Learning Analytics

Dr. Bowen Hui

Computer Science

University of British Columbia Okanagan

Last Class: General Decision Problems

- Probability theory
  - Estimate state of the world
- Utility theory
  - Quantifies preferences over outcomes
- Expected utility
  - Evaluates actions

Preferences

- A preference ordering $\succcurlyeq$ is a ranking of all possible states (worlds), S
- Specifically, for any two states *s* and *t*:
  - $s \succcurlyeq t$ means that *s* is at least as good as *t*
  - $s \succ t$ means that *s* is strictly preferred to *t*
  - $s \sim t$ means that the agent is indifferent between *s* and *t*

Axioms on Preferences

- Transitivity: Given 3 states A, B, C:

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

- Why is this important?
  - Suppose you strictly prefer coffee to tea, tea to OJ, OJ to coffee

What's wrong with this?

Axioms on Preferences

- Transitivity: Given 3 states A, B, C:
  $$(A \succ B) \land (B \succ C) \Rightarrow (A \succ C)$$
- Why is this important?
  - Suppose you strictly prefer coffee to tea, tea to OJ, OJ to coffee
  - If you prefer X to Y, you'll trade me Y + $1 for X

Axioms on Preferences

- Transitivity: Given 3 states A, B, C:

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

- Why is this important?
  - Suppose you strictly prefer coffee to tea, tea to OJ, OJ to coffee
  - If you prefer X to Y, you'll trade me Y + $1 for X
  - I can construct a "money pump" and extract lots of money from you!

- Others axioms: orderability, continuity, substitutability, monotonicity, decomposibility

Decision Problems: *Certainty*

- A decision problem under certainty consists of:
  - A set of decisions, D
    E.g. actions you can take in a plan
  - A set of outcomes or states, S
    E.g. states you can reach by executing a plan
  - An outcome function, $f: D \rightarrow S$
    E.g. the outcome of a decision
  - A preference ordering $\succcurlyeq$ over S

- A solution to a decision problem is any $d^* \in D$ s.t. $f(d^*) \succcurlyeq f(d)$ for all $d \in D$

Example: Decision Scenarios

		Variables	
		It rains	It doesn't rain
Actions	Take umbrella	Encumbered, Dry	Encumbered, Dry
	Leave umbrella	Wet	Free, Dry

- 2 variables x 2 actions = 4 possible outcomes
- Preference over outcomes:
  Free,Dry $\succ$ Encumbered,Dry $\succ$ Wet

- Later: quantify the strength of these preferences

Decision Making Under *Uncertainty*

- Suppose outcomes are not deterministic
  - E.g. bring umbrella, but 5% chance it gets stolen
- Or actions are not deterministic
  - E.g. pour coffee, but 20% of time it spills

- How to decide what to do?
  - Decision involves figuring out:
    - How likely is each outcome (probability)
    - How good is each outcome (utility)

Example: Decision Scenarios

		States	
		It rains	It doesn't rain
Actions	Take umbrella	Encumbered, Dry	Encumbered, Dry
	Leave umbrella	Wet	Free, Dry

- 2 states x 2 actions = 4 possible outcomes

Example: Decision Scenarios

		States	
		It rains	It doesn't rain
Actions	Take umbrella	Encumbered, Dry +7	Encumbered, Dry +5
	Leave umbrella	Wet -8	Free, Dry +10

- 2 states x 2 actions = 4 possible outcomes
- Utility = a real number value of each outcome

Example: Decision Scenarios

		States	
		It rains 0.4	It doesn't rain 0.6
Actions	Take umbrella	Encumbered, Dry +7	Encumbered, Dry +5
	Leave umbrella	Wet -8	Free, Dry +10

- 2 states x 2 actions = 4 possible outcomes
- Utility = a real number value of each outcome
- Each state has a probability

Example: Decision Scenarios

		States	
		It rains 0.4	It doesn't rain 0.6
Actions	Take umbrella	Encumbered, Dry +7	Encumbered, Dry +5
	Leave umbrella	Wet -8	Free, Dry +10

- 2 states x 2 actions = 4 possible outcomes
- Utility = a real number value of each outcome
- Each state has a probability

- EU(takeUmb) = Pr(rain)U(encDry) + Pr(noRain)U(encDry)
  = (0.4)(7) + (0.6)(5) = 5.8

Example: Decision Scenarios

		States	
		It rains 0.4	It doesn't rain 0.6
Actions	Take umbrella	Encumbered, Dry +7	Encumbered, Dry +5
	Leave umbrella	Wet -8	Free, Dry +10

- 2 states x 2 actions = 4 possible outcomes
- Utility = a real number value of each outcome
- Each state has a probability

- EU(leaveUmb) = Pr(rain)U(wet) + Pr(noRain)U(freeDry)
  = (0.4)(-8) + (0.6)(10) = 2.8

Example: Decision Scenarios

		States	
		It rains 0.4	It doesn't rain 0.6
Actions	Take umbrella	Encumbered, Dry +7	Encumbered, Dry +5
	Leave umbrella	Wet -8	Free, Dry +10

- 2 states x 2 actions = 4 possible outcomes
- Utility = a real number value of each outcome
- Each state has a probability

- EU(takeUmb) = 5.8
- EU(leaveUmb) = 2.8
- Which action will you choose?

