An introduction to the theory of coherent lower previsions

Enrique Miranda
Rey Juan Carlos University
Madrid, Spain

Overview, Part I

- Some considerations about probability.
- Coherent previsions and probabilities.
- Coherent lower and upper previsions.
- Sets of desirable gambles and linear previsions.
- Natural extension.

Which is the goal of probability?

Probability seeks to determine the plausibility of the different outcomes of an experiment when these cannot be predicted beforehand.

- What is the probability of guessing the 6 winning numbers in the lottery?
- What is the probability of arriving in 30' from the airport to the center of Prague by car?
- What is the probability of having a sunny day tomorrow?

Aleatory vs. epistemic probabilities

In some cases, the probability of an event A is a property of the event, meaning that it does not depend on the subject making the assessment. We talk then of aleatory probabilities.

However, and specially in the framework of decision making, we may need to assess probabilities that represent *our* beliefs. Hence, these may vary depending on the subject or on the amount of information he possesses at the time. We talk then of subjective probabilities.

The behavioural interpretation

One of the possible interpretations of subjective probability is the behavioural interpretation. We interpret the probability of an event A in terms of our betting behaviour: we are disposed to bet at most $P(A)$ on the event A.

If we consider the gamble I_A where we win 1 if A happens and 0 if it doesn't happen, then we accept the transaction $I_A - P(A)$, because the expected gain is

$$(1 - P(A)) * P(A) + (0 - P(A))(1 - P(A)) = 0.$$

Gambles

More generally, we can consider our betting behaviour on gambles.

A gamble is a bounded real-valued variable on $\mathcal{X}$, $f : \mathcal{X} \to \mathbb{R}$.

It represents a reward that depends on the outcome of the experiment modelled by X.

We shall denote the set of all gambles by $\mathcal{L}(\mathcal{X})$.

Example

Who shall win the next Euro Cup of Basketball?

Consider the outcomes a=Greece, b=Spain, c=Lituania, d=Other.

$\mathcal{X} = \{a, b, c, d\}$.

Consider the gamble $f(a) = 3, f(b) = -2, f(c) = 5, f(d) = 10$.

Depending on how likely I consider each of the outcomes I will accept the gamble or not.

Betting on gambles

Consider now a gamble f on $\mathcal{X}$. We may consider the supremum value μ such that we are disposed to pay μ for f, i.e., such that the reward $f - \mu$ is desirable: it will be the expectation $E(f)$.

- For any $\mu < E(f)$, we expect to have a gain.
- For any $\mu > E(f)$, we expect to have a loss.

Buying and selling prices

I may also give money in order to get the reward: if I am disposed to pay x for the gamble f, then the gamble $f - x$ is desirable to me.

I may also sell a gamble f, meaning that if I am disposed to sell it at a price x then the gamble $x - f$ is desirable to me.

In the case of probabilities, the supremum buying price for a gamble f coincides with the infimum selling price, and we have a fair price for f.

Existence of indecision

When we don't have much information, it may be difficult (and unreasonable) to give a fair price $P(f)$: there may be some prices μ for which we would not be disposed to buy or sell the gamble f.

In terms of desirable gambles, this means that we would be *undecided* between two gambles.

It is sometimes considered preferable to give different values $\underline{P}(f) < \overline{P}(f)$ than to give a precise (and possibly wrong) value.

Lower and upper previsions

The lower prevision for a gamble f, $\underline{P}(f)$, is my supremum acceptable *buying* price for f, meaning that I am disposed to buy it for $\underline{P}(f) - \epsilon$ (or to accept the reward $f - (\underline{P}(f) - \epsilon)$) for any $\epsilon > 0$.

The upper prevision for a gamble f, $\overline{P}(f)$, is my infimum acceptable *selling* price for f, meaning that I am disposed to sell f for $\overline{P}(f) + \epsilon$ (or to accept the reward $\overline{P}(f) + \epsilon - f$) for any $\epsilon > 0$.

Example (cont.)

