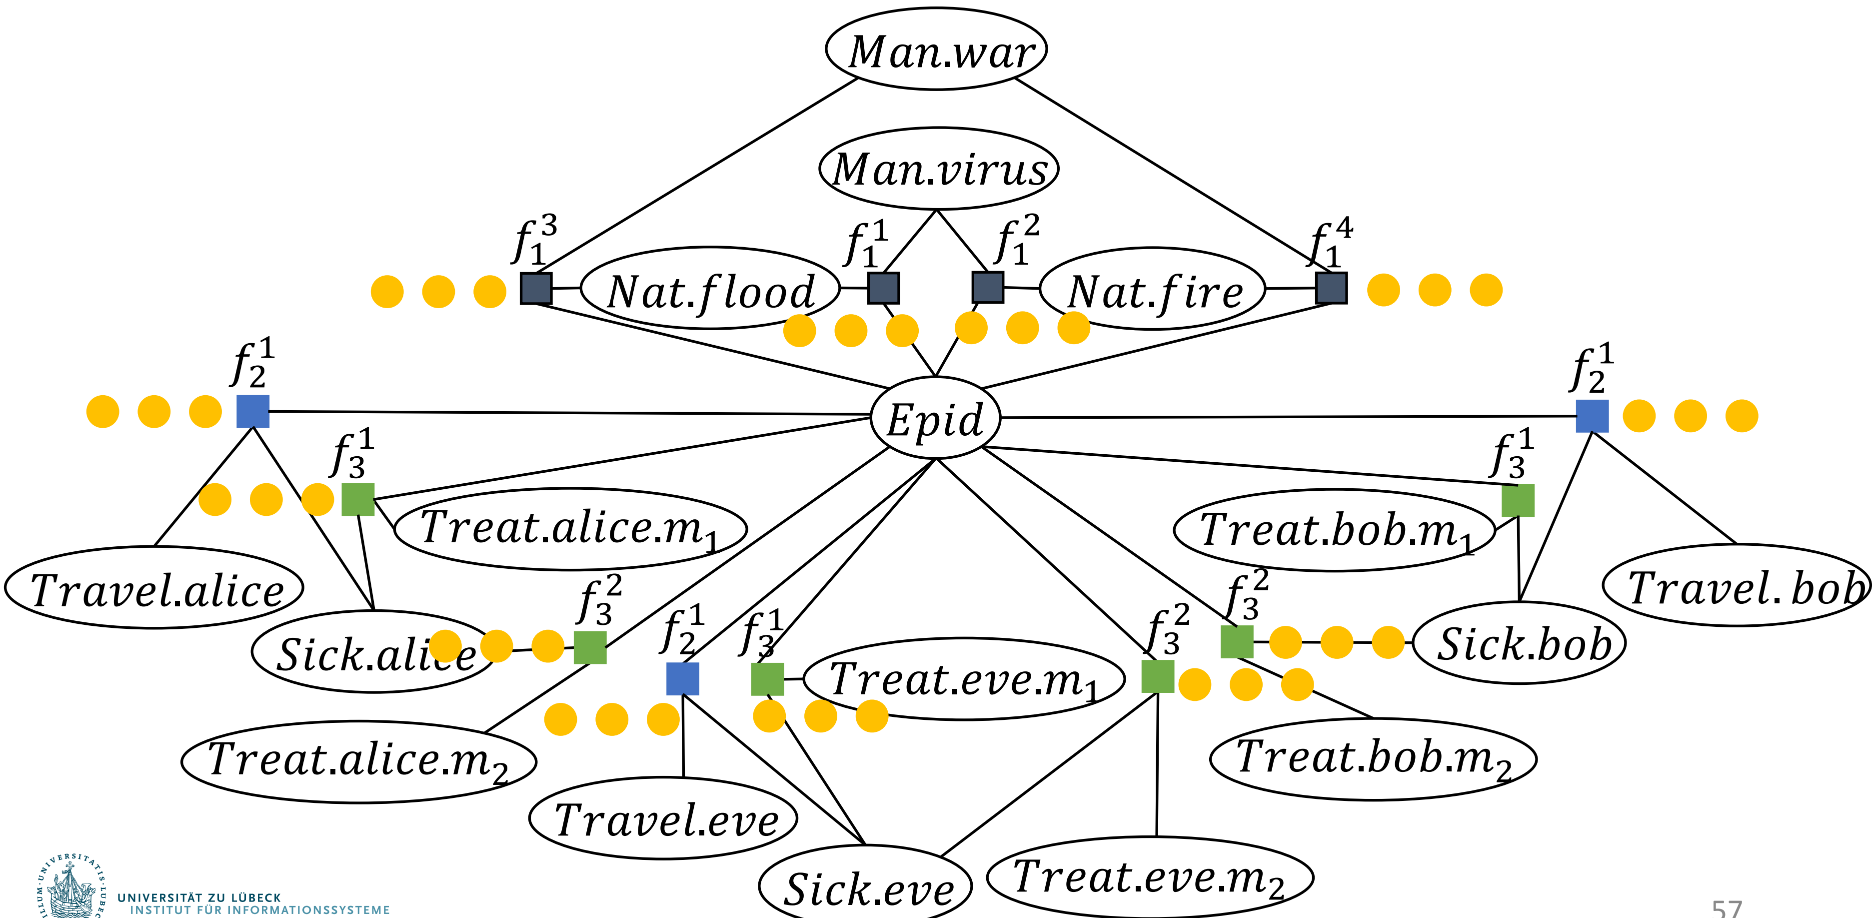

Compression

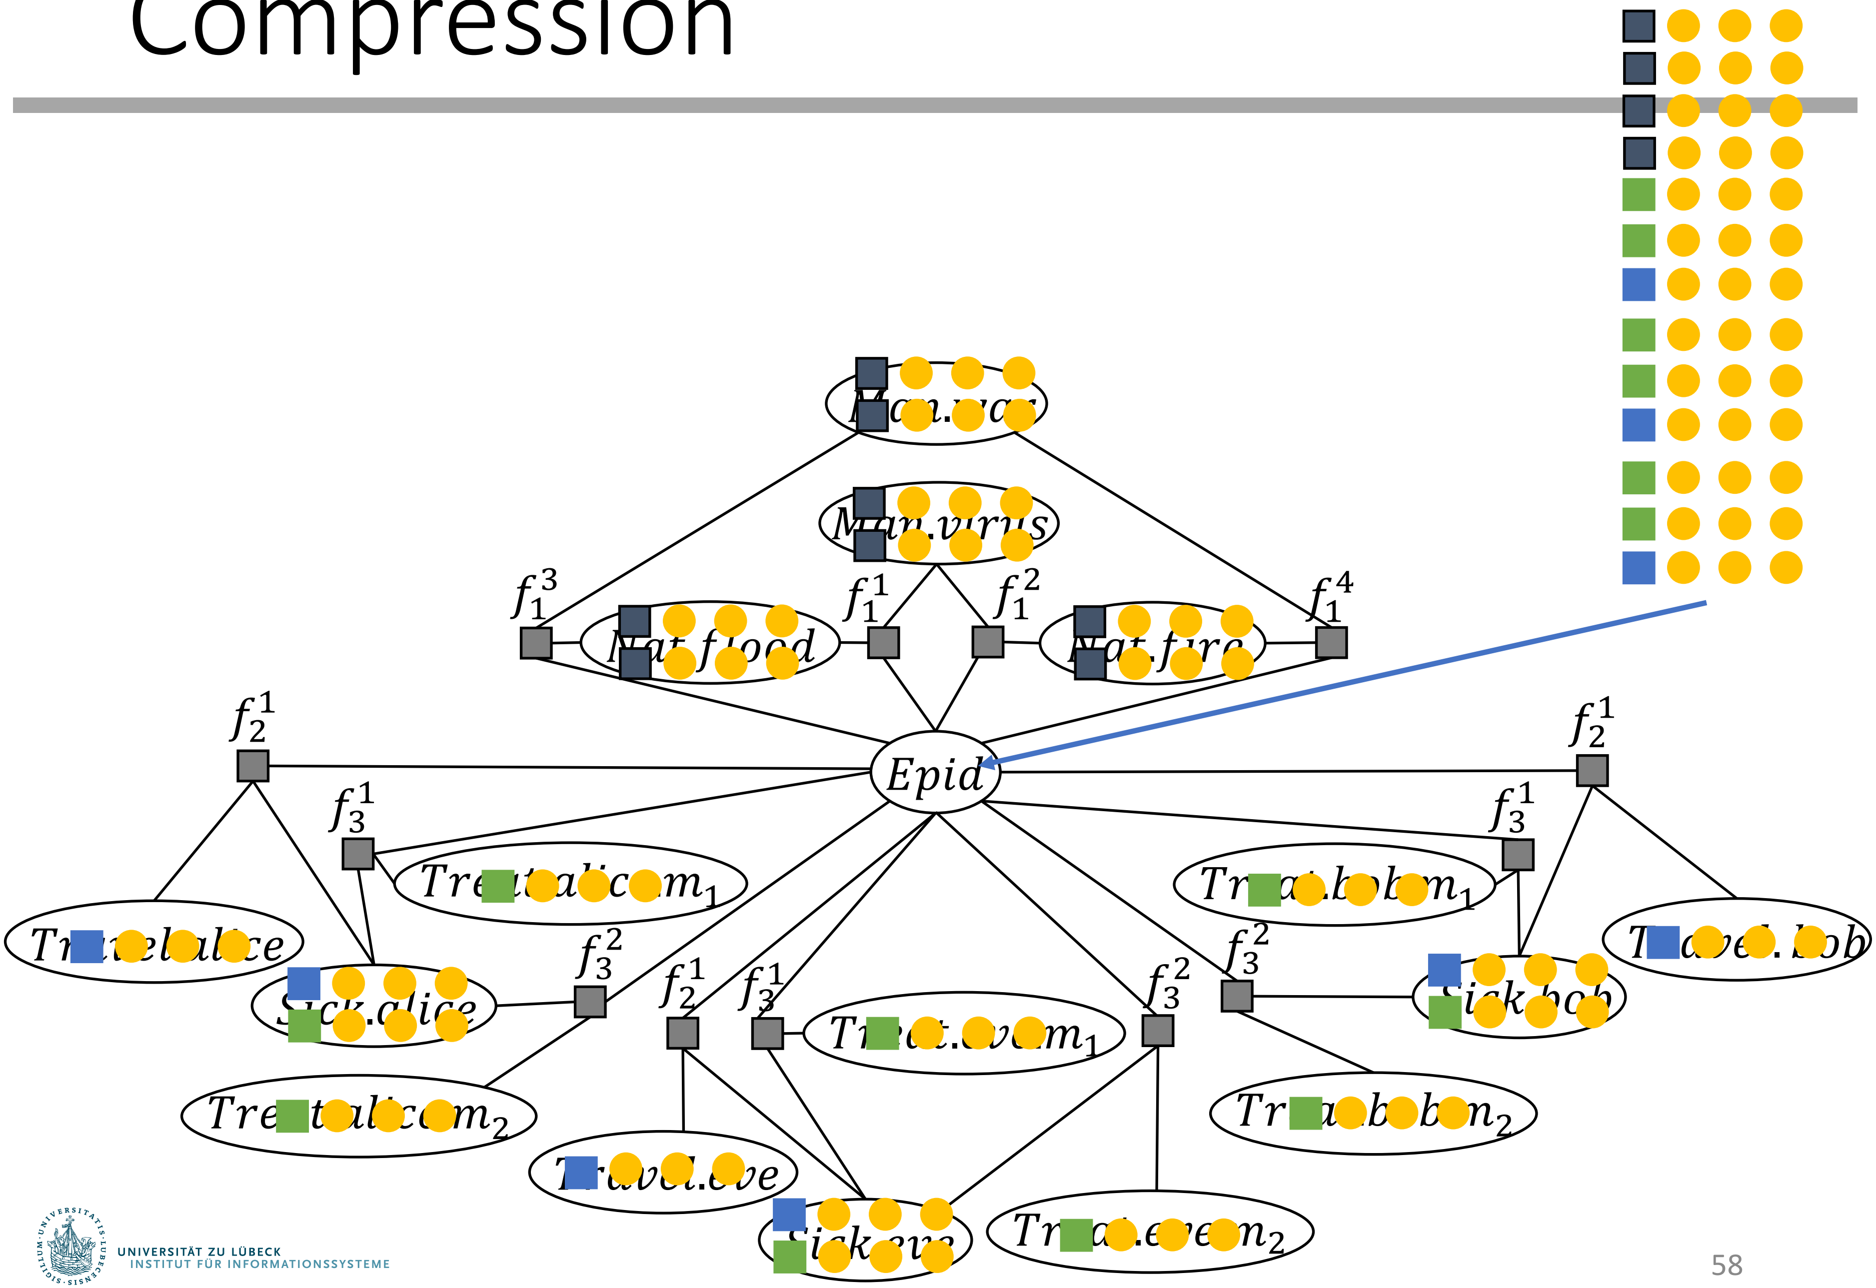

Compression

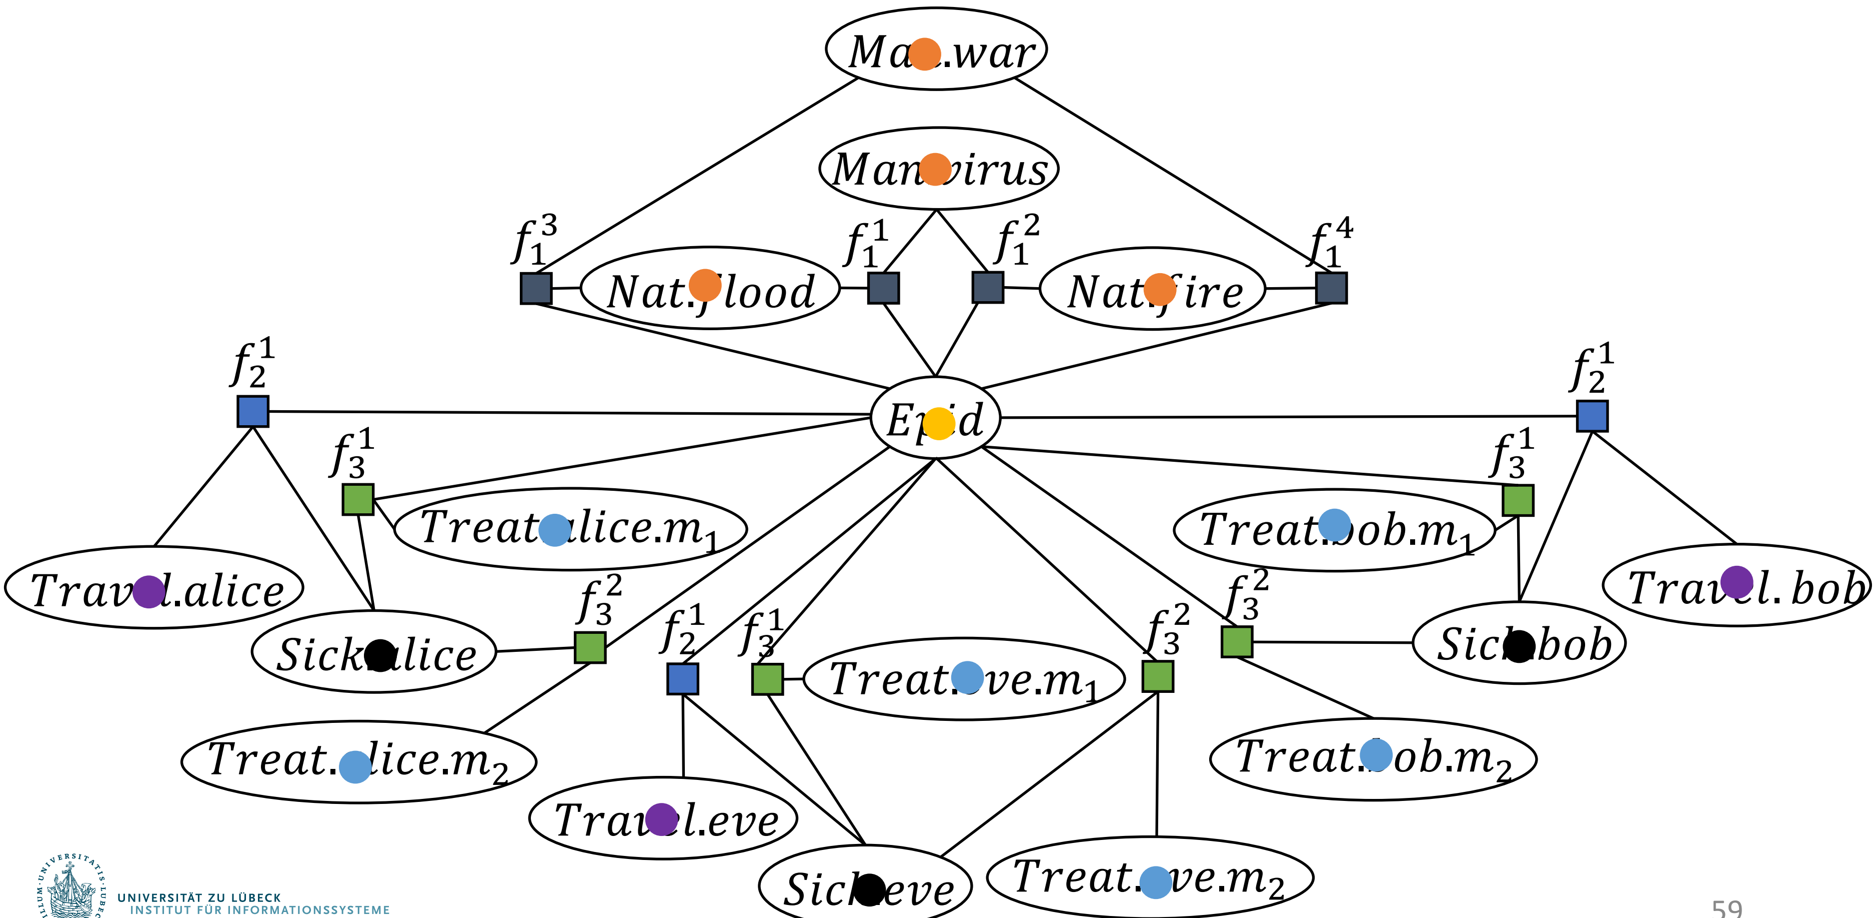

Compression

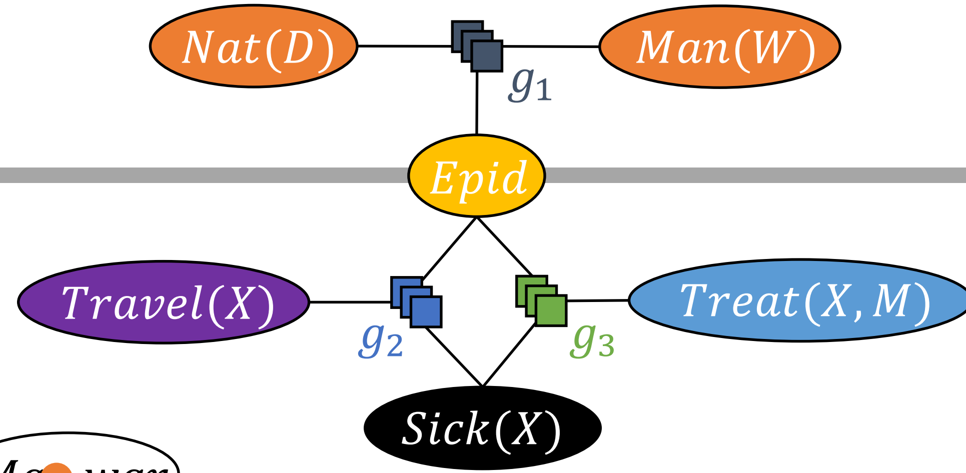

