Lecture 29: Decision Making Under Uncertainty

Last Lecture

Modeling temporal uncertainty (and failure) with DBNs

Combining Beliefs and Desires

Expected Utility:

$$EU(a|e) = \sum_{s'} P(Result(a) = s'|a, e)U(s')$$

Combining Beliefs and Desires

Expected Utility:

$$EU(a|e) = \sum_{s'} P(Result(a) = s'|a, e)U(s')$$

Combining Beliefs and Desires

Expected Utility:

$$EU(a|e) = \sum_{s'} P(Result(a) = s'|a, e)U(s')$$

Combining Beliefs and Desires

Expected Utility:

$$EU(a|e) = \sum_{s'} P(Result(a) = s'|a, e)U(s')$$

"A rational agent should choose actions which maximize its expected utility"

Constraints on Rational Preferences

$A \succ B$	Agent prefers A to B
$A \sim B$	Agent is indifferent to A and B
$A \succsim B$	Agent prefers A to B or is indifferent to them

A and B can be lotteries

Lotteries

$$L = [p_1, C_1; p_2, C_2; \ldots; p_n, C_n]$$

Constraints on Rational Preferences

- **Orderability**: $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \succ B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Constraints on Rational Preferences

- **Orderability**: $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \succ B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Constraints on Rational Preferences

- **Orderability**: $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \succ B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Monotonicity

$$A \succ B$$
$$B \succ C$$
$$C \succ A$$

Constraints on Rational Preferences

- **Orderability**: Exactly one of $(A \succ B), (B \succ A)$, or $(A \sim B)$ must hold
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \sim B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Constraints on Rational Preferences

- **Orderability**: Exactly one of $(A \succ B), (B \succ A)$, or $(A \sim B)$ must hold
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \sim B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Constraints on Rational Preferences

- **Orderability**: Exactly one of $(A \succ B), (B \succ A)$, or $(A \sim B)$ must hold
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \sim B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Constraints on Rational Preferences

- **Orderability**: Exactly one of $(A \succ B), (B \succ A)$, or $(A \sim B)$ must hold
- **Transitivity:** $(A \succ B) \land (B \succ C) \implies (A \succ C)$
- **Continuity:** $A \succ B \succ C \implies \exists p[p, A; 1-p, C] \sim B$
- **Substitutability:** $A \sim B \implies [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- **Monotonicity:** $A \succ B \implies (p \geq q \iff [p, A; 1-p, B] \succsim [q, A; 1-q, B])$
- **Decomposability:**
  $[p, A; 1-p, [q, B; 1-q, C]] \sim [p, A; (1-p)q, B; (1-p)(1-q), C]$

Decomposability

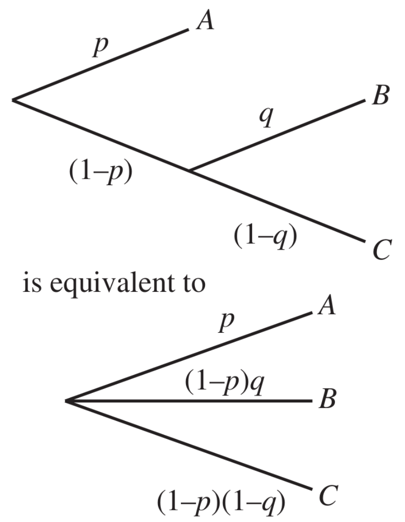

From Preferences to Utility

- **Existence of a Utility Function**: There exists a function U such that: $U(A) > U(B) \iff A \succ B, \quad U(A) = U(B) \iff A \sim B$
- **Expected Utility of a Lottery**: The utility of a lottery is the sum of the probability of each outcome times the utility of that outcome. $U([p_1, S_1; \dots; p_n, S_n]) = \sum_i p_i U(s_i)$

(**Proof?** *von Neumann, J. and Morgenstern, O. (1944). Theory of Games and Economic Behavior Princeton University Press*)

From Preferences to Utility

- **Existence of a Utility Function**: There exists a function U such that: $U(A) > U(B) \iff A \succ B, \quad U(A) = U(B) \iff A \sim B$
- **Expected Utility of a Lottery**: The utility of a lottery is the sum of the probability of each outcome times the utility of that outcome. $U([p_1, S_1; \ldots; p_n, S_n]) = \sum_i p_i U(s_i)$

