

Introduction to Computational Game Theory

CMPT 882

Simon Fraser University

Oliver Schulte

**Decision Making Under Uncertainty**

## Outline

- Choice Under Uncertainty: Formal Model
- Choice Principles
  - Expected Utility
  - Dominance
  - Maximin
  - Regret

## Choice and Uncertainty

- Most of decision theory is concerned with choice under uncertainty. We begin with a model of uncertainty.
- An agent faces a choice with uncertainty if there are several different ways the world might turn out to be that are relevant to his payoffs.
- Formally, we represent an agent's uncertainty with a set  $S$  of **possible states of the world**. That's it!
- **If** the agent knew which state of the world was the actual one, he would know what outcome would result from a given act. Thus we will think of **acts** as **functions** that assign an **outcome** to each possible state of the world.

## The Formal Model of Decision Under Uncertainty

We start with three sets.

1. A set  $A$  of acts that the agent is choosing among.
  2. A set  $O$  of possible outcomes.
  3. A set  $S$  of possible states of the world.
- An act  $a$  in  $A$  is a **function** from states of the world to outcomes. So for each state  $s$  in  $S$ , an act  $a$  assigns an outcome  $a(s)$  in  $O$ .

	<u>States of the World</u>			
<u>Acts</u>	$s_1$	$s_2$	....	$s_k$
$a_1$	$O_{11}$	$O_{12}$	....	$O_{1k}$
$a_2$	$O_{21}$	$O_{22}$	....	$O_{2k}$
...	....	....	....	....
$a_n$	$O_{n1}$	$O_{n2}$	....	$O_{nk}$

## Examples

<u>Acts</u>	<u>States of the World</u>	
	get caught	not get caught
Don't pay Skytrain fare	pay fine	pay nothing
Pay fare	pay fare	pay fare

*A* : pay fare, not pay fare

*O* : pay fine, pay fare, pay nothing

*S* : get caught, not get caught

<u>Acts</u>	<u>States of the World</u>	
	Accused is innocent	Accused is guilty
Secret trial	Likely wrong conviction	Likely true conviction
Public trial	Less likely wrong conviction	Less likely true conviction

*A* : secret trial, public trial

*O* : likely true conviction, likely false conviction, less likely wrong conviction, less likely true conviction

*S* : the accused is innocent/guilty

## Decision-making under Risk

- Suppose we know that Skytrain checks fares on 1/10 rides. Mathematically, we can represent this as a distribution that assigns a probability to each state of the world.

	1/10	9/10
	<u>States of the World</u>	
<u>Acts</u>	get caught	not get caught
Don't pay Skytrain fare	pay fine	pay nothing
Pay fare	pay fare	pay fare

- The fact that the probabilities of the state of the world are known makes this a decision problem under **risk**, as opposed to **uncertainty**. (This terminology is traditional in decision theory but not much used in computer science.)
- How does knowing the probabilities affect your decision?

## Decision-making under Risk: Expected Utility

Suppose that an agent can measure the value of each outcome. In the Skytrain example, the value may well be determined by the price of the fare and the fine.

<u>Acts</u>	1/10	9/10
	<u>States of the World</u>	
	get caught	not get caught
Don't pay Skytrain fare	-\$50	0
Pay fare	-\$3.50	-\$3.50

- Now we can assign to each option the average dollar cost as its **expected value**.

$$EU(\text{don't pay}) = -1/10 \times 50 - 9/10 \times 0$$

$$EU(\text{pay}) = -3.50$$

- Note the change in units: \$3.50 is a **dollar** amount, whereas  $EU(\text{pay})$  is an expected **utility** amount (sometimes called “utiles”). In this example we have assumed that dollar amount = utility.
- It actually suffices to assume only that utility is a positive linear function of money (e.g. utility = double dollar amount). You will often read in economics papers “we assume that money and utility are linearly related”.

## Notes on Expected Utility

- Expected utility is a score function for options:  $x \geq y$  if and only if  $EU(x) \geq EU(y)$ .
- Thus given a probability distribution over states of the world, expected utility extends a utility function on outcomes to a utility function over options.
- It may seem very natural to make decisions in this way. The famous von Neuman-Morgenstern **representation theorem** shows that an agent will choose as if they were maximizing an expected utility function if and only if the weak preference relation satisfies a set of axioms for choice under risk.
- The average payoff does not reflect the **variance** of the payoff. Given two options with the same expected monetary amount, most people prefer the one with less risk. A simple dramatic demonstration of this phenomenon is the **Ellsberg** paradox. Perhaps a rational choice theory should take variance into account as well.

## Subjective Expected Utility

- Von Neumann-Morgenstern expected utility theory assumes that the probabilities of outcomes are known. What if they are not?
- One answer is that in this case the agent should assign probabilities according to their best guess, and then maximize expected utility.
- The famous **representation theorem** due to Leonard Savage (1954) is a justification of this approach. It states a set of axioms like the ones we saw for weak preference such that an agent satisfies the axioms if and only if the agent acts if they are maximizing a personal probability assessment.
- Since the probabilities in question are not given objectively but constitute an agent's case, this approach is called person or subjective expected utility.