Utilities

- Rather than just ranking outcomes, we quantify our degree of preferences
  - How much more important is *s* over *t*

- A utility function, $U : S \to \mathbb{R}$ associates a real-valued utility with each outcome
  - U(s) measures your degree of preference for s

- Note: U induces a preference ordering $\succcurlyeq_U$ over S defined as: s $\succcurlyeq_U$ t iff $U(s) \geq U(t)$

Expected Utility

- Under conditions of uncertainty, each decision induces a distribution Pr_d over possible outcomes
  - $Pr_d(s)$ is the probability of outcome s under decision d

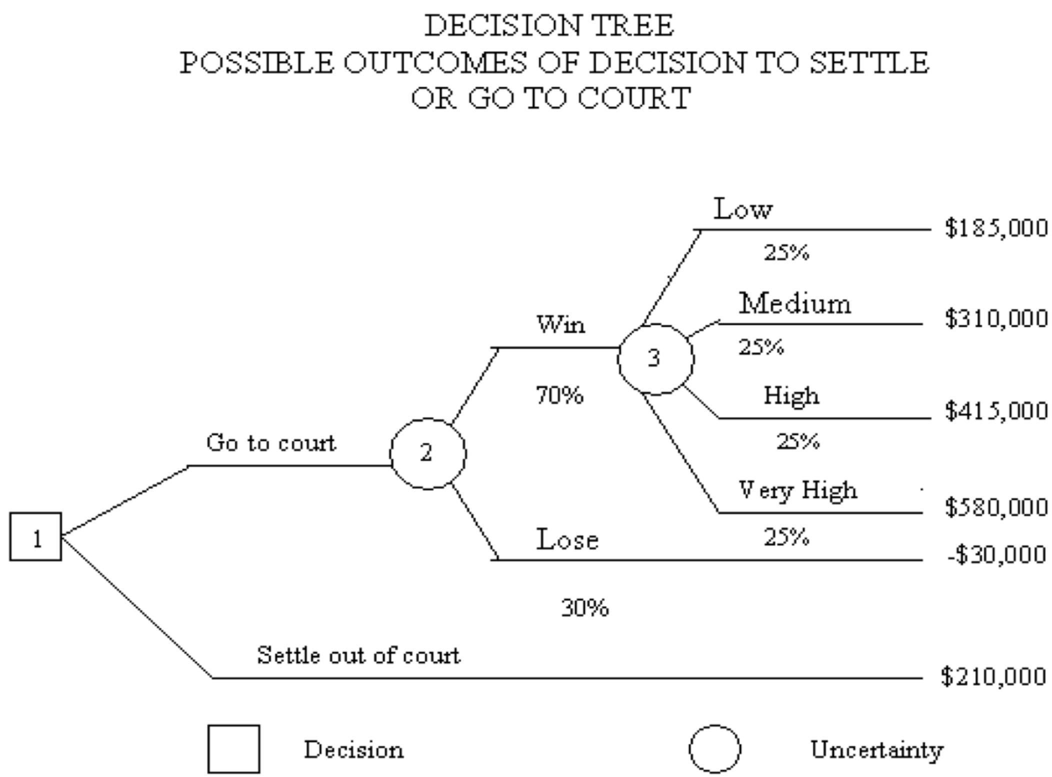

Image taken from www.mediate.com

Expected Utility

- Under conditions of uncertainty, each decision induces a distribution Pr_d over possible outcomes
  - $Pr_d(s)$ is the probability of outcome s under decision d

- The expected utility of decision d is defined as:

$$EU(d) = \sum_{s \in S} Pr_d(s) U(s)$$

Computing Expected Utility

DECISION TREE
POSSIBLE OUTCOMES OF DECISION TO SETTLE
OR GO TO COURT

Expected utility of winning:

