

ECON0106: Microeconomics

8a. Uncertainty^{*}

Duarte Gonçalves[†]

University College London

1. Overview

How likely is it that it is going to rain today? Should I take the umbrella with me or not? Different websites show different probabilities, but do we know the true, objective probability that it rains today? (Is there such a thing?) While this is an uninteresting example if you are in London — everyone knows the answers: (i) it is going to rain a.s., (ii) always take the umbrella — it points to a fundamental issue that we need to address: agents may not know the probability that an event realises.

More: from a frequentist perspective, it makes little sense to talk about the *objective* probabilities of singular, unrepeatable events. Instead, the modern conceptualisation of probability, based on measure theory, enables us to talk about *subjective* probability.

In this lecture, we will study models of choice under *uncertainty*. This can be seen as a decision-theoretic foundation to the very notions of subjective probability. We will introduce the two main models of subjective expected utility, due to [Savage \(1954\)](#) and [Anscombe and Aumann \(1963\)](#), and we will focus on the latter. We then turn to understanding a choice-based representation of belief updating, and we conclude with a discussion on uncertainty aversion.

2. Subjective Expected Utility

2.1. Anscombe–Aumann’s Framework

The main ingredients of [Anscombe and Aumann’s \(1963\)](#) framework are the following:

Ω : set of states of the world;

X : set of consequences or outcomes;

$f : \Omega \rightarrow \Delta(X)$: an act;

$\mathcal{F} := \Delta(X)^\Omega$: set of acts;

^{*}Last updated: 4 November 2025.

[†] Department of Economics, University College London; duarte.goncalves@ucl.ac.uk. Please do not share these notes with people outside of this class.

$\succsim \subseteq \mathcal{F}^2$: preference relation.

This framework differs from [Savage's \(1954\)](#) (which we will introduce later) by considering two sources of uncertainty: (i) subjective uncertainty (horse race) — which state $\omega \in \Omega$ will be realised — and (ii) objective uncertainty/risk (roulette wheel) — which consequence $x \in X$ will be realised in a lottery. That is, each can be seen as a compound lottery: a potentially different objective lottery is triggered by each (uncertain) state of the world.

Our goal will be to characterise properties of $\succsim$ that enable us to recover not only a Bernoulli utility function on consequences, $u : X \rightarrow \mathbb{R}$, but also a probability measure $\mu \in \Delta(\Omega)$ such that

$$f \succsim g \quad \text{if and only if} \quad \mathbb{E}_\mu[\mathbb{E}_{f(\omega)}[u]] \geq \mathbb{E}_\mu[\mathbb{E}_{g(\omega)}[u]],$$

where, for a given state ω , $f(\omega)$ is an objective probability distribution in $\Delta(X)$ and so $\mathbb{E}_{f(\omega)}[u]$ is our von Neumann – Morgenstern expected utility; and as μ represents the decision-maker's belief over the states, we then take expectations of vNM expected utility with respect to the beliefs μ to obtain the agent's subjective expected utility. Further, we want the representation to be sharp, in that beliefs should be uniquely pinned down from preferences.

For simplicity, we will assume that Ω and X are finite.

Definition 1. For acts $f, g \in \mathcal{F}$, we call their **mixture** given $\alpha \in [0, 1]$ the act $\alpha f + (1 - \alpha)g$ such that $(\alpha f + (1 - \alpha)g)(\omega) = \alpha f(\omega) + (1 - \alpha)g(\omega)$.

For lottery $p \in \Delta(X)$ we will abuse notation and call $\tilde{p} \in \mathcal{F}$ a **constant act** such that $\tilde{p}(\omega) = p \in \Delta(X)$ for every $\omega \in \Omega$.

That is, a constant act p is the act that delivers lottery p in every state of the world.

We define **continuity** of $\succsim$ as we did before, i.e. for all sequences $\{f_n, g_n\}_n$ such that $f_n \succsim g_n$ for every n and $f_n \rightarrow f, g_n \rightarrow g$, we have that $f \succsim g$. Analogously to what we did with preferences on objective probability distributions, we now say that $\succsim$ satisfies **independence** if for any acts $f, g, h \in \mathcal{F}$ and any $\alpha \in (0, 1]$, $f \succsim g \iff \alpha f + (1 - \alpha)h \succsim \alpha g + (1 - \alpha)h$.

We then get an intermediate result, a state-dependent expected utility representation.

Theorem 1. (*State-Dependent Expected Utility*) A preference relation $\succsim$ on $\mathcal{F}$ satisfies continuity and independence if and only if there exists a utility function $u : X \times \Omega \rightarrow \mathbb{R}$ and a uniform probability measure $\mu \in \Delta(\Omega)$ such that, for any $f, g \in \mathcal{F}$, $f \succsim g$ if and only if $\mathbb{E}_\mu[\mathbb{E}_{f(\omega)}[u]] \geq \mathbb{E}_\mu[\mathbb{E}_{g(\omega)}[u]]$.

Is this it? Not quite. Note that $u : X \times \Omega \rightarrow \mathbb{R}$, i.e. we have $u(x, \omega)$, not $u(x)$. This may be a virtue or a vice, depending on how you perceive it, but the main issue is that it is capturing both preferences over consequences and beliefs about states. This is why we have a uniform belief/prior μ , void of any empirical content.

Proof. The if part, as usual, is a matter of verification, and we then focus on the only if part.

Before proving **Theorem 1**, a stepping stone toward our target representation, we introduce some notation and a lemma:

$E \subseteq \Omega$: an event

fEg : a ‘conditional act,’ where for acts f, g and event E , $fEg \in \mathcal{F}$ is such that $(fEg)(\omega) = f(\omega)$ if $\omega \in E$ and $(fEg)(\omega) = g(\omega)$ if otherwise;

Null event E : an event such that for any $f, g, h \in \mathcal{F}$ for which $f \succ g$, $fEh \sim gEh$;

$\tilde{x}$: a constant act, $\tilde{x}(\omega) = x, \forall \omega \in \Omega$.