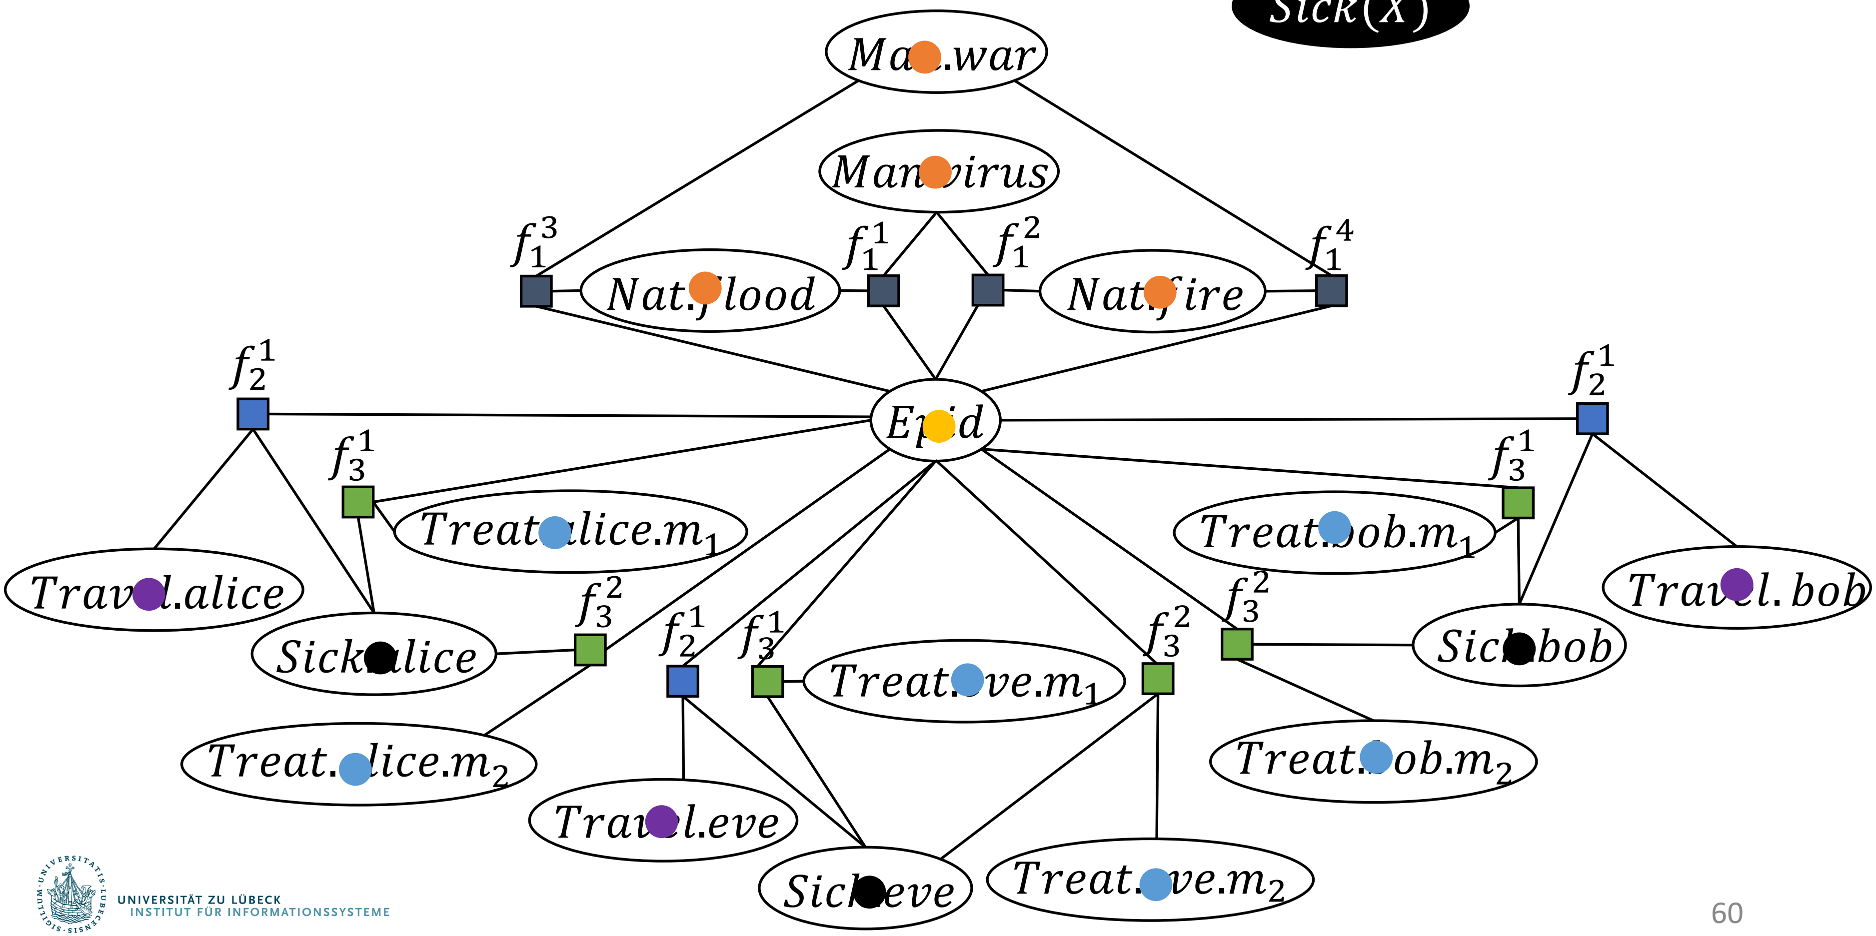

Lifting

- Factors with PRVs = parfactors
  - Undirected (graphical) Model G
  - E.g., g_2

$Travel(X)$	$Epid$	$Sick(X)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

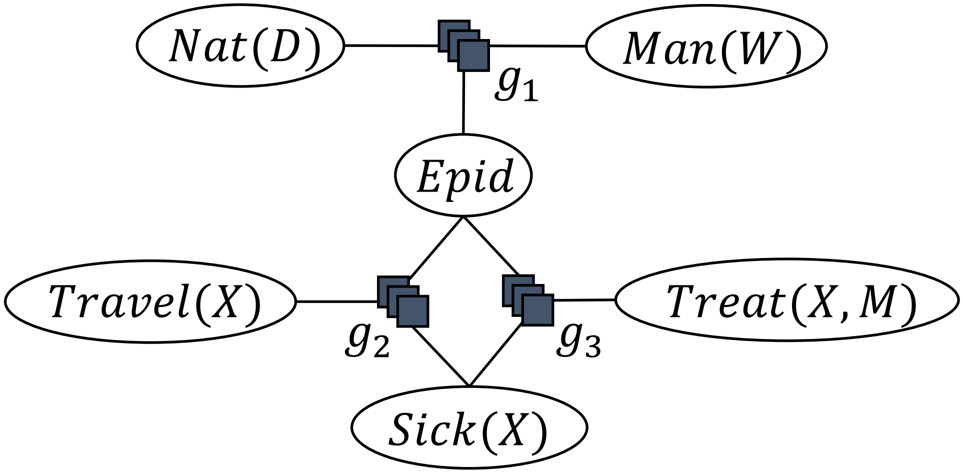