$p_1u_1 + p_2u_2 + p_3u_3 + p_4u_4$
$= \$372,500$

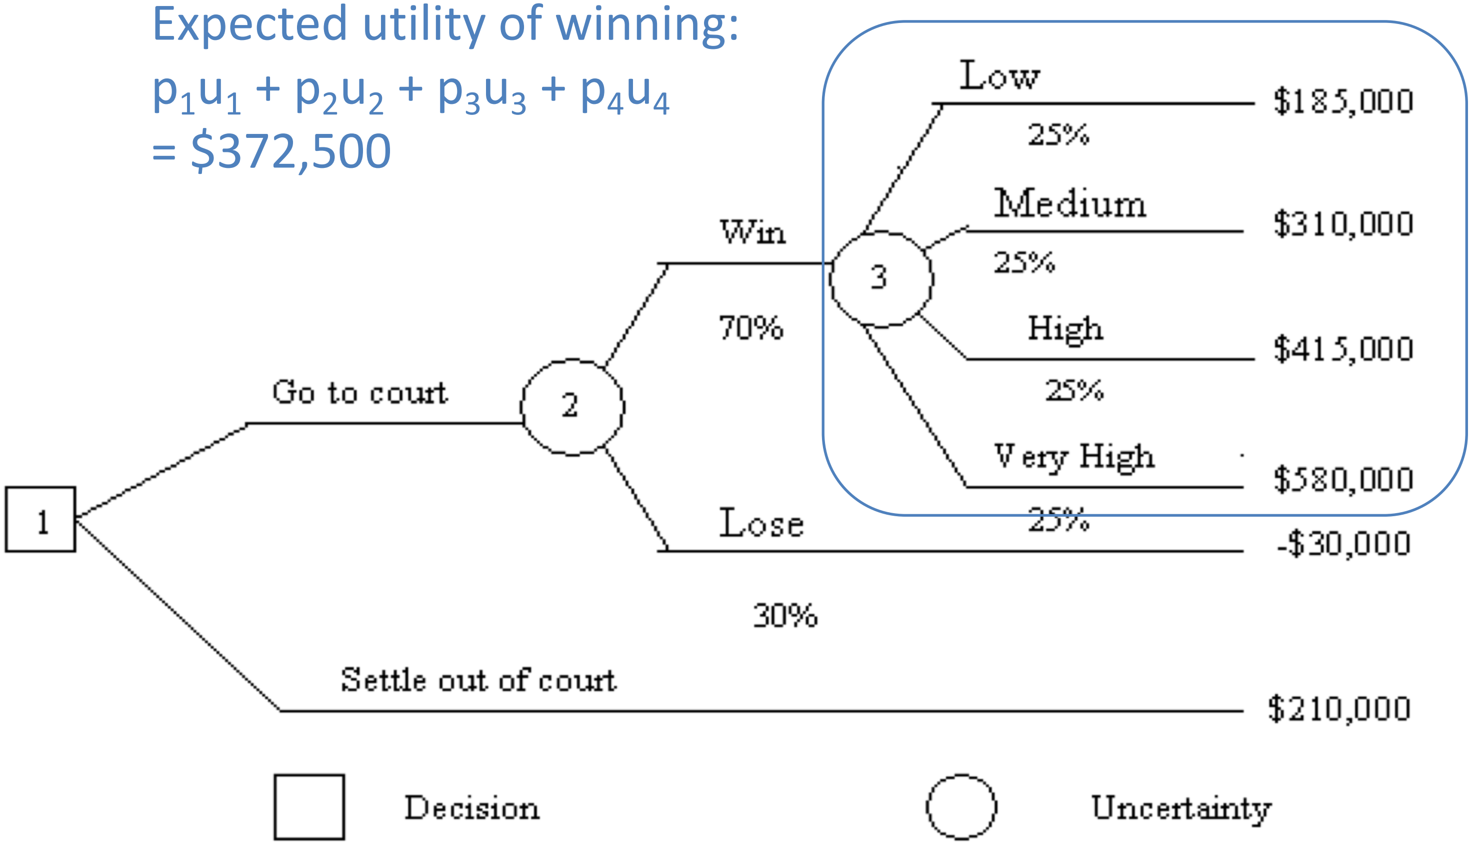

Image taken from www.mediate.com

Computing Expected Utility

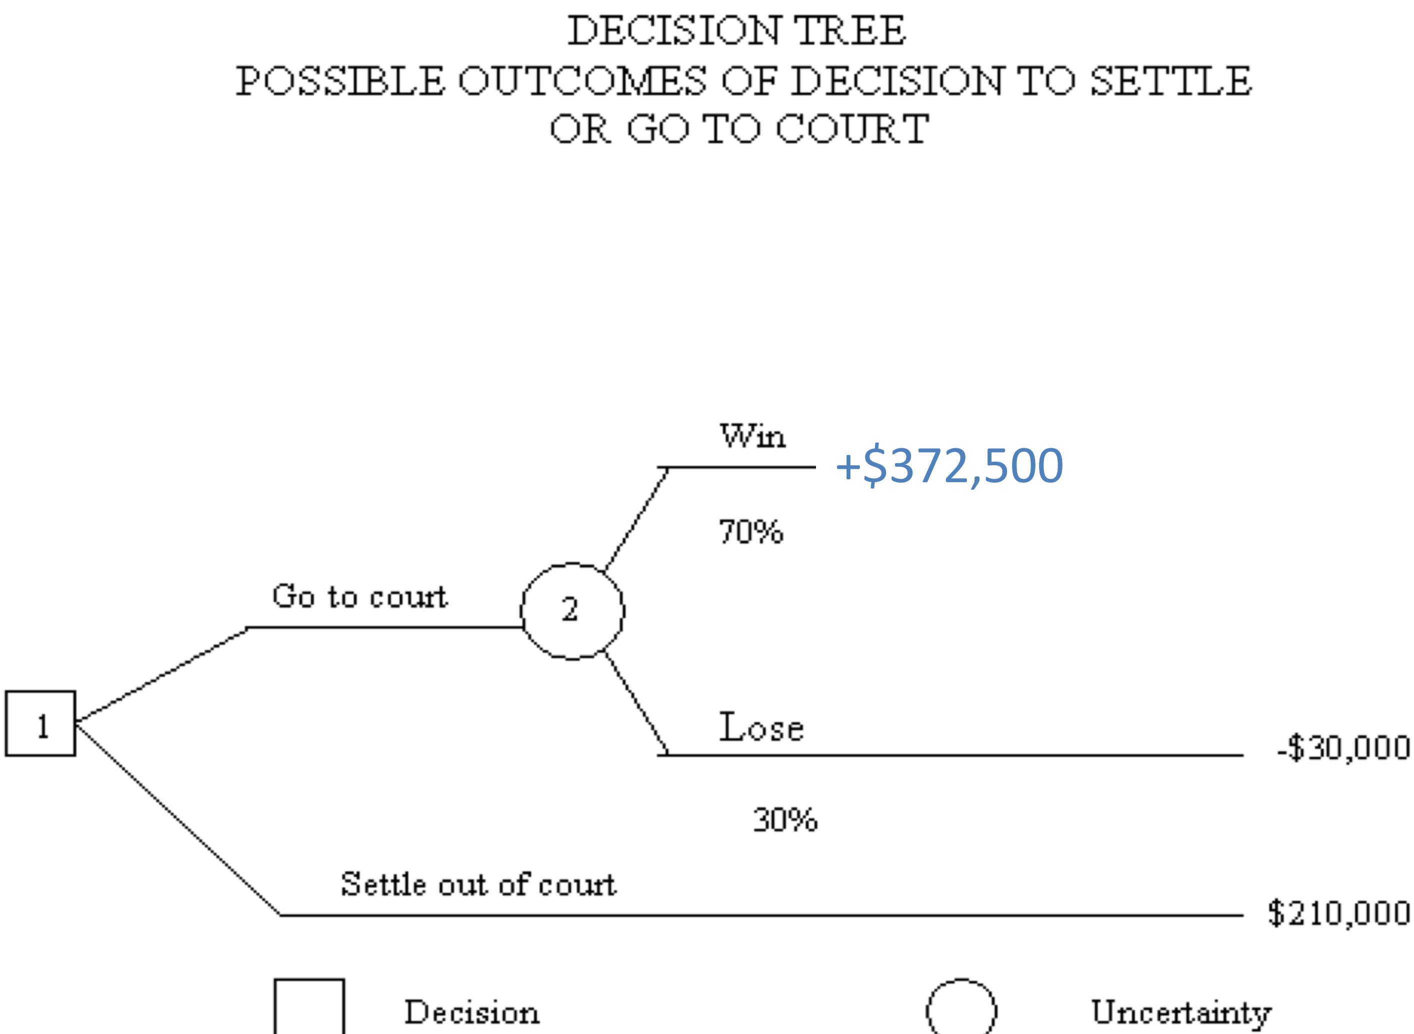

Image taken from www.mediate.com

Computing Expected Utility

DECISION TREE
POSSIBLE OUTCOMES OF DECISION TO SETTLE
OR GO TO COURT

Expected utility of going to court:

$p_1u_1 + p_2u_2$

$= \$251,750$

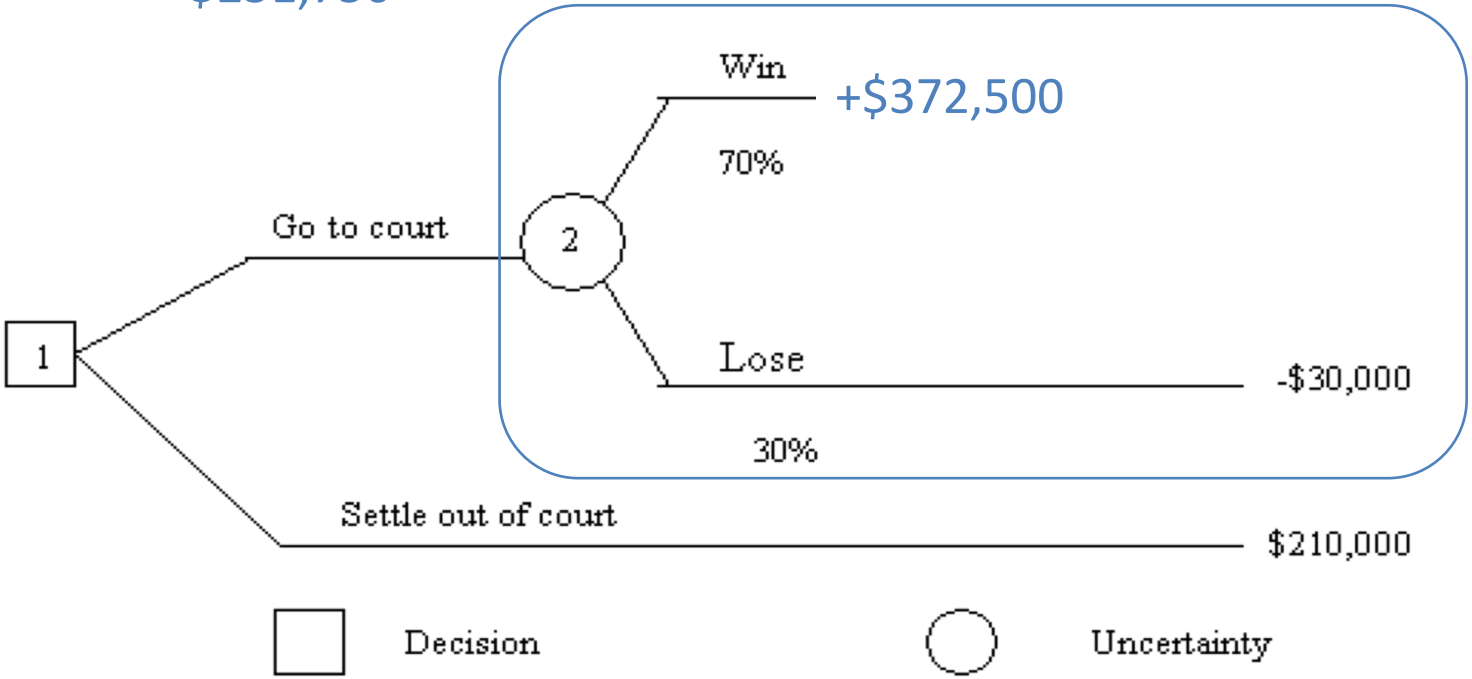

Image taken from www.mediate.com

Computing Expected Utility

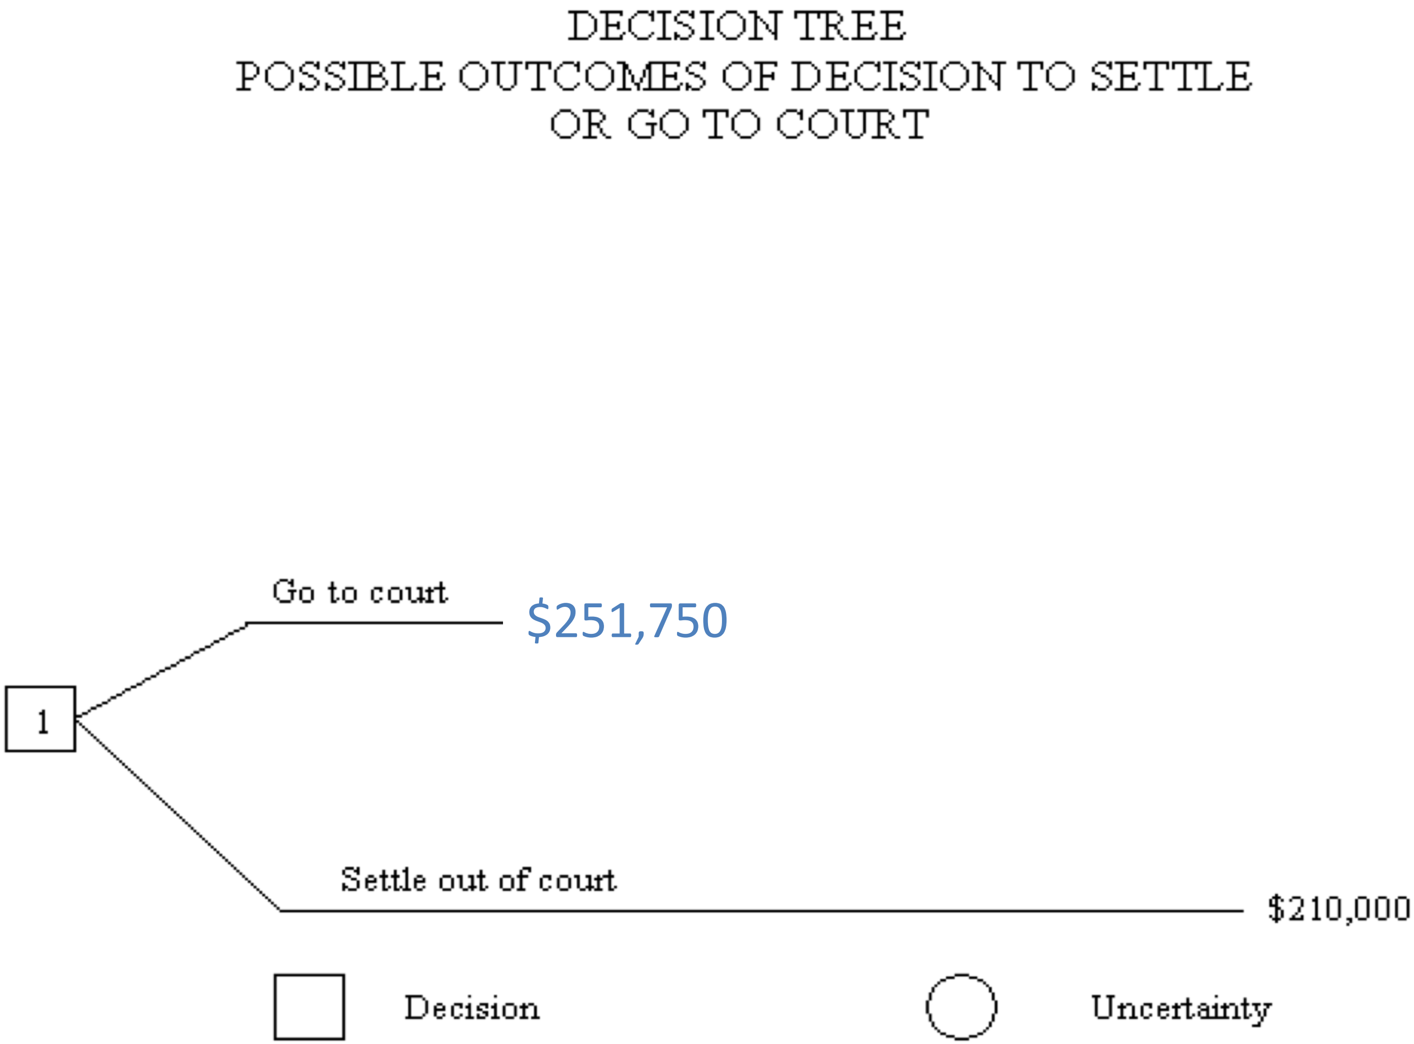

Image taken from www.mediate.com

Bayesian Action Selection