Consider the previous gamble
$f(a) = 3, f(b) = -2, f(c) = 5, f(d) = 10.$

- If I am certain that Spain is not going to win the Euro Cup, I should be disposed to accept this gamble, and even to pay as much as 3 for it. Hence, I would have $\underline{P}(f) \geq 3$.
- For the infimum selling price, if I think that the winner will be either Spain or Greece, I should sell f for anything greater than 3, because for such prices I will always win money with the transaction. Hence, I would have $\overline{P}(f) \leq 3$.

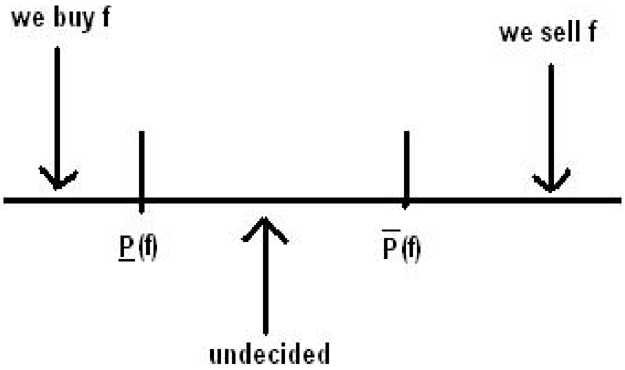

In the precise case we have $\underline{P}(f) = \overline{P}(f)$.

Conjugacy of $\underline{P}, \overline{P}$

Under this interpretation,

$$\begin{aligned}\underline{P}(-f) &= \sup\{x : -f - x \text{ acceptable }\} \\ &= -\inf\{-x : -f - x \text{ acceptable }\} \\ &= -\inf\{y : -f + y \text{ acceptable }\} \\ &= -\overline{P}(f)\end{aligned}$$

Hence, it suffices to work with one of these two functions.

Important remark

The domain $\mathcal{K}$ of $\underline{P}$:

- need not have any predefined structure.
- may contain indicators of events.

Lower probabilities of events

The lower probability of A, $\underline{P}(A)$

= lower prevision $\underline{P}(I_A)$ of the indicator of A.

= supremum betting rate on A.

= measure of the evidence supporting A.

= measure of the strength of our belief in A.

Upper probabilities of events

- The upper probability of A, $\overline{P}(A)$
  - $=$ upper prevision $\overline{P}(I_A)$ of the indicator of A.
  - $=$ measure of the lack of evidence against A.
  - $=$ measure of the plausibility of A.
- We have then a behavioural interpretation of upper and lower probabilities:

$$\text{evidence in favour of } A \uparrow \Rightarrow \underline{P}(A) \uparrow$$

$$\text{evidence against } A \uparrow \Rightarrow \overline{P}(A) \downarrow$$

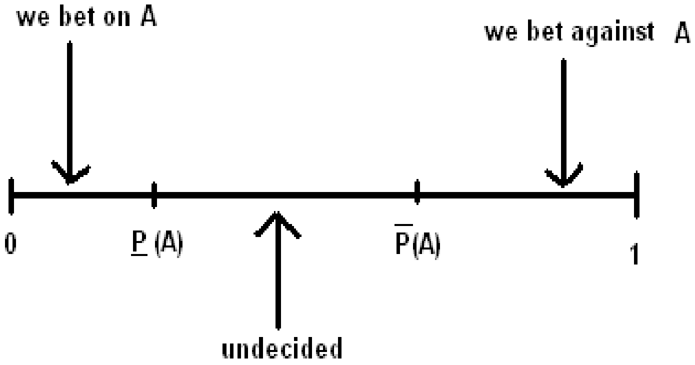
we bet on A
we bet against A
0
P (A)
P(A)
1
undecided

Example(cont.)

- The lower probability we give to Spain being the winner would be the lower prevision of I_b, where we get a reward of 1 if Spain wins and 0 if it doesn't.

- The upper probability of Lithuania or Greece winning would be the upper probability of $I_{\{a,c\}}$, or, equivalently, 1 minus the lower probability of Lithuania and Greece not winning.

Events or gambles?

In the case of probabilities, we are indifferent between betting on events or on gambles: our betting rates on events (a probability) determine our betting rates on gambles (its expectation).

However, in the imprecise case, the lower and upper previsions for events do not determine the lower and upper previsions for gambles uniquely.

Hence, lower and upper previsions are more informative than lower and upper probabilities.