Clearly null events are disregarded by the decision-maker and this, as we will see, has an important connection with their beliefs. Now, the lemma:

Lemma 1. *Let $V : \mathcal{F} \rightarrow \mathbb{R}$ be affine and continuous. Then, for every $\omega \in \Omega$, there is an affine and continuous function $V_\omega : \Delta(X) \rightarrow \mathbb{R}$ such that $V(f) = \sum_\omega V_\omega(f(\omega))$, for all $f \in \mathcal{F}$.*

Proof. Fix $f^* \in \mathcal{F}$. For any $f \in \mathcal{F}$, we can write

$$\begin{aligned} \frac{1}{|\Omega|}f &= \frac{1}{|\Omega|}f^* + \frac{1}{|\Omega|} \sum_\omega (f\{\omega\}f^* - f^*) \\ \iff \frac{1}{|\Omega|}f + \left(1 - \frac{1}{|\Omega|}\right)f^* &= \frac{1}{|\Omega|} \sum_\omega (f\{\omega\}f^*) \\ \iff V\left(\frac{1}{|\Omega|}f + \left(1 - \frac{1}{|\Omega|}\right)f^*\right) &= V\left(\frac{1}{|\Omega|} \sum_\omega (f\{\omega\}f^*)\right) \\ \iff \frac{1}{|\Omega|}V(f) + \left(1 - \frac{1}{|\Omega|}\right)V(f^*) &= \frac{1}{|\Omega|} \sum_\omega V(f\{\omega\}f^*) \\ \iff V(f) &= \sum_\omega [V(f\{\omega\}f^*) - (|\Omega| - 1)V(f^*)/|\Omega|], \end{aligned}$$

where we used the fact that V is affine. Then, we can define $V_\omega : \Delta(X) \rightarrow \mathbb{R}$ such that $V_\omega(f(\omega)) := V(f\{\omega\}f^*) - (|\Omega| - 1)V(f^*)/|\Omega|$. As V is continuous and affine, so is V_ω . $\square$

Now, for every $f \in \mathcal{F}$, let $\rho_f \in \Delta(X \times \Omega)$ as $\rho_f(\{x\} \times \{\omega\}) := \frac{1}{|\Omega|}f(\omega)(x)$, for all $(x, \omega) \in X \times \Omega$. ρ_f is then a joint distribution over $X \times \Omega$ with a uniform marginal over Ω .

Let $R := \{\rho_f \mid f \in \mathcal{F}\} \subseteq \Delta(X \times \Omega)$ and define $\triangleright \subseteq R^2$ as $\rho_f \triangleright \rho_g$ if and only if $f \succsim g$. As $\succsim$ is a continuous preference relation, so is $\triangleright$. Moreover, as $\succsim$ satisfies independence, then, for any $\alpha \in (0, 1]$ and every $h \in \mathcal{F}$, $\rho_f \triangleright \rho_g \iff \rho_{\alpha f + (1-\alpha)h} \triangleright \rho_{\alpha g + (1-\alpha)h}$. As

$$\begin{aligned} \rho_{\alpha f + (1-\alpha)h}(\{x\} \times \{\omega\}) &= (\alpha f + (1-\alpha)h)(\omega)(x) = \alpha f(\omega)(x) + (1-\alpha)h(\omega)(x) \\ &= \alpha \rho_f(\{x\} \times \{\omega\}) + (1-\alpha)\rho_h(\{x\} \times \{\omega\}), \end{aligned}$$

then we immediately have that $\triangleright$ (defined on R) satisfies independence.

Thus, by (an easy adaptation of) the von Neumann – Morgenstern expected utility representation theorem, we have that there is a unique function (up to positive affine transformations) $v : X \times \Omega \rightarrow \mathbb{R}$ such that $\rho_f \succeq \rho_g$ if and only if $\sum_{x,\omega} \rho_f(\{x\} \times \{\omega\})v(x,\omega) \geq \sum_{x,\omega} \rho_g(\{x\} \times \{\omega\})v(x,\omega)$.

Define $V : \mathcal{F} \rightarrow \mathbb{R}$ such that $V(f) := \sum_{x,\omega} \rho_f(\{x\} \times \{\omega\})v(x,\omega)$ and note that V is affine and continuous. Additionally, V represents $\succsim$ given that

$$f \succsim g \iff \rho_f \succeq \rho_g \iff V(f) \geq V(g).$$

By **Lemma 1**, there is a $V_\omega : \Delta(X) \rightarrow \mathbb{R}$ which is affine and continuous such that $V(f) = \sum_\omega V_\omega(f(\omega))$. Finally, define $u : X \times \Omega$ as $u(x,\omega) := V_\omega(\delta_x)|\Omega|$. Then, as for each ω , V_ω is affine, $V_\omega(p) = \sum_x p(x) \frac{1}{|\Omega|} u(x,\omega)$. $\square$

In order to establish a state-independent subjective expected utility representation we will need $\succsim$ to satisfy a property which we call **separability**: $\forall p, q \in \Delta(X)$, all $h \in \mathcal{F}$, and all $\omega, \omega' \in \Omega$ such that $\{\omega\}$ and $\{\omega'\}$ are non-null events, $\tilde{p}\{\omega\}h \succsim \tilde{q}\{\omega\}h$ if and only if $\tilde{p}\{\omega'\}h \succsim \tilde{q}\{\omega'\}h$.

We then have what we wanted, a state-independent representation:

Theorem 2. (*Anscombe–Aumann Subjective Expected Utility*) Let $\succsim$ be a preference relation on $\mathcal{F}$. $\succsim$ satisfies continuity, independence, and separability if and only if there exists $\mu \in \Delta(\Omega)$ and $u : X \rightarrow \mathbb{R}$ such that

$$f \succsim g \iff \mathbb{E}_\mu[\mathbb{E}_f[u]] \geq \mathbb{E}_\mu[\mathbb{E}_g[u]],$$

where, for all $f \in \mathcal{F}$, $\mathbb{E}_\mu[\mathbb{E}_f[u]] := \sum_{\omega \in \Omega} \mu(\omega) \sum_{x \in X} f(\omega)(x)u(x)$. Moreover, u is unique up to positive affine transformations and, if $\exists f, g \in \mathcal{F}$ such that $f \succ g$, μ is unique.

