University of California – Berkeley
Department of Economics
ECON 201A Economic Theory
Choice Theory
Fall 2023

Choice under uncertainty
Part II: subjective probability – Savage (1954) style theory
(Kreps Ch. 5 and Rubinstein Ch. 7)

Oct 5 and 12, 2023

vNM style theory

In vNM theory, the externally imposed objects are

1. a set X of (uncertain) prospects/prizes.
2. a set P of some sort of probability measures on X

(as part of the objective description of the prospect, probabilities are assigned to various prizes and/or sets of prizes so they are 'only' risky).

Inside such a setting, vNM provide conditions on $\succ$ on P that give an expected-utility representation:

$$p \succ q \Leftrightarrow \sum_{x \in X} p(x)u(x) > \sum_{x \in X} q(x)u(x). \qquad (*)$$

– The summation in $(*)$ makes sense if either X is finite or if P is the set of simple/discrete probability measures on some (arbitrary) X.

– For more general types of P, well-defined integrals need to replace the summations in $(*)$.

! But however we ‘finesse’ the definition the idea of the vNM representation is the same:

1. u is an index of how good each prize x is
2. p is indexed by the expected value of that index.

Savage (1954)

In vNM, the objects of choice are probability distributions over prizes. But in many contexts, the odds of various outcomes are not at all clear...

As a result, what a $\mathcal{DM}$ chooses depends critically on what s/he subjectively assesses as the odds of the outcomes.

h win $1,000 if Liverpool wins the Premier League ($0 otherwise).

h' win $1,000 if Man UTD wins the Premier League ($0 otherwise).

? But what if also add win $1,000 if four fair coin flips all come up heads ($0 otherwise)?!

- The vNM model would regard the two gambles as lotteries with objectively specified probabilities.

- Any (reasonable) $\mathcal{DM}$ would bet the same way – pick whichever has the better chance of getting the $1,000 prize.

! Not all $\mathcal{DM}$s will bet the same way on this $\Longrightarrow$ a model of choice under uncertainty that develops within itself the subjective probabilities.

The classic formal treatment of choice where there is subjective uncertainty is that of Savage (1954).

The basics of the Savage formulation:

– a set of X of prizes/consequences

– a set S of states of the world (or nature)

Each $s \in S$ is a compilation of all characteristics/factors about which the $\mathcal{DM}$ is uncertain and which are relevant to the consequences that will result from her/his choice.

The set S is an exhaustive list of mutually exclusive states – only one $s \in S$ will be the realized state.

There are three possible outcomes when Liverpool plays Man UTD: Liverpool wins (s_1), Man UTD wins (s_2), or they draw (s_3).

In the Premier League (or a horse race), each $s \in S$ will describe the order of finish of all teams, and S is the set of all such orders of finish.

From X and S we construct the choice space, denoted by H, as the set of all functions from S to X (formally written $H = X^S$).

Being less fanciful, $s \in S$ are states and $x \in X$ are prizes so $h \in H$ are state-contingent claims, which is the set of objects of choice (Savage calls them acts).

Savage seeks (and most economists employ) a representation of the following form: There exist a function $\pi : S \longrightarrow [0,1]$ such that $\sum_{s \in S} \pi(s) = 1$ (a probability distributions) and a function $u : X \longrightarrow \mathbb{R}$ such that

$$h \succ h' \Leftrightarrow \sum_{s \in S} \pi(s) u(h(s)) > \sum_{s \in S} \pi(s) u(h'(s)). \qquad (**)$$

This is just like vNM expected utility, except that the probabilities of the various prizes are obtained by a two-step process:

(i) probabilities are subjectively assessed for the various states of nature $s \in S$, and

(ii) the probability of getting a prize $x \in X$ if $h \in H$ is chosen is the sum of the probabilities of those states $s \in S$ such that $h(s) = x$.

! In Savage, S is infinite so the summations in $(**)$ are appropriately defined integrals but we take S and X to be finite so that no mathematical difficulties get in the way of the conceptual content...

Three things to note about Savage's representation $(**)$:

1. Both tastes (the utility function u), and beliefs (the probability measure π) are subjective.

2. The utility of a particular prize and the probability measure are independent of the action taken – we do not write $\pi(s; h)$ or anything like that...

3. The utility of a prize does not depend on the state in which it is received (nor does the prize received in a state affect the probability of that state).

(1) and (2) are obvious... To understand points (3), consider an example that is hard to fit into the Savage setup $(**)$:

h win a ticket to a Liverpool game if its standing in the Premier League improves over the next month (nothing otherwise).

h' win a ticket to a Liverpool game if its standing in the Premier League worsens over the next month (nothing otherwise).

We can create a model where the states $(s \in S)$ are all possible Premier League standings and the prizes $(x \in X)$ are "ticket" and "no ticket."

But winning the ticket when Liverpool is a favourite is (much) better than winning it when a miracle is required for Liverpool to win the title. (This should be very clear. If not ask Michael...)