Consistency requirements

The assessments made by a lower prevision on a set of gambles should satisfy a number of consistency requirements:

- A combination of the assessments should not produce a net loss, no matter the outcome: avoiding sure loss.
- Our supremum buying price for a gamble f should not depend on our assessments for other gambles: coherence.

Avoiding sure loss

I represent my beliefs about the possible winner of the EuroCup saying that

$$\overline{P}(a) = 0.55, \overline{P}(b) = 0.25, \overline{P}(c) = 0.4, \overline{P}(d) = 0.1$$
$$\underline{P}(a) = 0.45, \underline{P}(b) = 0.2, \underline{P}(c) = 0.35, \underline{P}(d) = 0.05$$

where $\{a, b, c, d\} = \{Greece, Spain, Lithuania, Other\}$.

This means that the gambles $I_a - 0.44$, $I_b - 0.19$, $I_c - 0.34$ and $I_d - 0.04$ are desirable for me. But if I accept all of them I get the sum

$$[I_a + I_b + I_c + I_d] - 1.01 = -0.01$$

which produces a net loss of 0.01, no matter the outcome of the Cup.

Avoiding sure loss: general definition

Let $\underline{P}$ be a lower prevision defined on a set of gambles $\mathcal{K}$. It avoids sure loss iff

$$\sup_{\omega \in \mathcal{X}} \sum_{i=1}^{n} f_k(\omega) - \underline{P}(f_k) \geq 0$$

for any $f_1, \ldots, f_n \in \mathcal{K}$.
Otherwise, there is some $\epsilon > 0$ such that

$$\sum_{i=1}^{n} f_k - (\underline{P}(f_k) - \epsilon) < -\epsilon$$

no matter the outcome.

Consequences of avoiding sure loss

- $\underline{P}(f) \leq \sup f$.
- $\underline{P}(\mu) \leq \mu \leq \overline{P}(\mu)\ \forall \mu \in \mathbb{R}$.
- If $f \geq g + \mu$, then $\overline{P}(f) \geq \underline{P}(g) + \mu$.
- $\underline{P}(\lambda f + (1-\lambda)g) \leq \lambda \overline{P}(f) + (1-\lambda)\overline{P}(g)$.
- $\underline{P}(f + g) \leq \overline{P}(f) + \overline{P}(g)$.

Coherence

After reflecting a bit, I come up with the assessments:

$$\overline{P}(a) = 0.55, \overline{P}(b) = 0.25, \overline{P}(c) = 0.4, \overline{P}(d) = 0.1$$
$$\underline{P}(a) = 0.45, \underline{P}(b) = 0.15, \underline{P}(c) = 0.30, \underline{P}(d) = 0.05$$

These assessments avoid sure loss. However, they imply that the transaction

$$I_a - 0.44 + I_c - 0.29 + I_d - 0.04 = 0.23 - I_b$$

is acceptable for me, which means that I am disposed to bet against Spain at a rate 0.23, smaller that $\overline{P}(b)$. This indicates that $\overline{P}(b)$ is too large.

Coherence: general definition

A lower prevision $\underline{P}$ is called coherent when given gambles $f_0, f_1, \ldots, f_n$ in its domain and $m \in \mathbb{N}$,

$$\sum_{i=1}^{n}[f_i(\omega) - \underline{P}(f_i)] \geq m[f_0(\omega) - \underline{P}(f_0)]$$

for some $\omega \in \mathcal{X}$.

Otherwise, there is some $\epsilon > 0$ such that

$$\sum_{i=1}^{n} f_i - (\underline{P}(f_i) - \epsilon) < m(f_0 - \underline{P}(f_0) - \epsilon),$$

and $\underline{P}(f_0) + \epsilon$ would be an acceptable buying price for f_0.

Coherence on linear spaces

Suppose the domain $\mathcal{K}$ is a linear space of gambles:

- If $f, g \in \mathcal{K}$, then $f + g \in \mathcal{K}$.
- If $f \in \mathcal{K}, \lambda \in \mathbb{R}$, then $\lambda f \in \mathcal{K}$.