Proof. Again, we focus on the only if part. From **Theorem 1**, there is $u : X \times \Omega$ such that $f \succsim g$ if and only if $\sum_{\omega,x} f(\omega)(x)u(x,\omega) \geq \sum_{\omega,x} g(\omega)(x)u(x,\omega)$.

Let $U : \Delta(X) \times \Omega \rightarrow \mathbb{R}$ be defined as $U(p,\omega) := \sum_{x \in X} p(x)u(x,\omega)$ for all $\omega \in \Omega$, $p \in \Delta(X)$. Take any $p, q \in \Delta(X)$ and non-null $\{\omega\}$, $\omega \in \Omega$ such that $U(p,\omega) \geq U(q,\omega)$. From separability we have that for any non-null $\{\omega'\}$, $\omega' \in \Omega$, and any $h \in \mathcal{F}$,

$$\begin{aligned} U(p,\omega) \geq U(q,\omega) &\iff U(p,\omega) + \sum_{\omega'' \in \Omega \setminus \{\omega\}} U(h(\omega''), \omega'') \geq U(q,\omega) + \sum_{\omega'' \in \Omega \setminus \{\omega\}} U(h(\omega''), \omega'') \\ &\iff V(\tilde{p}\{\omega\}h) \geq V(\tilde{q}\{\omega\}h) \\ &\iff \tilde{p}\{\omega\}h \succsim \tilde{q}\{\omega\}h \iff \tilde{p}\{\omega'\}h \succsim \tilde{q}\{\omega'\}h \\ &\iff V(\tilde{p}\{\omega'\}h) \geq V(\tilde{q}\{\omega'\}h) \\ &\iff U(p,\omega') + \sum_{\omega'' \in \Omega \setminus \{\omega'\}} U(h(\omega''), \omega'') \geq U(q,\omega') + \sum_{\omega'' \in \Omega \setminus \{\omega'\}} U(h(\omega''), \omega'') \\ &\iff U(p,\omega') \geq U(q,\omega'). \end{aligned}$$

Let $\Omega^* \subseteq \Omega$ be the set of ω such that $\{\omega\}$ is non-null and fix $\omega^* \in \Omega^*$. As u is unique up to positive affine transformations, for all $\omega \in \Omega^*$, $\exists \alpha_\omega > 0, \beta_\omega \in \mathbb{R}$ such that $u(x, \omega) = \alpha_\omega u(x, \omega^*) + \beta_\omega$, for all $x \in X$.¹ Let $u : X \rightarrow \mathbb{R}$ be defined as $u(x) := u(x, \omega^*)$. Then, we have that

$$\begin{aligned} f \succsim g &\iff \sum_{\omega \in \Omega} U(f(\omega), \omega) = \sum_{\omega \in \Omega} \sum_{x \in X} f(\omega)(x) u(x, \omega) \geq \sum_{\omega \in \Omega} \sum_{x \in X} g(\omega)(x) u(x, \omega) \\ &\iff \sum_{\omega \in \Omega^*} \sum_{x \in X} f(\omega)(x) \alpha_\omega u(x) + \beta_\omega \geq \sum_{\omega \in \Omega^*} \sum_{x \in X} g(\omega)(x) \alpha_\omega u(x) + \beta_\omega \\ &\iff \sum_{\omega \in \Omega^*} \frac{\alpha_\omega}{\sum_{\omega' \in \Omega^*} \alpha_{\omega'}} \sum_{x \in X} f(\omega)(x) u(x) \geq \sum_{\omega \in \Omega^*} \frac{\alpha_\omega}{\sum_{\omega' \in \Omega^*} \alpha_{\omega'}} \sum_{x \in X} g(\omega)(x) u(x). \end{aligned}$$

Now define $\mu : \Omega \rightarrow [0, 1]$ such that $\mu(\omega) := \mathbf{1}_{\omega \in \Omega^*} \frac{\alpha_\omega}{\sum_{\omega' \in \Omega^*} \alpha_{\omega'}}$.²

Uniqueness of $u : X \rightarrow \mathbb{R}$ up to positive affine transformations follows by the fact that $u : X \times \Omega \rightarrow \mathbb{R}$ is too.

Now take any $f > g$; we will show that μ is uniquely defined. If $f > g$ then it must be the case that u is non-constant (otherwise $\mathbb{E}_\mu[\mathbb{E}_{f(\omega)}[u]] = \mathbb{E}_\mu[\mathbb{E}_{g(\omega)}[u]]$). Then, there exists $z, y \in X$ such that $u(z) > u(y)$. By taking the acts $\tilde{\delta}_z\{\omega'\}\tilde{\delta}_y$ for each $\omega' \in \Omega^*$, and the constant act $h = \mu(\omega')\tilde{\delta}_z + (1 - \mu(\omega'))\tilde{\delta}_y$,³ we have that $h \sim \tilde{\delta}_z\{\omega'\}\tilde{\delta}_y$, and therefore

$$\begin{aligned} \sum_{\omega \in \Omega} \mu(\omega) \sum_{x \in X} h(\omega)(x) u(x) &= \sum_{x \in X} h(\omega')(x) u(x) \\ &= \mu(\omega') u(z) + (1 - \mu(\omega')) u(y) = \sum_{\omega \in \Omega} \mu(\omega) \sum_{x \in X} (\tilde{\delta}_z\{\omega'\}\tilde{\delta}_y)(\omega)(x) u(x), \end{aligned}$$

which shows uniqueness of μ (supposing existence of some other $\nu \in \Delta(\Omega)$ such that $\nu \neq \mu$ will give a contradiction, as h is a constant act and therefore the ν would not play a role in determining its expected utility). $\square$

Exercise 1. We say that $\succsim$ on $\mathcal{F}$ satisfies **monotonicity** if $\forall f, g \in \mathcal{F}$, whenever $\widetilde{f(\omega)} \succsim \widetilde{g(\omega)}$ for all $\omega \in \Omega$ such that $\{\omega\}$ is non-null, then $f \succsim g$. Show that if $\succsim$ is preference relation that satisfies independence and continuity, then separability and monotonicity are equivalent.