Grounding

- Grounding: replace logical variables with constants
  - e.g., $gr(g_2) = \{f_2^1, f_2^2, f_2^3\}$
- Semantics: ground + build full joint

$$P_G = \frac{1}{Z} \prod_{f \in gr(G)} f$$

$Travel(eve)$	$Epid$	$Sick(eve)$	f_2^1
false	false	false	5
false	false	true	0
false	true	false	4
false	true	true	6
true	false	false	4
true	false	true	6
true	true	false	2
true	true	true	9

$Travel(alice)$	Epi		
false	fals		
false	fals		
false	tru		
false	tru		
true	fals		
true	fals		
true	true	false	2
true	true	true	9

(bob)	$Epid$	$Sick(bob)$	f_2^3
se	false	false	5
se	false	true	0
se	true	false	4
se	true	true	6
ue	false	false	4
ue	false	true	6
ue	true	false	2
true	true	true	9

Lifted Decision Networks

- Decision parameterised model
  - Parfactor graph + utility nodes + action nodes
  - Example
    - Condition of water retention (C) correlated with weight (W)
      - LT = living together, S = scale works
    - Ranges for PRVs: true/false for $S(X)$, $LT(X, X')$; normal/deviation/retains water for $C(X)$; steady/falling/rising for $W(X)$
    - Action range: do_{not}, do_{vis} for do nothing, visit patient