- Expected utility accommodates for uncertainty modeled in system's belief

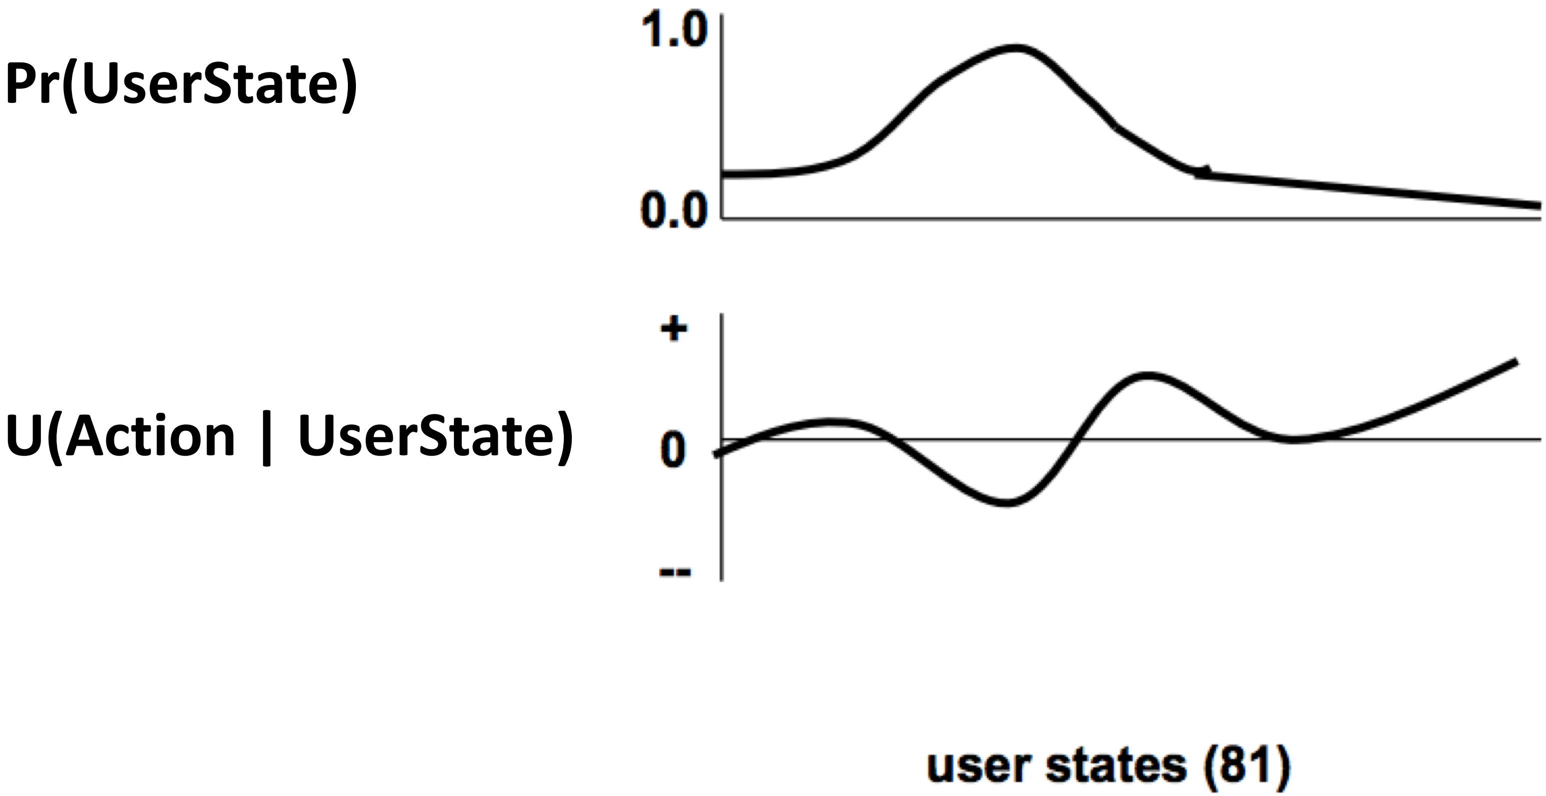

The MEU Principle

- Principle of maximum expected utility (MEU) states that the optimal decision under conditions of uncertainty is that the one with the highest expected utility