While [Anscombe and Aumann's \(1963\)](#) is a more tractable approach to modeling subjective uncertainty, it has been pointed out as conceptually less appealing than [Savage's \(1954\)](#) exactly because of state-independence. We did want state-independence as this was the only way to derive from preferences the decision-maker's beliefs.

State-independence imposes some constraints on how we should think about states, acts, and consequences. A usual criticism to [Anscombe and Aumann's \(1963\)](#) framework goes as fol-

¹Recall that for any $\omega \in \Omega$ such that $\{\omega\}$ is non-null, $V_\omega : \Delta(X) \rightarrow \mathbb{R}$ is an expected utility representation of preferences over $\Delta(X)$. Separability gives us that V_ω all represent the same preferences on $\Delta(X)$ and we know they are vNM representations; hence, they must be affine transformations of each other.

²Technically speaking, μ is a probability measure and is therefore defined over a σ -algebra Σ on Ω . But as Ω is finite, we can redefine μ on 2^Ω (a σ -algebra) by doing $\mu(E) := \cup_{\omega \in E} \mu(\omega)$ for every $E \subseteq \Omega$, with the understanding that $\mu(\{\emptyset\}) = 0$.

³This is the constant act $h := \tilde{p}$, such that $p = \mu(\omega')\tilde{\delta}_z + (1 - \mu(\omega'))\tilde{\delta}_y$, where $p \in \Delta(X)$.

lows: Suppose we want to express that a decision-maker prefers to carry an umbrella when it is raining and not carrying it when it is not raining. It would be tempting to model this as $\Omega = \{\text{rain, no rain}\}$ and $X = \{\text{carrying umbrella, not carrying umbrella}\}$, but then we are ruling out that carrying umbrella $>$ not carrying umbrella when $\omega = \text{rain}$ and carrying umbrella $<$ not carrying umbrella when $\omega = \text{no rain}$.

The issue with this example is that it is using Anscombe–Aumann and thinking Savage where we can have $u(\text{umbrella}(\text{rain})) > u(\text{umbrella}(\text{no rain}))$. We can see this example as how *not* to think about states, acts, and consequences in [Anscombe and Aumann’s \(1963\)](#) framework. Carrying an umbrella or not is an *act*; we want to have it when it rains because we don’t want to get wet; and we don’t want to have it when it doesn’t rain because we don’t want to carry it around. So, instead, you need to rethink the set of consequences: $X = \{\text{having to carry an umbrella, not having to carry an umbrella}\} \times \{\text{getting wet, not getting wet}\}$. This solves the apparent problem:

$\mathcal{F}$	Ω	
	rain	no rain
taking an umbrella	not wet but carrying the umbrella	not wet, carrying the umbrella
not taking an umbrella	wet, not carrying umbrella	not wet, not carrying the umbrella

In other words: once you specify the set of consequences appropriately, you come to realise that the reason for why the decision-maker wants the umbrella when it rains is to not get wet. An umbrella as an act then induces two state-dependent degenerate lotteries: getting wet if it rains for sure and not getting wet but having to carry the bloody thing for sure if it doesn’t.

Exercise 2. Suppose that you want to learn what a person thinks is more likely, E or E^C (not E) and can’t ask them directly. You decide to put what you learned about Anscombe–Aumann to work.

Assume that the person strictly prefers more money over less. Show how to define a decision problem where the person has to choose between two acts f, g and, from their choices alone, you can infer whether they think E is more likely than E^C or not.

2.2. Savage’s Framework

We now introduce [Savage’s \(1954\)](#) framework. We discuss the main postulates of the model but we won’t provide any proofs — you should see [Kreps \(1988, ch. 8\)](#) if you’re interested.

The basic ingredients are as follows:

Ω : set of states of the world;

X : set of consequences or outcomes;

$f : \Omega \rightarrow X$: an act;

$\mathcal{F} := X^\Omega$: set of acts;

$\succsim \subseteq \mathcal{F}^2$: preference relation.

As you can see from the above, the crucial difference is that acts map directly to consequences and not to state-dependent lotteries on the set of consequences. To better appreciate the difference, let us put both representations next to each other:

- **Savage:** $\int_{\Omega} u(f(\omega))d\mu(\omega)$;
- **Anscombe–Aumann:** $\int_{\Omega} \int_X f(\omega)(x)u(x)d\mu(\omega)$.

The representation theorem goes as follows:

Theorem 3. (*Savage, 1954*) $\succsim$ satisfies P1-P7 if and only if there exist

- (i) a unique nonatomic and finitely additive probability measure μ on Ω , for which $\mu(E) = 0$ if and only if E is a null event; and
- (ii) a bounded function $u : X \rightarrow \mathbb{R}$, unique up to positive affine transformations

such that for every $f, g \in \mathcal{F}$,

$$f \succsim g \quad \text{if and only if} \quad \mathbb{E}_{\mu}[u \circ f] := \int_{\Omega} u(f(\omega))d\mu(\omega) \geq \int_{\Omega} u(g(\omega))d\mu(\omega) = \mathbb{E}_{\mu}[u \circ g].$$

where P1-P7 are defined in [Appendix A](#).

3. Bayesian Updating

Let's continue with the example What happens when our decision-maker, who is deciding whether or not to take an umbrella, looks out the window? They learn whether it is or it is not raining at that moment and this tells them something about whether it is going to rain later on.