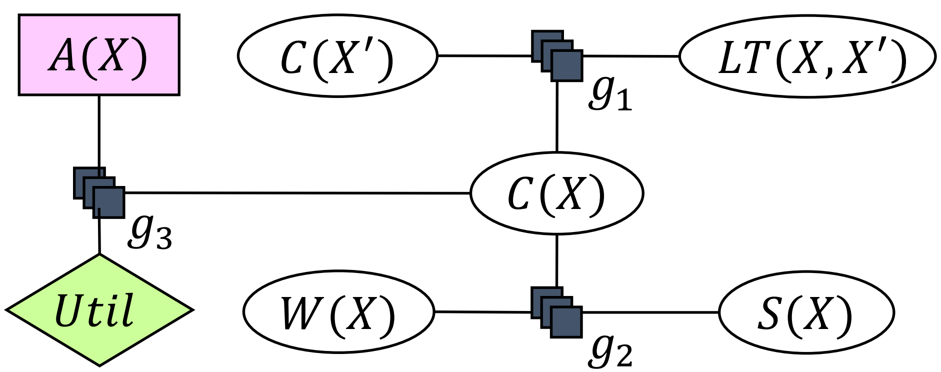

Evaluation: Example

- Evaluation as with propositional decision networks
  - Using lifted inference for eliminations
  - Compute actions at once for group of indistinguishable constants
    - No evidence = no distinguishable features
  - With $W(X) = true$ as evidence for some $\hat{X}$: two groups, four action "plans"
    - 2 actions · 2 groups

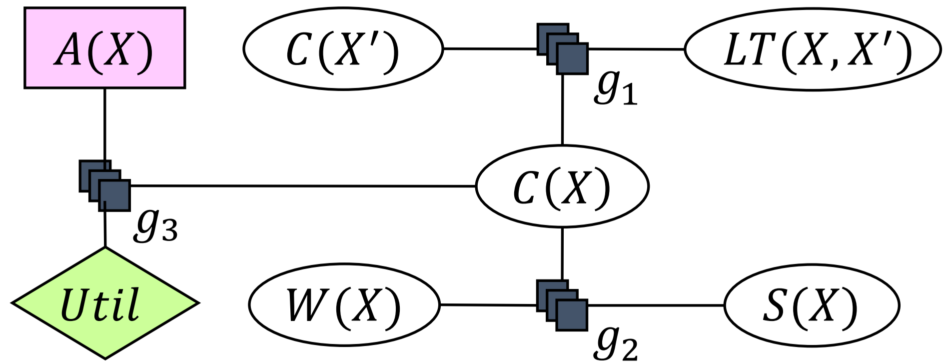

$$EU(A(X) = do_{not})$$

$$\propto \left(\sum_{c,c' \in \mathcal{R}(C(X))} g_3(A(X) = do_{not}, C(X) = c) \sum_{l \in \mathcal{R}(LT(X,X'))} g_1(C(X) = c, LT(X,X') = l, C(X') = c') \sum_{w \in \mathcal{R}(W(X))} \sum_{s \in \mathcal{R}(S(X))} g_2(W(X) = w, S(X) = s) \right)^{|dom(X)|}$$

$EU(A(X) = do_{vis})$ **MEU: same action for all X**

Lifted Decision Making

- Solving MEU/query answering problems intractable in general
  - Exponential in tree width of the graphical model
- Explicitly encoded symmetries allows for tractable inference in terms of domain sizes for logical variables
  - Polynomial in domain size

    Guy Van den Broeck: On the Completeness of First-order Knowledge Compilation for Lifted Probabilistic Inference, NIPS-11.

    Nima Taghipour, Jesse Davis, and Hendrik Blockeel: First-order Decomposition Trees, NIPS-13.
    - Of course: the goal should be linear and better
  - Tractability through exchangeability

    Mathias Niepert and Guy Van den Broeck: Tractability through Exchangeability: A New Perspective on Efficient Probabilistic Inference, AAAI-14.

Intermediate Summary

- Decision networks
  - Utilities, actions, random variables
  - Evaluation: for each action setting, eliminate everything else
- Value of information
  - How much is a piece of information worth?
- Relational domains
  - First-order constructs for compact representation
  - Same action for sets of indistinguishable constants

Outline

Utility theory

- Preferences
- Utilities
- Dominance
- Preference structure

Decision theory

- Decision networks
- Value of information
- Relational domains

$\Longrightarrow$ Next: Making Complex Decisions