- Decision making under rational choice
  - "Subjective" variables to model in utility function

Decision Problems: *Uncertainty*

- A decision problem under certainty consists of:
  - A set of decisions, D
  - A set of outcomes or states, S
  - An outcome function, Pr: D → $\Delta(S)$
    - $\Delta(S)$ is the set of distributions over S (e.g. Pr_d)
    - Under certainty, the outcome function was f: D → S
  - A utility function U over S
    - Under certainty, we had a preference ordering $\succcurlyeq$ over S

- A solution to a decision problem is any d* ∈ D s.t. EU(d*) ≥ EU(d) for all d ∈ D

EU Modeling Uncertainty

- EU accounts for both
  - Uncertainty in action outcomes
  - Uncertainty in state of knowledge
  - Combination of the two

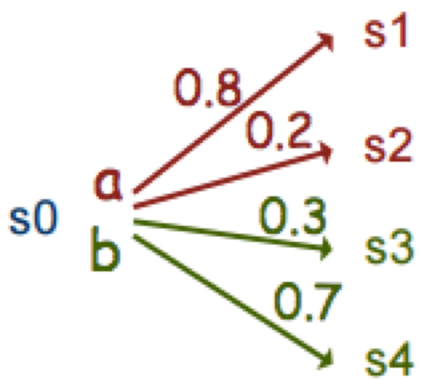

Stochastic actions

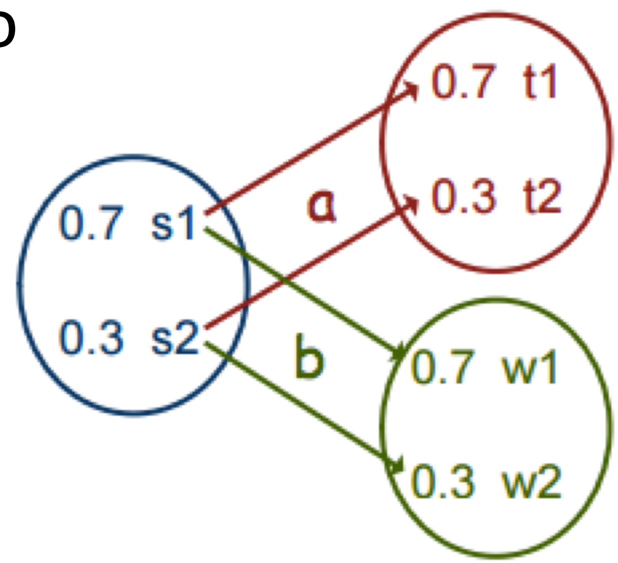

Uncertain knowledge

Where Do Utilities Come From?

- Utility theory tightly coupled with action/choice

- Preference elicitation – ask (as few as possible) structured queries to determine preferences between specific scenarios

- Utility functions needn't be unique
  - U is unique up to positive affine transformation
  - All decisions have the same relative utility
  - Adding a constant to U or multiplying U by a positive constant doesn't change the modeled preferences

Computational Bottlenecks

- Outcome space is large
  - State space can be huge
  - Don’t spell out distributions explicitly
  - Solution: Use Bayes nets to exploit independence

Computational Bottlenecks

- Outcome space is large
  - State space can be huge
  - Don't spell out distributions explicitly
  - Solution: Use Bayes nets to exploit independence

- Decision space is large
  - Single-shot actions are pretty simple
  - More commonly, problems involve sequential choices (e.g. plans)
  - Viewing each plan as a distinct decision results in huge decision space
  - Solution: Use dynamic programming methods to construct optimal plans

Illustrative Example

- Suppose we have two actions: a, b
- We have to decide to execute two actions in sequence, which is one of:
  - *[a,a], [a,b], [b,a], [b,b]*
- These actions are stochastic:
  - Action *a* induces distribution $Pr_a(s_i|s_j)$ over states
  - E.g. $Pr_a(s_2|s_1) = 0.9$ means probability of moving to state s_2 when *a* is performed at s_1 is 0.9
- Which sequence of actions to take?