## Choice Under Uncertainty: Strict Dominance

Consider an agent facing a decision under uncertainty with actions  $A$ , outcomes  $O$  and states of the world  $S$ .

We assume that the agent has a rational preference relation  $\succ$  over the outcomes  $O$ .

### Definition

An act  $a$  **strictly dominates** another act  $a'$  iff for all states of the world  $s \in S$ , it is the case that  $a(s) \succ a'(s)$ .

To put it in a slogan, with a strictly dominant act “you can only win”.

### Example

Consider the Prisoner’s Dilemma from the point of view of the row player only.

<u>Acts</u>	<u>States of the World</u>	
	Cooperate	Defect
Cooperate	2	0
Defect	3	1

No matter what the column player does, Row prefers Defect.

## Choice Under Uncertainty: Weak Dominance

1. Consider an agent facing a decision under uncertainty with actions  $A$ , outcomes  $O$  and states of the world  $S$ .
2. We assume that the agent has a rational preference relation  $\circ$  over the outcomes  $O$ .

### Definition

An act  $a$  **weakly dominates** another act  $a'$  iff

1. for each state of the world  $s \in S$ ,  
it is the case that  $a(s) \circ a'(s)$

**and**

2. for some state of the world  $s \in S$ ,  
it is the case that  $a(s) \succ a'(s)$ .

To put it in a slogan, with a weakly dominant act “you can’t lose and you might win”.

## Examples

Somebody sends you some books to check out, “with no obligation to buy”.

<u>Acts</u>	<u>States of the World</u>	
	like book	don't like book
check out book	got good book 👍	no book, no money spent
return book	no book, no money spent	no book, no money spent

Checking out the book weakly dominates returning it immediately.

## Exercise

1. Jennifer is not satisfied with her grade on an assignment. She considers two options, accepting the grade, and taking up the instructor's office hour with a discussion of her grade. Model Jennifer's situation with a decision matrix with two acts and two states of the world. Assume that Jennifer prefers a better grade, and does not mind spending an hour discussing with her instructor. Show that going to discuss the grade weakly dominates not discussing it.

## The Maximin Principle

The maximin principle directs us to choose the act whose **worst-case** performance is the **best**.

### Examples

<u>Options</u>	<u>States of the World</u>	
	Heads	Tails
Bet 1	100	60
Bet 2	70	100

Maximin says to choose Bet 2.

<u>Options</u>	<u>States of the World</u>	
	catch two birds in the bush	don't catch two birds in the bush
take bird in the hand	1	1
chase two birds in the bush	2	0

maximin says to take the bird in the hand

In CS terms, maximin is a **worst-case** criterion.

## Minimax Regret

- The regret criterion can be intuitively looked at as a hindsight criterion: if you make a choice and then find out the true state of the world, how much better could you have done?
- Formally, suppose we have a utility function over the outcomes  $O$ . Recall that  $a(s)$  is the outcome of act  $a$  in state  $s$ .
- The regret for option  $a$  in state  $s$  is  $\max_{a'} u(a'(s)) - u(a(s))$ .
- The max regret is  $\max_s \max_{a'} u(a'(s)) - u(a(s))$ .
- The minimax regret criterion recommends keeping the worst-case regret as small as possible.  $\operatorname{argmin}_a \max_s \max_{a'} u(a'(s)) - u(a(s))$ .

### Example

<u>Options</u>	<u>States of the World</u>	
	Heads	Tails
Bet 1	100	60
Bet 2	70	100

Max regret of Bet 1 = 40

Max regret of Bet 2 = 30

## Exercise

	Blue	White	Green
X	0	50	100
Y	100	0	50
Z	50	100	0

1. Compute the max regrets of Options X,Y,Z.
2. What happens if we eliminate Option Z?
3. What happens if we eliminate Option X?

## Review and Closing Comments

- The formal model of uncertainty uses states of the world to represent the fact that acts may have several non-deterministic consequences.
- If the probabilities of the different states are known, the most common choice criterion is expected utility. However:
  - Expected utility is insensitive to the variance of the payoff distribution.
  - In many experiments, **behavioral decision theorists** have shown that people systematically violate expected utility theory. Time permitting, we can reproduce some of these experiments in class. I guarantee some of you will fall into irrational behaviour---even Savage did!
- If the probabilities of outcomes are unknown, most economists and some statisticians think that an agent should choose according to their **subjective probabilities**.
- Computer scientists and engineers often prefer to look at objective guarantees like maximin and regret which give worst-case bounds on performance.
- Dominance also does not assume any probabilities and is uncontroversial. However, in most single-agent problems there is no dominated act. For reasons we can discuss later, dominance plays a more useful role in game theory. One example is the Prisoner's Dilemma.