Then, they should update their beliefs (by Bayes rule) and act accordingly: for any event $E \subseteq \Omega$, the posterior belief given $A \subseteq \Omega$ is given by

$$\mu \mid A(E) = \frac{\mu(A \cap E)}{\mu(A)}$$

for any A such that $\mu(A) > 0$. The issue is that their beliefs were *deduced* from their preferences, their behaviour. Is it the case that these beliefs we deduced are updated according to Bayes rule? Yes — under some additional restrictions.

Let's continue with the Anscombe–Aumann setup⁴ and enrich it in the following way. Events

⁴You can also do this within the Savage framework and the requirements on preferences will be similar; see [Ghirardato \(2002\)](#).

$A \subseteq \Omega$ are information, i.e., the decision-maker is learning that the state ω lies in A . Our new primitive is no longer just a single preference relation $\succsim$ but a collection of preferences $\{\succsim_A\}_{A \subseteq \Omega}$.⁵ Each $\succsim_A$ is a preference relation on acts $\mathcal{F}$ and the idea is that $\succsim_A$ describes their behaviour upon obtaining information A . We write $\succsim \equiv \succsim_\Omega$ to denote the decision-maker's preference in absence of any information.

We will start by assuming that for each $A \subseteq \Omega$, $\succsim_A$ is a preference relation on the set of acts satisfying independence, continuity, and monotonicity (do **exercise 1!**).

We need to impose some consistency requirements.

Definition 2. We say that $\{\succsim_A\}_{A \subseteq \Omega}$ satisfy

- (i) **constant-act consistency** if preferences over constant acts are consistent: for all lotteries $p, q \in \Delta(X)$ and events $A, B \subseteq \Omega$, $\tilde{p} \succsim_A \tilde{q}$ if and only if $\tilde{p} \succsim_B \tilde{q}$;
- (ii) **dynamic consistency** if, for all non-null events (with respect to $\succsim$) $A \subseteq \Omega$ and all acts $f, g \in \mathcal{F}$, $fAg \succsim_\Omega g$ if and only if $f \succsim_A g$;
- (iii) **consequentialism** if, for event $A \subseteq \Omega$, two acts $f, g \in \mathcal{F}$ deliver the same lottery $f(\omega) = g(\omega)$ for every $\omega \in A$, then $f \sim_A g$.

Some remarks on why these seem sensible assumptions to make. First, constant-act consistency says that if you take two acts whose (distribution over) consequences is completely orthogonal to the state, then whatever you learn should not change how you compare them. Together with independence, continuity, and monotonicity, we get that $\exists \alpha_A > 0, \beta_A \in \mathbb{R}$ such that $u_A = \alpha_A u + \beta_A$. In other words, constant-act consistency is doing the work in tying together the utility functions over consequences.

If constant-act consistency deal makes preferences over consequences remain the same, then the other two assumptions are working on rendering the inferred beliefs consistent across the different information sets. Dynamic consistency is saying that if you take two acts that differ only when A occurs, then if you learn that A occurs, you compare them in the same way. This makes the decision-maker evaluate the 'relevant acts' under A the same way as they did before they had the information. We said that consistency of preferences over consequences was taken care of; then what is this assumption doing? It is making the decision-maker keep the same relative beliefs across any two non-null events.

Finally, consequentialism is taking care of rendering events B such that $B \cap A = \emptyset$ null events with respect to $\succsim_A$. So, in other words, one can interpret this assumption as 'the decision-maker believes the information received.'

The representation theorem follows:

⁵Please read this as requiring $A \neq \emptyset$.

Theorem 4. Let $\{\succsim_A\}_{A \subseteq \Omega}$ be a collection of preference relations on $\mathcal{F}$ and assume that there are $f, g \in \mathcal{F}$ such that $f \succ_{\Omega} g$.

$\{\succsim_A\}_{A \subseteq \Omega}$ satisfy constant-act consistency, dynamic consistency, and consequentialism and $\succsim_A$ satisfies continuity, independence, and monotonicity for every $A \subseteq \Omega$ if and only if there exist $u : X \rightarrow \mathbb{R}$ and a collection of probability measures $\{\mu_A\}_{A \subseteq \Omega}$, $\mu_A \in \Delta(\Omega) \forall A \subseteq \Omega$ ($A \neq \emptyset$), such that

- (i) $\forall f, g \in \mathcal{F}$, $f \succsim_A g$ if and only if $\mathbb{E}_{\mu_A}[\mathbb{E}_{f(\omega)}[u]] \geq \mathbb{E}_{\mu_A}[\mathbb{E}_{g(\omega)}[u]]$, and
- (ii) for all non-null events with respect to $\succsim_{\Omega}$, $A \subseteq \Omega$, $\mu_A(B) = \frac{\mu_{\Omega}(A \cap B)}{\mu_{\Omega}(A)}$ for all $B \subseteq \Omega$.

Moreover, u is unique up to positive affine transformations and μ_{Ω} is unique.

Proof. By **Theorem 2** (AA SEU theorem), for all non-null event $A \subseteq \Omega$, there is $u_A : X \rightarrow \mathbb{R}$ and a prior $\mu_A \in \Delta(\Omega)$ such that $f \succsim_A g$ if and only if $\mathbb{E}_{\mu_A}[\mathbb{E}_{f(\omega)}[u_A]] \geq \mathbb{E}_{\mu_A}[\mathbb{E}_{g(\omega)}[u_A]]$.

We break the proof into steps that you should prove yourself:

Exercise 3.

Step 1. Show that constant-act consistency implies that for every non-null event $A \subseteq \Omega$, $\exists \alpha_A > 0, \beta_A \in \mathbb{R}$ such that $u_A = \alpha_A u_{\Omega} + \beta_A$.

Step 2. Show that there are $x, y \in X$ such that $\tilde{\delta}_x \succ_A \tilde{\delta}_y$ for all non-null events $A \subseteq \Omega$.

Step 3. Use the previous step to prove that for all non-null events $A \subseteq \Omega$, $\mu_{\Omega}(A) > 0$.

Step 4. Prove that, for any acts $f, g, h \in \mathcal{F}$ and any non-null event $A \subseteq \Omega$, $f \succsim_A g \iff fAh \succsim_{\Omega} gAh$.

