PGM 2004/05 Tirgul7
Foundations of Decision Theory

(mostly from Pearl)

Motivation

- You can choose between
  1) Get 3,000,000 $ with probability 1
  2) Get 4,000,000 $ with probability 0.8
- Though most people choose (1) the expectancy of profit is higher is (2)
- The reason is we assign Utilities for each prize.

Introduction

I'm planning a party and having a hard time to decide whether to stage it indoors or outdoors. In despair, I summarize what I can say in a table:

indoors	dry (0.7)	=	regret
indoors	wet (0.3)	=	relief
outdoors	dry (0.7)	=	perfect
outdoors	wet (0.3)	=	disaster

It is clear I prefer perfect to regret, but is it enough to make a decision? Probabilities quantify the *likely-hood* of events. We are looking for a measure which will quantify *desirability* and will be the basis of all decision making.

Utilities, consequences and MEU

Given some evidence, we are given a set of *actions* that we can take. Each action has several *consequences*. For each consequence c we assign a utility (or desirability) measure U(c). The expected utility of the action is then:

$$U(a) = \sum_{c} U(c)P(c \mid a, e)$$

Following the Bayesian approach, we want to maximize this expected utilitiy (MEU)

Lotteries

At each stage of the decision process, we are faced with a dilemma of which action to take. Each action has a payoff or a loss. Thus, much like gambling we are faced with a lottery situation. Lotteries are defined as pairs of consequences and probabilities L(C,P).

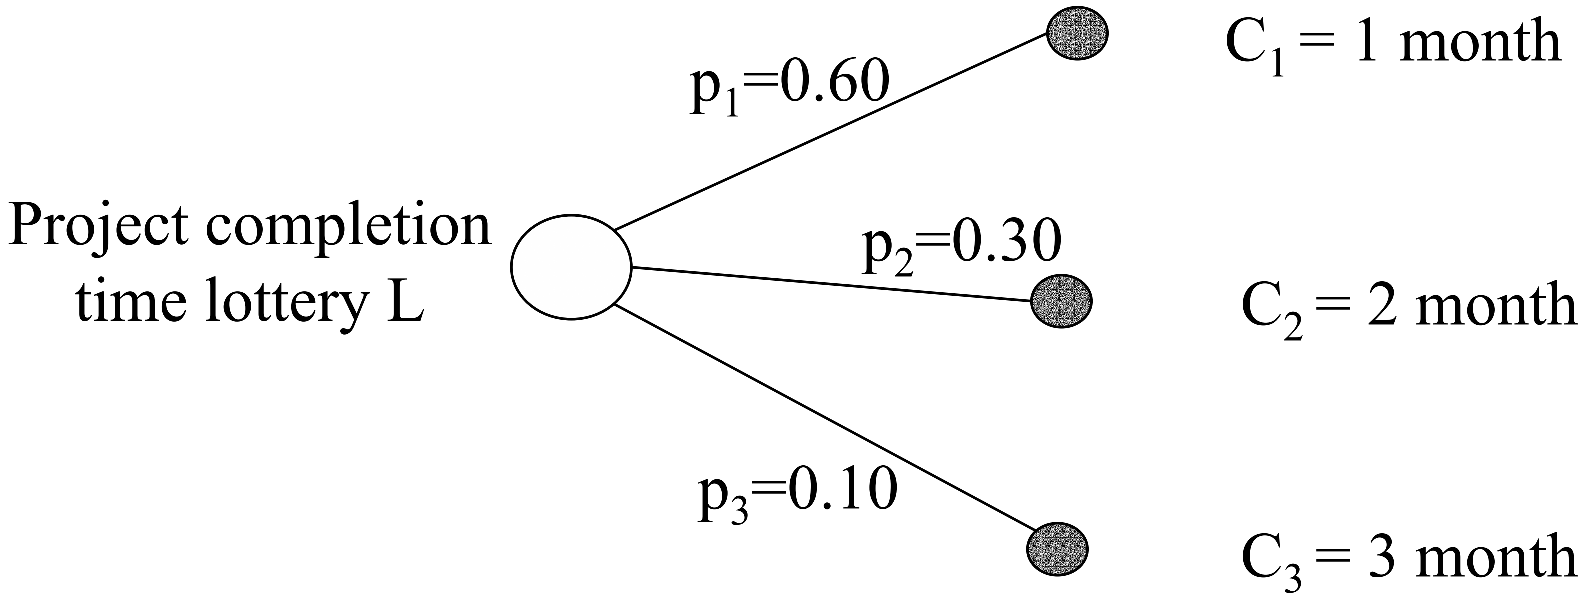

Lotteries (cont.)