– One 'cure' is to give up on $(**)$ and to seek, instead, a (weaker) "state-dependent" u-representation $u : X \times S \rightarrow \mathbb{R}$ of the form:

$$h \succ h' \Leftrightarrow \sum_{s \in S} \pi(s) u(h(s), s) > \sum_{s \in S} \pi(s) u(h'(s), s).$$

In the decision theory literature, this is called a "state-dependent expected utility."

– We can also go one level further... Given π and u, define $v : X \times S \to \mathbb{R}$ as

$$v(x, s) = \pi(s)u(x, s).$$

Then the representation just given becomes

$$h \succ h' \Leftrightarrow \sum_{s \in S} v(h(s), s) > \sum_{s \in S} v(h'(s), s) \qquad (***)$$

and it is called the "additively-separable-across-states" representation.

Savage's sure thing principle

Suppose we are comparing two acts h and g which are identical on subset T of S, that is

$$h(s) = g(s) \text{ for all } s \in T \subset S$$

then whether the $\mathcal{DM}$ prefers h or g depends only on how h and g compare on states $s \in S \setminus T$.

Formally, if $h \succ g$ and if $h'$and g' are are two other acts such that

$$h(s) = h'(s) \text{ and } g(s) = g'(s) \text{ for all } s \in S \setminus T$$

and

$$h'(s) = g'(s) \text{ for all } s \in T$$

then $h' \succ g'$.

$$
h := \begin{array}{ll} \overset{s_1}{\nearrow} & x \\ \xrightarrow{s_2} & y \\ \underset{s_3}{\searrow} & z \end{array} \quad \succ \quad g := \begin{array}{ll} \overset{s_1}{\nearrow} & x' \\ \xrightarrow{s_2} & y' \\ \underset{s_3}{\searrow} & z \end{array}
$$

$$
\Updownarrow
$$

$$
h' := \begin{array}{ll} \overset{s_1}{\nearrow} & x \\ \xrightarrow{s_2} & y \\ \underset{s_3}{\searrow} & z' \end{array} \quad \succ \quad g' := \begin{array}{ll} \overset{s_1}{\nearrow} & x' \\ \xrightarrow{s_2} & y' \\ \underset{s_3}{\searrow} & z' \end{array}
$$

In words, the ranking of h and g does not depend on the specific way that they agree on T – that they agree is enough!

Savage's sure thing principle is clearly implied by representation $(**)$ and even by representation $(***)$:

$$\begin{aligned} U(h) &= \sum_{s \in S} v(h(s), s) \\ &= \sum_{s \in T} v(h(s), s) + \sum_{s \in S \setminus T} v(h(s), s) \end{aligned}$$

and write a similar expression for $U(g)$. As $h(s) = g(s)$ for all $s \in S \setminus T$, a comparison of $U(h)$ and $U(h)$ depends on how

$$\sum_{s \in T} v(h(s), s) \text{ compares with } \sum_{s \in T} v(g(s), s).$$

Anscombe and Aumann (1963) (A-A)

1. Obtaining the representation $(**)$ is quite a hard task so we will continue with a different, easier, formalization of A-A.

2. What A-A have done is to enlarge the domain of choice in the Savage formulation in the hope that this will make matters easier...

3. A-A avoid the complexities that Savage encounters by enriching the space over which preferences are defined:

   - $i.$ take as given a prize space X and a (finite) state space S
   - $ii.$ let P be the set of all (simple) probability distributions on X
   - $iii.$ redefine H so that it is the set of all functions from S to P

Formalizing, let P be the set of probability distributions on X and take H to be the set of all functions from S to probability distributions over prizes, $H = P^S$.

A-A seek a representation of the following form: There exist a function $\pi : S \to [0,1]$ such that $\sum_{s \in S} \pi(s) = 1$ and a function $u : X \to \mathbb{R}$ such that

$$\begin{array}{c} h \succ h' \\ \Updownarrow \\ \sum_{s \in S} \pi(s) \left[\sum_{x \in X} h(x \,|\, s) u(x) \right] > \sum_{s \in S} \pi(s) \left[\sum_{x \in X} h'(x \,|\, s) u(x) \right] \end{array}$$

where $h(x \,|\, s)$ for the (objective) probability of winning prize $x \in X$ under h, conditional on the state being s.

The "story" of A-A:

– There is a set of roulette wheel gambles or lotteries – randomizing devices that allows to construct any (objective) probability distribution $p \in P$.

– For each possible outcome of the horse race $h \in H$, a roulette wheel lottery is won by the holder of the betting ticket (a degenerate roulette wheel in Savage).

– The $\mathcal{DM}$ has preferences given by $\succ$ defined on this fancy H where $h(\cdot\,|s)$ is the element of P that is the prize under h in state s.

Suppose h and $h'$are two horse race lotteries (in this fancy/expanded sense).

For any $\alpha \in [0.1]$, define a new horse race lottery, $\alpha h + (1 - \alpha)h'$, as a new horse race lottery that gives as prize the roulette wheel lottery

$$\alpha h(\cdot\,|s) + (1 - \alpha)h'(\cdot\,|s) \text{ for all } s \in S.$$

Then, (A1)-(A3) above hold — (A2) and (A3) do not depend on the fact that p, q, and r are probability distributions (only that convex combinations of the objects of choice can be taken).