(**Proof?** *von Neumann, J. and Morgenstern, O. (1944). Theory of Games and Economic Behavior Princeton University Press*)

From Preferences to Utility

- **Existence of a Utility Function**: There exists a function U such that: $U(A) > U(B) \iff A \succ B, \quad U(A) = U(B) \iff A \sim B$
- **Expected Utility of a Lottery**: The utility of a lottery is the sum of the probability of each outcome times the utility of that outcome. $U([p_1, S_1; \ldots; p_n, S_n]) = \sum_i p_i U(s_i)$

(**Proof?** *von Neumann, J. and Morgenstern, O. (1944). Theory of Games and Economic Behavior Princeton University Press*)

What the Axioms Don't give you 1: Guidance on Arbitrary Preference

"I prefer to have a prime number in my bank account; When i have £10, I will give away £3"

Utility of Money

Monotonic preference towards money

Utility of Money

Monotonic preference towards money

What about lotteries?

Utility of Money

A: I give you £1,000,000

B: Toss a coin: Heads, I give you £3,000,000; Tails, you get nothing

Utility of Money

A: I give you £1,000,000

B: Toss a coin: Heads, I give you £3,000,000; Tails, you get nothing

Expected *monetary* value **A**: 1,000,000

Expected *monetary* value **B**: $0.5 \times 0 + 0.5 \times 3,000,000 = 1,500,000$

Utility of Money

$S_k -$ state of possessing £k

$$EU(A) = U(S_{k+1M})$$
$$EU(B) = \frac{1}{2}U(S_k) + \frac{1}{2}U(S_{k+3M})$$

Utility of Money

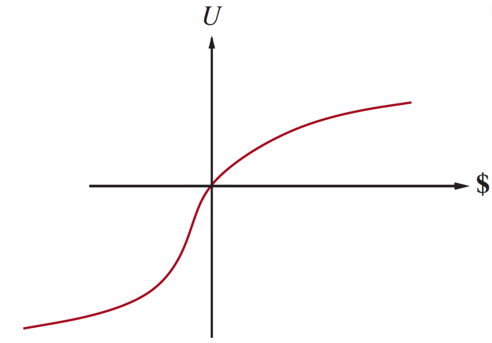

What the Axioms Don't Give You 2: Uniqueness

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

What the Axioms Don't Give You 2: Uniqueness

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

$$EU(a_1|e) = \sum_{s'} P(Result(a_1) = s'|a, e)U'(s')$$
$$EU(a_2|e) = \sum_{s'} P(Result(a_2) = s'|a, e)U'(s')$$

What the Axioms Don't Give You 2: Uniqueness

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

$$EU(a_1|e) = \sum_{s'} P(Result(a_1) = s'|a, e)(k_1 + k_2 U(s'))$$
$$EU(a_2|e) = \sum_{s'} P(Result(a_2) = s'|a, e)(k_1 + k_2 U(s'))$$

What the Axioms Don't Give You 2: Uniqueness

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

$$EU(a_1|e) = k_1 + k_2 \sum_{s'} P(Result(a_1) = s'|a, e) U(s')$$
$$EU(a_2|e) = k_1 + k_2 \sum_{s'} P(Result(a_2) = s'|a, e) U(s')$$

A Strategy to Construct Utility Functions

Normalized Utility:

"Best Possible Outcome" ($u^{\top} = 1$)

"Worst Possible Catastrophe" ($u^{\perp} = 0$)

Figuring out the utility of S:

- Offer a lottery: $[p, u^{\top}; 1 - p, u^{\perp}]$
- Adjust p until $[p, u^{\top}; 1 - p, u^{\perp}] \sim S$
- Set $U(S) = p$

Decision Networks (Influence Diagrams)

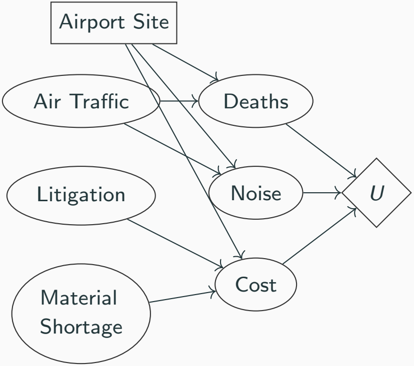

- Chance Nodes (ovals)
- Decision Nodes (rectangles)
- Utility Nodes (diamonds)