When trying to assess the desirability of a lottery, we encounter several problems:

♦ What is a good calibration commodity? Is money stable enough? Does the commodity depend on time and circumstances?

♦ Once a commodity is found, we need to be able to deduce complex lotteries from simpler ones (or deal with infinitely many possibilities).

Calibrating a Lottery

Solution: calibrate a lottery against a standard best/worst lottery which is equivalent in desirability to ours (calibrate a lottery using it's inherent uncertainty).

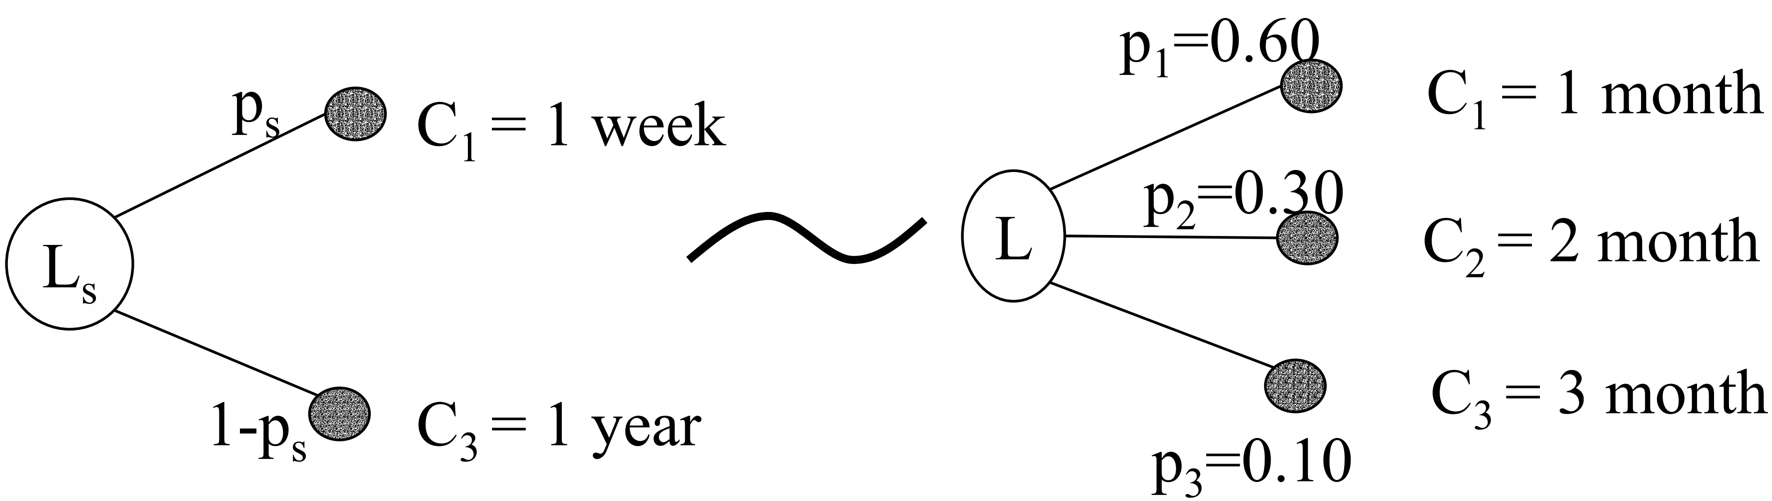

Axioms of Utility Theory

We are still left with the problem of how to define our preference over different consequences. We want a preference pattern that will allow us to prefer one lottery over another in a consistent way. The following axioms define some constraints on what kind of preference we are allowed to use as a basis for making decisions.

Axiom 1 - Orderability: A linear and transitive preference relation must exist between the prizes. In simpler words, an agent must know what he wants:

$$(c_1 \prec c_2) \vee (c_1 \succ c_2) \vee (c_1 \sim c_2)$$

$$(c_1 \prec c_2) \wedge (c_2 \prec c_3) \Rightarrow (c_1 \prec c_3)$$

Axioms of Utility Theory (cont.)