Step 5. Conclude by showing that for any acts $f, g, h \in \mathcal{F}$ and any non-null event $A \subseteq \Omega$, $f \succsim_A g \iff fAh \succsim_{\Omega} gAh$ implies that $\mu_A(\omega) = \mu_{\Omega}(\omega)/\mu_{\Omega}(A)$.

Step 6. Argue the uniqueness claims based on the previous steps and on **Theorem 2**.

Step 7. Verify the ‘if’ part by showing that the representation implies the assumptions on $\{\succsim_A\}_{A \subseteq \Omega}$.

□

Theorem 4 imposes conditions on choices that agents must satisfy if they are *behaving like* Bayesian subjective expected utility maximisers. There are two points to make here. On the one hand, note that we cannot really observe people’s beliefs; we can only infer them from their behaviour. So, if someone’s behaviour is not consistent with these axioms, it does not mean that they are not updating beliefs in a Bayesian fashion (which is not falsifiable per se). On the other hand, it is interesting to point out that **Savage’s (1954)** book is called *The*

Foundations of Statistics. That is to say, there is a view that Bayesian statistics is *founded upon* statistical decision theory.

You don't have to take a stance on this: you may assume agents have and update beliefs in accordance to Bayes' rule and/or adhere (or not) to Bayesian statistics despite acknowledging that their behaviour may not be in line with this model. There are other interesting models that provide rationales for Bayesian updating that are not subjective expected utility (e.g. [Cripps, 2019](#)), and models that use Bayesian updating and are not rationalisable with subjective expected-utility maximising behaviour (e.g. [Alós-Ferrer and Mihm, 2021](#)).

Bayesian updating and subjective expected utility remain the main framework in economic models: they are very appealing principles and their vices and virtues when it comes to behavioural implications are well-known. All-in-all, a model meant to be an approximation to reality that captures all the relevant aspects of the situation being represented. And only if people think something truly crucial is missing, that either Bayesian updating or subjective expected utility are making you unable to account for, do they seek alternatives.

This being said, let's see one well-known issue with subjective expected utility and some approaches that have been suggest to dealing with it.

4. Uncertainty Aversion

4.1. Ellsberg Paradox

This is an experiment by [Ellsberg \(1961\)](#):⁶

A box contains 60 balls: 20 are black and the rest are either red or green. Which would you prefer:

- A £20 if a black ball is drawn;
- B £20 if a red ball is drawn; or
- C £20 if a green ball is drawn.

Most people choose A.

Now consider choosing among the following:

- a £20 if a black or a green ball is drawn;
- b £20 if a black or a red ball is drawn; or
- c £20 if a red or a green ball is drawn.

⁶While I am not aware of such an interesting story surrounding this paradox as the one involving Allais and Savage, the Ellsberg paradox does feature Daniel Ellsberg, known world-wide because of the story of the pentagon papers: https://en.wikipedia.org/wiki/Daniel_Ellsberg.

	Ω	
	ω_1	ω_2
f	p	0
g	0	p

And, to this question, most people choose c .

You should convince yourself that this is incompatible with SEU.

Exercise 4. *Show that a decision-maker strictly preferring A over B and C , and c over a and b is incompatible with subjective expected utility.*

4.2. A Set of Probability Measures: Maxmin Expected Utility

There have been multiple ways to approach this issue. One has been to relax the independence (always to blame) in a way such that μ no longer has to add up to one (i.e. is not a probability measure). This approach, followed in [Schmeidler \(1989\)](#), but many find it unappealing, as we know well the properties of probability measures and we want to keep them for modeling convenient.

Further, in the end what may be behind this is the intuitive idea that individuals may want to hedge their bets to avoid uncertainty. Suppose that the decision-maker is indifferent between f and g . While this implies that $\mu(\omega_1) = \mu(\omega_2) = 1/2$, it should not need to be the case that the decision-maker is also indifferent between f, g , and $h := \tilde{q}$, where $q := 1/2p + 1/2\delta_0$. To capture this idea we say that $\succsim$ on the set of (Anscombe–Aumann) acts $\mathcal{F}$ is **GS uncertainty averse** (neutral/seeking) if for any acts f, g such that $f \sim g$, we have that $\frac{1}{2}f + \frac{1}{2}g \succsim f$ ($\sim/\succ$).

This led to the following relaxation of independence (always taking the blame): we say that a preference relation $\succsim$ on $\mathcal{F}$ satisfies **C-independence** (C for certainty) if, for any two acts $f, g \in \mathcal{F}$ independence holds with respect to mixtures with a constant act $\tilde{p} \in \mathcal{F}$, with $p \in \Delta(X)$, that is, $\forall \alpha \in (0, 1]$,

$$f \succsim g \iff \alpha f + (1 - \alpha)\tilde{p} \succsim \alpha g + (1 - \alpha)\tilde{p}.$$

Clearly, C-independence is implied but does not imply independence. The reasoning is that we want to hedge, but hedging is only valuable when it can eliminate uncertainty, which is not the case if it uses a constant act (think about the example above).

This next exercise makes the point that independence is exactly C-independence plus with neutral attitudes toward uncertainty.

Exercise 5. *Show that if $\succsim$ is a preference relation on $\mathcal{F}$ satisfying continuity and monotonicity, then independence is equivalent to C-independence and uncertainty neutrality.*

These assumptions enable a representation of a different kind: **maxmin (subjective) ex-**

pected utility (Gilboa and Schmeidler, 1989).