Decision Networks (Influence Diagrams)

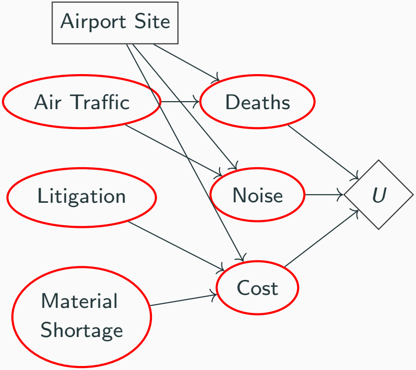

- Chance Nodes (ovals)
- Decision Nodes (rectangles)
- Utility Nodes (diamonds)

Decision Networks (Influence Diagrams)

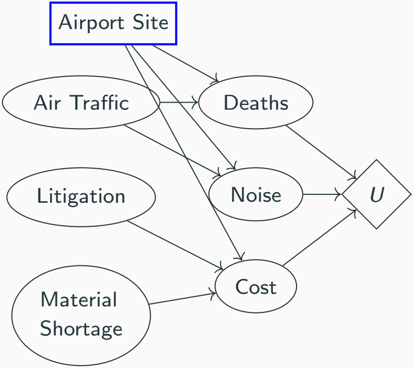

- Chance Nodes (ovals)
- Decision Nodes (rectangles)
- Utility Nodes (diamonds)

Decision Networks (Influence Diagrams)

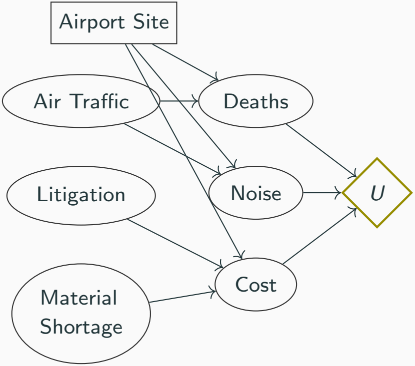

- Chance Nodes (ovals)
- Decision Nodes (rectangles)
- Utility Nodes (diamonds)

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
    2.1 Set decision node to that value
    2.2 Calculate posterior probabilities for parents of utility node
    2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

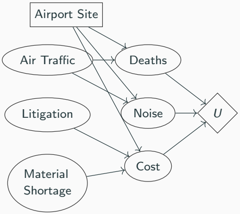

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
    2.1 Set decision node to that value
    2.2 Calculate posterior probabilities for parents of utility node
    2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

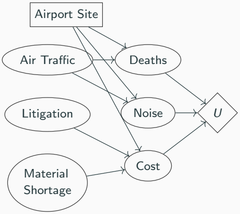

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
    2.1 Set decision node to that value
    2.2 Calculate posterior probabilities for parents of utility node
    2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

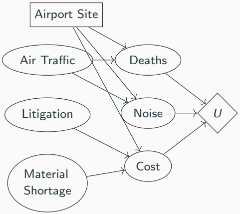

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
    2.1 Set decision node to that value
    2.2 Calculate posterior probabilities for parents of utility node
    2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

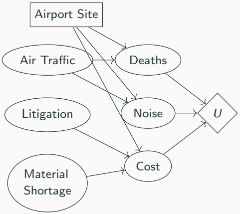

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
   2.1 Set decision node to that value
   2.2 Calculate posterior probabilities for parents of utility node
   2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

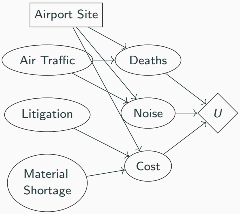

Evaluating Expected Utility with a Decision Network

1. Set evidence variables for current state
2. For each value of decision node
    2.1 Set decision node to that value
    2.2 Calculate posterior probabilities for parents of utility node
    2.3 Calculate resulting expected utility for action
3. Return action with highest expected utility

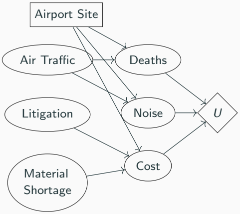

Summary

- Utility Theory, Axioms and Criticisms
- Decision Networks for Expected Utility