Axiom 2 - Continuity: If C_2 is in between C_1 and C_3 then we could compare C_2 to some lottery with only C_1 and C_3:

$$(c_1 \prec c_2 \prec c_3) \Rightarrow \exists p : c_2 \sim [c_1, p ; c_3, (1-p)]$$

Axiom 3 - Substitutability: If two things look equivalent, they will also be equivalent in a larger context:

$$L_1 \sim L_2 \quad iif \quad [L_1, p ; L_3, (1-p)] \sim [L_2, p ; L_3, (1-p)]$$

Axioms of Utility Theory (cont.)

Axiom 4 - Monotonicity: We prefer a better prize with higher probability:

$$p \succ p', C_1 \succ C_2 \Rightarrow [C_1, p ; C_2, (1-p)] \succ [C_1, p' ; C_2, (1-p')]$$

Axiom 5 - Decomposability: there is no fun in gambling (just the outcome matters):

$$L_2 = [C_1, q ; C_2, (1-q)] \Rightarrow$$

$$[L_1, p ; L_2, (1-p)] \sim [L_1, p ; C_1, (1-p)q ; C_2, (1-p)(1-q)]$$

Constructing Utilities

Theorem: If a preference pattern obeys axioms 1-5 then by specifying the utility measure of each consequence we can faithfully prefer one lottery to another using:

$$L_1 \succ L_2 \quad iif \quad U(L_1) \succ U(L_2)$$

$$where \quad U(L) = \sum_i p_i U(C_i)$$

This allows us to do just what we set out for: define desirability for elementary consequences and then use the MEU principle to evaluate complex lotteries.

Proof

“Surprisingly”, the essence of the proof rests on exactly the essence of the axioms: the ability to reduce complex lotteries to simple ones and that compare these simple two prize lotteries to each other. The outline of the proof follows:

Let’s us examine two lotteries L_1 and L_2 that we want to compare with a joint set of prizes $C_1 \ldots C_n$ (ordered from best prize to worst using axiom 1). Using substitution (axiom 3) we can write both lotteries in terms of all prizes (with probability 0 where needed).

$\Rightarrow$

Since any prize is between C_1 and C_n, using continuity (axiom 2):

$$C_i \sim \left[C_1, u_i \,;\, C_n, (1-u_i)\right]$$

and we can now substitute (axiom 3) and get:

Proof (cont.)

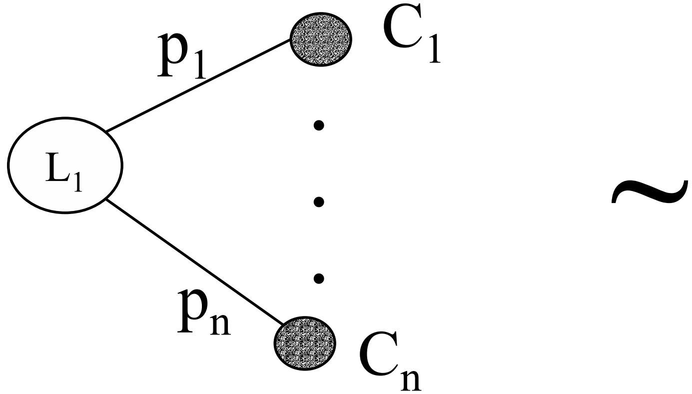

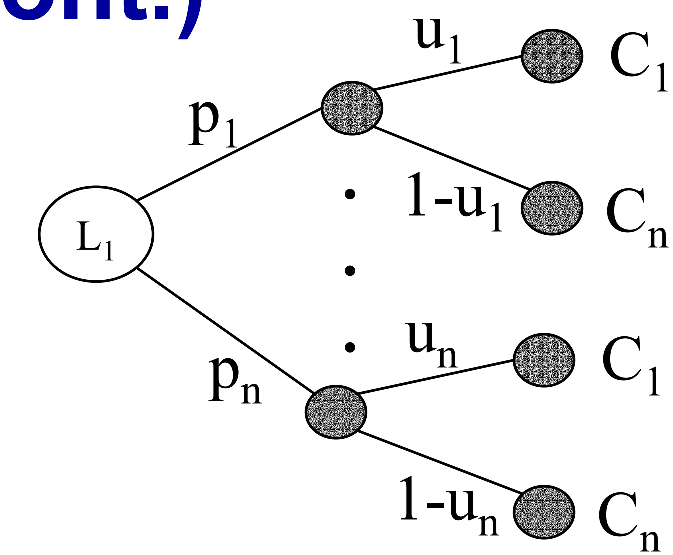

Now, applying decomposability, we can present our lotteries in the following equivalent way:

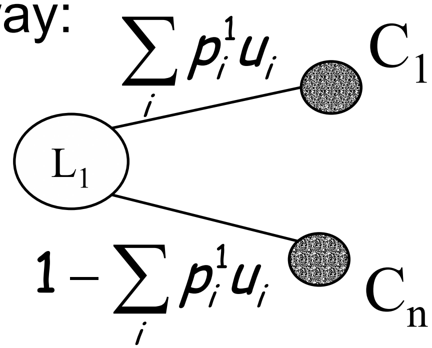

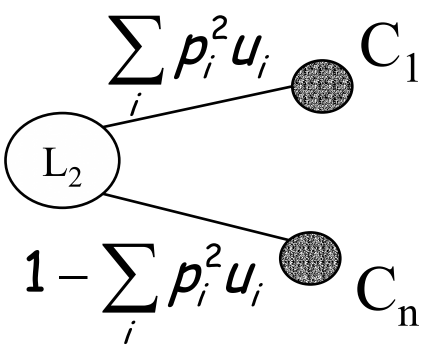

Using monotonicity (axiom 4), since $C_1 > C_n$, the preference between L_1 and L_2 is determined by the probability assigned to C_1. Thus:

$$L_1 \succ L_2 \quad iif \quad \sum_i p_i^1 u_i \succ \sum_i p_i^2 u_i$$

Real-life Utilities

Given all the axioms of decision theory, we are still left with the actual problem of assigning utilities to consequences. In fact, any utility we come up with is only as unique as a linear transformation of it.

Can we use monetary values as utilities?

Do you prefer $3,000,000 for sure or $4,000,000 with 80% chance?

This just means that our utilities is not money directly. If, we assign a utlity of 12 to the case of adding $3,000,000 to our bank account and 10 to that of adding $4,000,000 than our choice makes sense.

Utility of Money

St. Petersburg paradox (by Bernoulli): in the game you toss a coin until it comes up as heads. If the head appeared on the nth toss you get \$2n. How much are you willing to pay to participate?

The expected winnings are:

$$\sum_n P(H_n)U(H_n) = \sum_n \frac{1}{2^n} 2^n = 1 + 1 + 1 + \ldots = \infty$$

if we the utility of money is logarithmic then:

$$\sum_n P(H_n)U(H_n) = \sum_n \frac{1}{2^n} \log(2^n) = \sum_n \frac{n}{2^n} \to 2$$

Bounded Rationality

Finally, we hope that at least we are rational compared to some utility function (decision theory gives us freedom in choosing the function). Unfortunately we are all irrational:

Test 1: Would you participate in a lottery A where you get 15000 $ with 50% and loose 10000$ in 50%.

Test 2: Assess the value of all your properties M. Would you participate in a lottery B where you have M+15000 with 50% and M-10000 with 50%.

Kahneman & Tversky

- We often decide based on changes instead of on absolute conditions.
- We often decide based on assessments on average and not summation.
- This is based on the way our intuitive thinking works.

Decision Trees

When we want to make optimal decisions, we need to plan. We will ultimately merge Bayesian networks and decisions into influence diagrams. Let's start with a model which represent all possible scenarios:

A *decision tree* has two types of nodes: decision and chance. The leaves of the tree carry the utility associated with the scenario at that leaf.

We will find an optimal strategy by starting with the leaves and propagating towards the root:

- a chance node receives the expected utility
- a decision node is assigned the best utility (MEU) and the corresponding branch is marked.

Decision Trees: An example

The buyer of a used car can decide to carry out two tests. t_1 at the cost of \$50 and t_2 at the cost of \$20. There are two candidate cars c_1 costs \$1500 and it's market value is \$2000 but will cost \$700 to repair if the car is bad. c_2 costs \$1150 with a market value of \$1400 but a repair cost of only \$150.

The buyer needs to buy a car and has time for one test. From experience, c_1 has 70% chance of being good while c_2 has 80% chance of being good (as an ex, verify that using no test, it is better to buy c_1 with EMV of \$290).

Test t_1 checks car c_1. If the car is good there is a 90% chance that the test will confirm. If it is bad, the test will discover it in 65% of the cases.

Test t_2 checks car c_2. If will confirm good quality with 25% probability and discover flaws with 70% probability.

Decision Trees: An example (cont.)

