SELF INTERESTED AGENT:

- Knows what states of the world likes more than others
- His decisions seek to bring the world in those states he likes more

WORLD = ENVIRONMENT — Its state changes depending on each agent action and even by chance

Our objective:

- formalize this concept of like/dislike states
- quantify it
- consider uncertainties

$\Bigg\}$ UTILITY THEORY

- Utility theory is grounded in the concept of PREFERENCES

$O = \{$ finite set of outcomes $\}$ (outcomes = world states)

for any $o_1, o_2 \in O$ (only this is needed: $\succeq$)

$o_1 \succeq o_2$ AGENT WEAKLY PREFERS o_1 TO o_2

$o_1 \sim o_2$ AGENT IS INDIFFERENT

$o_1 \succ o_2$ AGENT STRICTLY PREFERS o_1 TO o_2

this is how we denote preferences

This formalism allows us to give preferences over outcomes. It will be convenient to reason also uncertain outcomes which can be described by lotteries

$$[p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k] \qquad o_i \in O, \quad p_i \geq 0, \quad \sum_{i=1}^{k} p_i = 1$$

call $\mathcal{L}$ the set of all possible lotteries

extend $\succeq$ also to element of $\mathcal{L}$, so it now can be applied with the same interpretation to any pair of lotteries

(notice that (with a slight abuse of notation) $\mathcal{L} \supseteq O$)

UP TO NOW WE JUST HAVE A FORMALISM WITH WHICH WE CAN ENCODE AN AGENT'S PREFERENCE OVER THE OUTCOMES. WE CANNOT JUST GIVE "ANY" PREFERENCE RELATION. IN ORDER TO BE MEANINGFUL WE MUST REQUIRE SOME PROPERTY TO BE SATISFIED

(MEANINGFUL: THE RELATION WE SPECIFY IS SOUND WITH OUR IDEA OF PREFERENCES)

(REFLEXIVITY, NOT ANTISIMMETRY ALREADY IMPLIED)

- Completeness (C)

$$\forall o_1, o_2 \quad o_1 \succ o_2 \text{ or } o_2 \succ o_1 \text{ or } o_1 \sim o_2$$

the agent has a complete idea of his tastes.
He never answers "Dunno" when asked which one of two alternatives prefers

- Transitivity (T)

$$\forall o_1, o_2, o_3 \in O \text{ if } o_1 \succeq o_2 \wedge o_2 \succeq o_3 \text{ then } o_1 \succeq o_3$$

transitivity: we feel it is reasonable right? Why?
If it does not hold, then there exists $o_1, o_2, o_3 \in O$ such that

$$o_1 \succeq o_2, \quad o_2 \succeq o_3, \quad o_3 \succ o_1$$

$$o_3 \xrightarrow{m_{32}} o_2 \xrightarrow{m_{21}} o_1 \xrightarrow{m_{13}} o_3$$

m_{ij} = AMOUNT OF MONEY I AM WILLING TO PAY TO PASS FROM i TO j
(always ≥ 0)

then the agent is willing to pay
$(m_{32} + m_{21} + m_{13}) > 0$ to pass from o_3 to o_3 (NOT VERY WISE IN OUR FRAMEWORK!)

- Substitutability (S)

if $o_1 \sim o_2$ then

$$[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k]$$

for any $o_3, \dots, o_k$ s.t. $p + \sum_{i=3}^{k} p_i = 1$

· two lotteries that differ only in offering indifferent outcomes with the same probability

- Decomposability (D)

let p_i^e := PROB. WITH WHICH lottery e picks outcome i

if $\forall o_i \in O$ $p_i^{e_1} = p_i^{e_2}$, then $e_1 \sim e_2$

"no fun in gambling" → applies to nested lotteries: what matters is only the final induced probabilities

- Monotonicity (M)

if $o_1 \succ o_2$ and $p > q$ then $[p : o_1, (1-p) : o_2] \succ [q : o_1, (1-q) : o_2]$

THESE ABOVE ENCODE OUR IDEA OF "CONSISTENT" PREFERENCES WHEN HAVING UNCERTAINTY ON THE TABLE

- Continuity (C)

if $o_1 \succ o_2$ and $o_2 \succ o_3$ then $\exists p \in [0,1]$ such that

$$o_2 \sim [p : o_1, (1-p) : o_3]$$

* von Neumann and Morgenstern

If $C \wedge T \wedge S \wedge D \wedge M \wedge \mathcal{C}$ then there exists

$$u : \mathcal{L} \longmapsto [0,1] \quad \text{such that}$$

- $u(o_1) \geq u(o_2)$ iff $o_1 \succeq o_2$
- $u([p_1 : o_1 \ldots p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$

IT'S A NICE RESULT! IF OUR PREFERENCE RELATION IS WELL FORMED THEN WE CAN QUANTIFY THE PREFERENCE DEGREE WITH AN UTILITY FUNCTION u THAT

- ASSIGNS TO EACH OUTCOME A VALUE
- IT TAKES THE FORM OF EXPECTANCY WHEN UNCERTAINTY IS PRESENT

→ [Our self-interest agent wants to maximize the expected value of a function u called the "utility function"]

LET'S PUT THIS FRAMEWORK TO WORK WITH A SIMPLE EXAMPLE

Single agent decision problem

- We have an agent that can undertake two actions: <u>a</u> and <u>b</u>
- the agent is initially in state S_0
- actions are prone to errors: we are not sure where they would lead

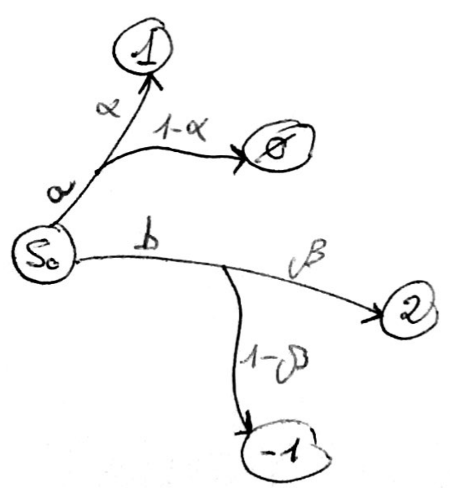

state s: $u(s)$

$u(a) = u([\alpha : ①, (1-\alpha) : ⓪]) = \alpha$

$u(b) = 2\beta - (1-\beta) = 3\beta - 1$

WHAT TO CHOOSE?

$$\sigma^* : \begin{cases} a & \text{if } \alpha \geq 3\beta - 1 \\ b & \text{if } \alpha \leq 3\beta - 1 \end{cases}$$

AN INTERESTING DIGRESSION: is u = MONEY?

LET'S MAKE THIS QUESTION UNDER A UTILITY = MONEY SCOPE:

BOX 2: 0,5: 2 MILLION € CASH (20 €)
0,5: FREE VEGAN DINNER

BOX 1: 500,000 € CASH (9 €)

REASONING WITH EXPECTED UTILITY WE WOULD DETERMINE THAT BOX 1 IS WORTH 1 MILLION WHICH IS STRICTLY BETTER THAN BOX 2

Utility and money are not linearly related. To understand this we should talk about risk aversion.

LAST IMPORTANT PROPERTY: (P)

each positive affine transformation of u is still an utility function

$$u() \longrightarrow a\,u() + b \qquad a, b \text{ constant}, \; a > 0$$

Agent $i \longrightarrow u_i()$