Theorem 5. (Gilboa–Schmeidler Maxmin Expected Utility) Let $\succsim$ be a preference relation on $\mathcal{F}$. Then $\succsim$ satisfies continuity, monotonicity, C-independence, and GS uncertainty aversion if and only if there is a utility function $u : X \rightarrow \mathbb{R}$ and a convex and compact set of probability measures $\mathcal{M} \subseteq \Delta(\Omega)$ such that

$$f \succsim g \quad \text{if and only if} \quad \min_{\mu \in \mathcal{M}} \mathbb{E}_\mu[\mathbb{E}_f[u]] \geq \min_{\mu \in \mathcal{M}} \mathbb{E}_\mu[\mathbb{E}_g[u]].$$

Note that GS uncertainty aversion is a statement regarding an agent’s attitudes toward uncertainty. The decision-maker has not a prior, but a set of probability measures $\mathcal{M} \subseteq \Delta(\Omega)$ that is *endogenous* to the representation. Different preferences can induce representations with different sets of probability measures. The maxmin model implicitly assumes that agents are extremely uncertainty averse, behaving as if expecting the worst to happen among all the probability distributions over the state space that they entertain.

4.3. Beliefs over Unknown Probabilities: Smooth Uncertainty Aversion

We may then wonder whether we can’t just get something like standard risk aversion but for uncertainty, where agents don’t think about the worst possible outcome but instead trade-off uncertainty for better consequences. And the answer is yes, we can. [Klibanoff et al. \(2005\)](#) — and, more recently, [Denti and Pomatto \(2021\)](#) — prove conditions on preferences that allow us to have a utility representation that looks like this:

$$U(f) := \int_{\Delta(\Omega)} \phi \left(\int_{\Omega} u(f(\omega)) d\mu(\omega) \right) d\pi(\mu)$$

where $f : \Omega \rightarrow X$ is a Savage act,

$u : X \rightarrow \mathbb{R}$ a von Neumann – Morgenstern utility,

$\phi : \mathbb{R} \rightarrow \mathbb{R}$ a strictly increasing and continuous function, $\mu \in \Delta(\Omega)$ is a probability measure on the state space, and

$\pi \in \Delta(\Delta(\Omega))$ is the decision-maker’s *prior*, capturing their uncertainty about how the state is actually distributed.

So, we go one level up: the agent doesn’t only have beliefs about what is the true state; they are uncertain about how it is distributed to start with! How can this be reasonable? You can of this as a situation where you know that there is randomness — e.g. you know that machine impression produces biased coins (the probability of flipping heads and tails isn’t exactly 1/2) — but you face uncertainty about this bias, i.e. don’t really know how this bias is distributed.

It is called smooth uncertainty aversion because if the curvature of u captures risk attitudes,

that of ϕ captures uncertainty attitudes in an analogous way.⁷ That is, we separate preferences over risk and over uncertainty. Klibanoff et al. (2005) also show that we get equivalent characterisations of uncertainty aversion⁸ in the same way as we saw equivalent characterisations of risk aversion. For instance, the decision-maker is uncertainty averse if and only if ϕ is concave, and we can make statements comparing uncertainty attitudes of different agents just like we did with risk aversion. Further, we recover the maxmin model as a limit case, in which the decision-maker is extremely uncertainty-averse.

Exercise 6. We often observe behaviour that seems to be driven by regret avoidance. The purpose of this exercise is to show that, in the presence of regret, preferences need not be transitive.

Let $\Omega := \{\omega_1, \omega_2, \dots, \omega_n\}$ be the states of the world and $\mu \in \Delta(\Omega)$ — i.e., $\mu(\omega)$ denotes the belief that state ω occurs. As before, we have $X \subseteq \mathbb{R}$ be a set of consequences (e.g., money). We define acts as $f : \Omega \rightarrow X$, with $\mathcal{F} := \{f : \Omega \rightarrow X\}$ denoting the set of acts. In other words, μ specifies the probability with which each state ω occurs and $f(\omega)$ specifies the consequences. Fix some $\phi : \mathbb{R} \rightarrow \mathbb{R}$, strictly increasing.

Given a finite menu $F \subseteq \mathcal{F}$, an act $f \in F$, and a state ω , let us define **ex-post regret** of choosing f as $r(f, F, \omega) := \max_{g \in F} \phi(g(\omega) - f(\omega))$.

Define the **minmax regret** associated with choosing f from F as $\rho(f, F) := \max_{\omega \in \Omega} r(f, F, \omega)$; let $C_\rho(F) := \operatorname{argmin}_{f \in F} \rho(f, F)$.

Define $\succsim_\rho \subseteq \mathcal{F}^2 : f \succsim_\rho g \iff \rho(f, \{f, g\}) \leq \rho(g, \{f, g\})$.

1. For simplicity, suppose X is finite. Does $\succsim_\rho$ satisfy completeness? What about independence?
2. For simplicity, suppose X is finite. Prove or provide a counterexample: C_ρ satisfies (i) Sen's α and (ii) Sen's β .
3. Suppose X is convex. Prove or provide a counterexample: $\succsim_\rho$ satisfies uncertainty aversion, i.e., $f \sim_\rho g \implies \frac{1}{2}f + \frac{1}{2}g \succsim_\rho f \sim_\rho g$.
4. Do you find C_ρ reasonable? Why or why not? Does it capture some notion of uncertainty aversion?
5. Now define the **anticipated regret** associated with choosing f from F as $R(f, F) := \mathbb{E}_{\omega \sim \mu}[r(f, F, \omega)]$; let $C_R(F) := \operatorname{argmin}_{f \in F} R(f, F)$. Does C_R sound like a better or a worse model of choice than C_ρ ?

⁷Note the analogy with compound lotteries and having different preferences over the compound lotteries and the reduced lotteries.

⁸This is a concept which we haven't really defined, as we made a point of distinguishing it from GS uncertainty aversion given the latter is defined on Anscombe–Aumann and not Savage acts. However, uncertainty aversion is defined for preferences over acts in an analogous manner as risk aversion is defined for preferences over lotteries; i.e. it means what you would expect.

5. Further Reading

Standard References: Mas-Colell et al. (1995, Chapter 6F), Kreps (2012, Chapter 5.3-5.5), Kreps (1988, Chapters 8, 9, 10).

Related questions/topics: There is an immense literature (both theoretical and experimental) on subjective uncertainty and alternatives, which we've alluded to above. Below I discuss some other related topics.