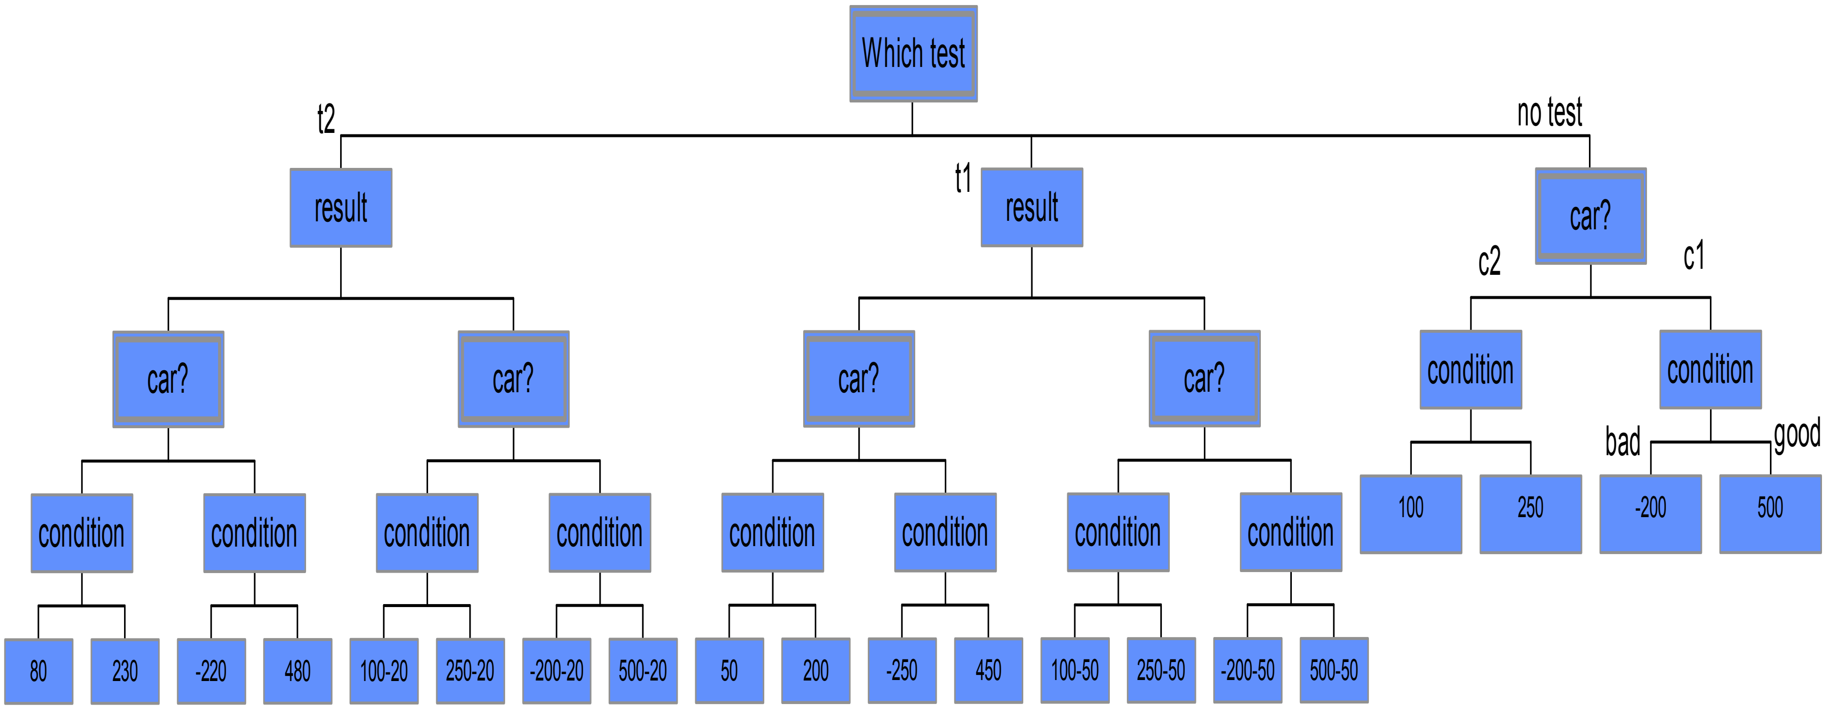

Plugging in the probabilities and propagating for optimal plan, we find that we should do t_1. If it passes we should bye c_1 and if it fails we should by c_2. Over all, we expect a utility of 303.77