Distributions for Action Sequences

- More actions means more branches
- Longer horizon means longer sequence

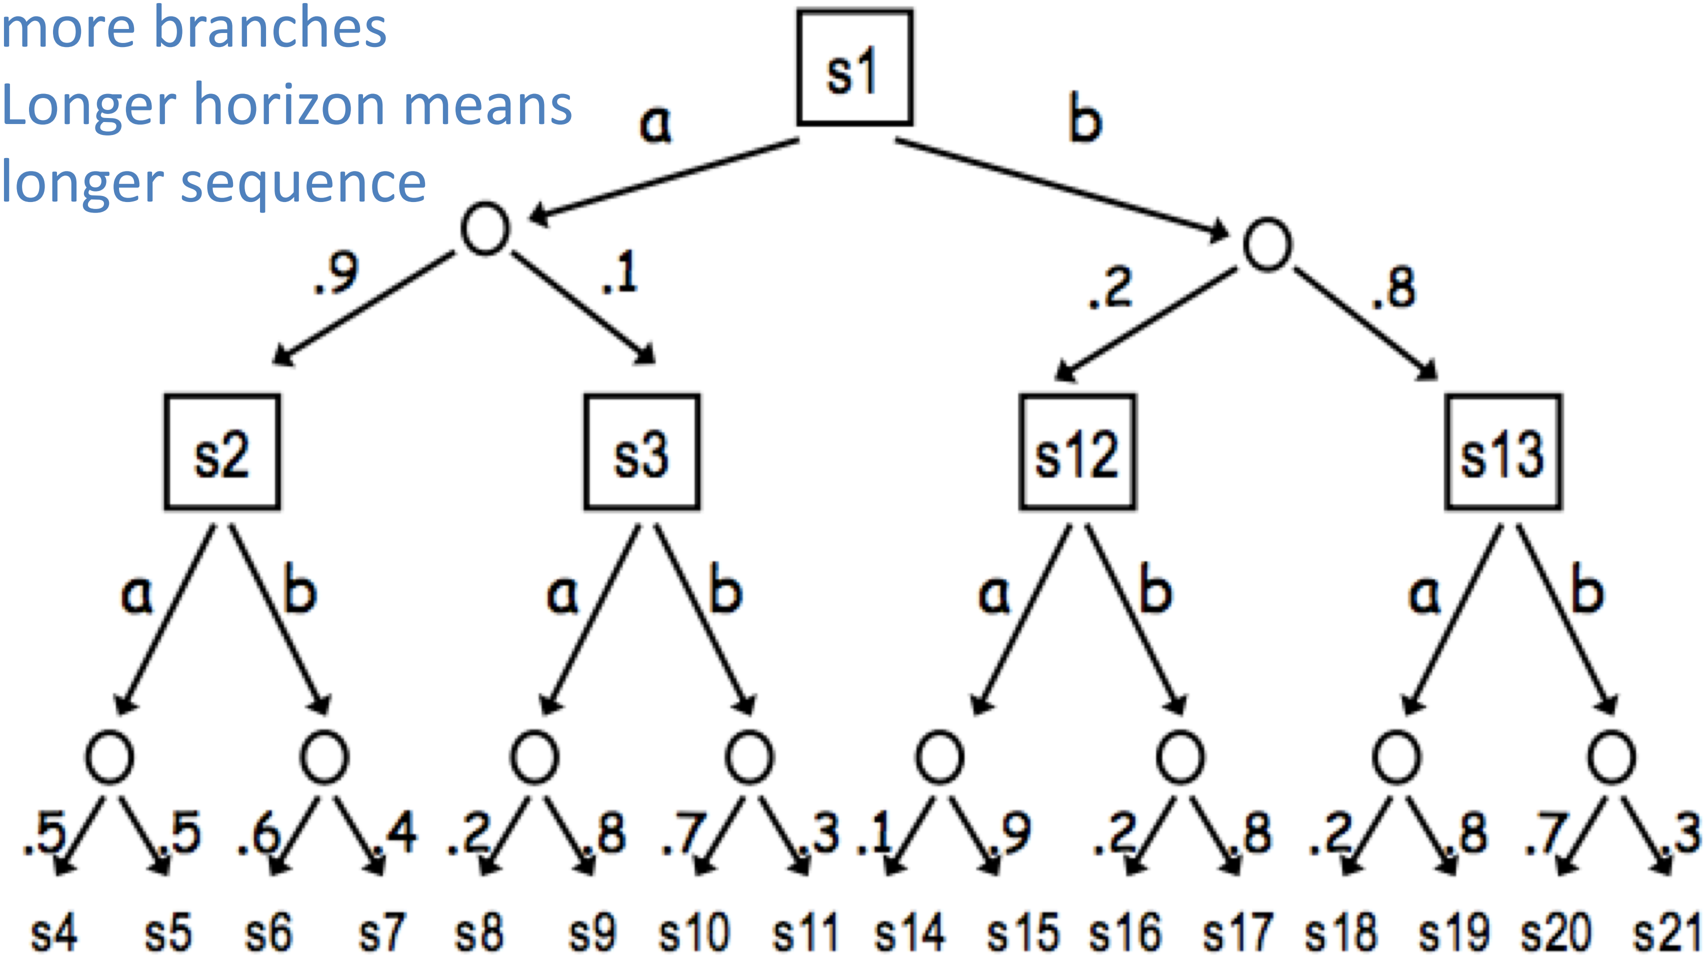

EU Exercises

- Complete handout of exercises
  - Individually or in groups

- Good practice for Quiz 2!

Key Ideas

- Main concepts
  - Utility theory enables us to express the strength of our preferences over outcomes
  - Maximum expected utility (MEU) principle
- Representation:
  - Decision making under certainty: Preference ranking of outcomes
  - Decision making under uncertainty:
    - Probability distribution of outcomes
    - Utility function to express preference of outcomes
- Computational issues:
  - Outcome space is too large
  - Decision space is too large