As we've seen, subjective expected utility is a useful model that recovers beliefs from choice (under some stark assumptions). There is a large literature on how to elicit (subjective) beliefs, both theoretical and experimental, that is, both defining methods that are effective in eliciting beliefs under the broadest assumptions possible, and testing these models. There are methods to elicit beliefs about probability, variance, and confidence intervals, in settings with objective and subjective probabilities, strategic or not; these are typically covered in experimental economics (both the theoretical and the practical side of it).

There is a vast literature on belief updating. We know people are not perfect Bayesian updaters, but there is far less evidence on the issue than is socially optimal. While many terms such as confirmation bias, base-rate neglect, and conservatism have become mainstream some decades ago, we still know little about how pervasive and meaningful these are, what drives them, how it depends on the kind of information presented, on memory, time pressure. Moreover, more needs to be done in modelling these features, both in developing useful, portable models, and in understanding how common deviations from Bayesian updating change our models. Some examples of situations where this may be particularly relevant: herding, social learning (e.g. technology diffusion, bank runs); pricing data and value of information; information diffusion through social networks (e.g. effect of polls, polarisation, fact-checking).

References

- Alós-Ferrer, Carlos, and Maximilian Mihm.** 2021. "Updating Stochastic Choice." *Working Paper* 1–36. 10.2139/ssrn.3811298.
- Anscombe, Frank J., and Robert J. Aumann.** 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34 (1): 199–205. 10.1214/aoms/1177704255.
- Cripps, Martin.** 2019. "Divisible Updating." *Working Paper* 1–39, https://drive.google.com/file/d/1YgfJI6p4E9sH8IoxmB1uL5eLBMJ_WZ7-/view.
- Denti, Tommaso, and Luciano Pomatto.** 2021. "Model and Predictive Uncertainty: A Foundation for Smooth Ambiguity Preferences." *Working Paper* 1–50, http://www.its.caltech.edu/~lpomatto/smooth_model.pdf.
- Ellsberg, Daniel.** 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75 (4): 643–669. 10.2307/1884324.

- Ghirardato, Paolo.** 2002. “Revisiting Savage in a conditional world.” *Economic Theory* 20 83–92. 10.1007/s001990100188.
- Gilboa, Itzhak, and David Schmeidler.** 1989. “Maxmin expected utility with non-unique prior.” *Journal of Mathematical Economics* 18 (2): 141–153. 10.1016/0304-4068(89)90018-9.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji.** 2005. “A Smooth Model of Decision Making under Ambiguity.” *Econometrica* 73 (6): 1849–1892. 10.1111/j.1468-0262.2005.00640.x.
- Kreps, David M.** 1988. *Notes on the Theory of Choice*. Westview Press.
- Kreps, David M.** 2012. *Microeconomic Foundations I. Choice and Competitive Markets*. Princeton, NJ: Princeton University Press.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green.** 1995. *Microeconomic Theory*. Oxford University Press.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. John Wiley and Sons.
- Schmeidler, David.** 1989. “Subjective Probability and Expected Utility without Additivity.” *Econometrica* 57 (3): 571–587. 10.2307/1911053.

Appendix A. Details on Savage’s Framework

To define [Savage’s \(1954\)](#) postulates (the properties we require of $\succsim$ and, indirectly, of Ω) we restate some definitions in the Savage framework:

$E \subseteq \Omega$: an event;

fEg : a ‘conditional act,’ where for acts f, g and event E , $fEg \in \mathcal{F}$ is such that $(fEg)(\omega) = f(\omega)$ if $\omega \in E$ and $(fEg)(\omega) = g(\omega)$ if otherwise;

Null event E : an event such that for any $f, g, h \in \mathcal{F}$ for which $g \succ f$, $fEh \sim gEh$;

$\tilde{x}$: a constant act, $\tilde{x}(\omega) = x, \forall \omega \in \Omega$.

Now the postulates:

P1 (Ordering): $\succsim$ is complete and transitive (a preference relation).

P2 (Sure-Thing Principle): For any acts f, g, h, h' , and any event E , $fEh \succsim gEh$ if and only if $fEh' \succsim gEh'$.

(P2 gives a form of independence.)

P3 (Monotonicity): For every non-null event E and all constant acts, $\tilde{x}$ and $\tilde{y}$, $\tilde{x} \succ \tilde{y}$ if and only if $\tilde{x}Eh \succ \tilde{y}Eh$ for any act h .

(P3 allows us to rank acts based on the ranking of constant acts.)

P4 (Weak Comparative Probability): For all events A, B and constant acts $\tilde{x}, \tilde{x}', \tilde{y}, \tilde{y}'$, such that $\tilde{x} \succ \tilde{y}$ and $\tilde{x}' \succ \tilde{y}'$, then $\tilde{x}A\tilde{y} \succ \tilde{x}B\tilde{y}$ if and only if $\tilde{x}'A\tilde{y}' \succ \tilde{x}'B\tilde{y}'$.

(P4 is crucial to infer from preferences alone whether an event A is more likely than another event B .)

P5 (Nondegeneracy): There are constant acts $\tilde{x}, \tilde{y}$ such that $\tilde{x} > \tilde{y}$.

(P5 just makes it a nontrivial preference relation.)

P6 (Small Event Continuity): For all acts f, g such that $f > g$ and all consequences degenerate acts $\tilde{x}, \tilde{y}$, there is a finite partition $\{E_i\}_{i \in [n]}$ of Ω such that $\tilde{x}E_i f > g$ and $f > \tilde{y}E_i g$ for every $i \in [n]$.

(P6 is a form of Archimedean property)

P7 (Uniform Monotonicity): For every event E and acts f, g , (i) if $fEh > \tilde{g}(\omega)Eh$ for any $\omega \in E$ — i.e., $\tilde{g}(\omega)$ is a constant act equal to x in every state, where $x = g(\omega)$ — and any act h , then $fEh \succsim gEh$; (ii) if $\tilde{f}(\omega)Eh > gEh$ for all $\omega \in E$, then $fEh \succsim gEh$.