(A1)-(A3) are necessary and sufficient for $\succ$ to have a u-representation of the A-A form:

$$h \succ h'$$
$$\Updownarrow$$
$$\sum_{s \in S} \sum_{x \in X} h(x \,|\, s) u_s(x) > \sum_{s \in S} \sum_{x \in X} h'(x \,|\, s) u_s(x)$$

where for each state $s \in S$ there is state-dependent utility function u_s.

That is, to evaluate h, first, for each state s, compute the expected-utility of the roulette gamble $h(x \,|\, s)$ using the utility function for state u_s (that need not bear any relationship to any $u_{s'}$).

We need to tie together the various u_s...

(A4) For all $p, q \in P$ and state s^*, construct horse race lotteries $\cdot h$ and h' as follows:

$$h(s) = \begin{cases} r & \text{if} \quad s \neq s^* \\ p & \text{if} \quad s = s^* \end{cases} \quad \text{and} \quad h'(s) = \begin{cases} r & \text{if} \quad s \neq s^* \\ q & \text{if} \quad s = s^* \end{cases}$$

for an arbitrarily $r \in P$. Then $h \succ h' \Leftrightarrow p \succ q$.

The difference in utilities between h and h' is 'just' the difference between

$$\sum_{x \in X} p(x) u_{s^*}(x) \text{ and } \sum_{x \in X} q(x) u_{s^*}(x).$$

Hence, the $\mathcal{DM}$'s $\succ$ are independent of the state in which a pair of roulette lotteries are compared:

$$h := \begin{array}{c} \overset{s_1}{\nearrow} \\ \xrightarrow{s_2} \\ \underset{s_3}{\searrow} \end{array} \begin{array}{c} p \\ r \\ r \end{array} \quad \succ \quad h' := \begin{array}{c} \overset{s_1}{\nearrow} \\ \xrightarrow{s_2} \\ \underset{s_3}{\searrow} \end{array} \begin{array}{c} q \\ r \\ r \end{array}$$

$$\Updownarrow$$

$$g := \begin{array}{c} \overset{s_1}{\nearrow} \\ \xrightarrow{s_2} \\ \underset{s_3}{\searrow} \end{array} \begin{array}{c} r \\ r \\ p \end{array} \quad \succ \quad g' := \begin{array}{c} \overset{s_1}{\nearrow} \\ \xrightarrow{s_2} \\ \underset{s_3}{\searrow} \end{array} \begin{array}{c} r \\ r \\ p \end{array}$$

(A1)-(A4) are necessary and sufficient for $\succ$ to have a u-representation of the following form:

$$h \succ h'$$
$$\Updownarrow$$
$$\sum_{s \in S} \pi(s) \left[\sum_{x \in X} h(x\,|\,s) u(x) \right] > \sum_{s \in S} \pi(s) \left[\sum_{x \in X} h'(x\,|\,s) u(x) \right]$$

In words, a horse race gamble is as good as its (subjective) expected utility. Except for the objective lottery part of these gambles, we have Savage's representation $(**)$.

Michael loves Liverpool and hates Man UTD so he has three states in mind:

s_1	Liverpool wins the Premier League
s_2	Man UTD wins the Premier League
s_3	some other team wins...

Obviously...

$$h := \begin{array}{lll} \overset{s_1}{\nearrow} & \xrightarrow{1} & 100 \\ \xrightarrow{s_2} & \xrightarrow{1} & 0 \\ \underset{s_3}{\searrow} & \xrightarrow{1} & 0 \end{array} \quad \succ \quad h' := \begin{array}{lll} \overset{s_1}{\nearrow} & \xrightarrow{1} & 0 \\ \xrightarrow{s_2} & \xrightarrow{1} & 100 \\ \underset{s_3}{\searrow} & \xrightarrow{1} & 0 \end{array}$$

If Michael's $\succ$ conform to the A-A axioms (and if he prefers more money to less), what can we tell about how he assesses the probability the Liverpool will win the Premier League if $h \succ h'' \succ h'$ where

$h'' :=$

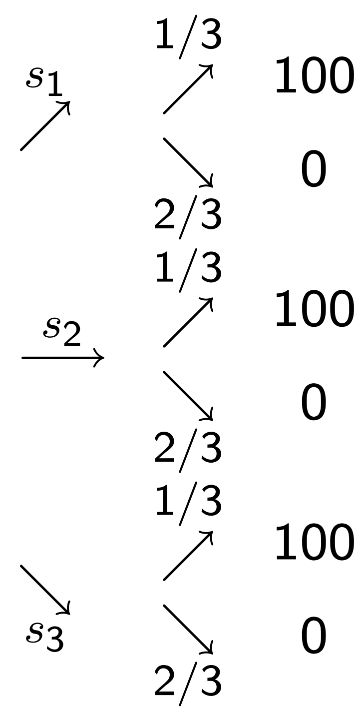


Women
WE ARE LIVERPOOL.
YOU'LL NEVER WALK ALONE.