Then, $\underline{P}$ is coherent if and only if for any $f, g \in \mathcal{K}, \lambda \geq 0$,

- $\underline{P}(f) \geq \inf f$.
- $\underline{P}(\lambda f) = \lambda \underline{P}(f)$.
- $\underline{P}(f+g) \geq \underline{P}(f) + \underline{P}(g)$.

Consequences of coherence

Whenever the gambles belong to the domain of $\underline{P}, \overline{P}$:

- $\underline{P}(\emptyset) = \overline{P}(\emptyset) = 0, \underline{P}(\mathcal{X}) = \overline{P}(\mathcal{X}) = 1.$
- $A \subseteq B \Rightarrow \underline{P}(A) \leq \underline{P}(B), \overline{P}(A) \leq \overline{P}(B).$
- $\underline{P}(f) + \underline{P}(g) \leq \underline{P}(f+g) \leq \underline{P}(f) + \overline{P}(g) \leq \overline{P}(f+g) \leq \overline{P}(f) + \overline{P}(g).$
- $\underline{P}(\lambda f) = \lambda \underline{P}(f), \overline{P}(\lambda f) = \lambda \overline{P}(f)$ for $\lambda \geq 0$.
- If $f_n \to f$ uniformly, then $\underline{P}(f_n) \to \underline{P}(f)$ and $\overline{P}(f_n) \to \overline{P}(f)$.

Consequences of coherence (II)

- $\lambda \underline{P}(f) + (1-\lambda)\underline{P}(g) \leq \underline{P}(\lambda f + (1-\lambda)g) \ \forall \lambda \in [0,1]$.
- $\underline{P}(f+\mu) = \underline{P}(f) + \mu \ \forall \mu \in \mathbb{R}$.
- The lower envelope of a set of coherent lower previsions is coherent.
- A convex combination of coherent lower previsions (with the same domain) is coherent.
- The point-wise limit (inferior) of coherent lower previsions is coherent.

Linear previsions

When $\mathcal{K} = -\mathcal{K} := \{-f : f \in \mathcal{K}\}$ and $\underline{P}(f) = \overline{P}(f)$ for all $f \in \mathcal{K}$, then $P = \underline{P} = \overline{P}$ is a called a linear or precise prevision on $\mathcal{K}$. If $\mathcal{K}$ is a linear space, this is equivalent to

- $P(f) \geq \inf f$.
- $P(f + g) = P(f) + P(g)$,

for all $f, g \in \mathcal{K}$.

These are the previsions considered by de Finetti. We shall denote by $\mathbb{P}(\mathcal{X})$ the set of all linear previsions on $\mathcal{X}$.

Linear previsions and probabilities

A linear prevision P defined on indicators of events only is a finitely additive probability.

Conversely, a linear prevision P defined on the set $\mathcal{L}(\mathcal{X})$ of all gambles is characterised by its restriction to the set of events, which is a finitely additive probability on $\mathcal{P}(\mathcal{X})$, through the expectation operator.

Coherence and precise previsions

Given a lower prevision $\underline{P}$ on $\mathcal{K}$, we can consider

$$\mathcal{M}(\underline{P}) := \{P \in \mathbb{P}(\mathcal{X}) : P(f) \geq \underline{P}(f)\ \forall f \in \mathcal{K}_k\}.$$

- $\underline{P}$ avoids sure loss $\iff \mathcal{M}(\underline{P}) \neq \emptyset$.
- $\underline{P}$ coherent $\iff \underline{P} = \min \mathcal{M}(\underline{P})$.

There is a 1-to-1 correspondence between coherent lower previsions and (closed and convex) sets of linear previsions.

This correspondence establishes a sensitivity analysis interpretation to coherent lower previsions.

Example (cont.)

Consider the coherent assessments:

$$\overline{P}(a) = 0.5, \overline{P}(b) = 0.2, \overline{P}(c) = 0.35, \overline{P}(d) = 0.1$$
$$\underline{P}(a) = 0.45, \underline{P}(b) = 0.15, \underline{P}(c) = 0.30, \underline{P}(d) = 0.05$$

The equivalent set of coherent previsions represents the possible models for the probabilities of each team being the winner:

$$\mathcal{M}(\underline{P}) := \{(p_a, p_b, p_c, p_d) : p_a + p_b + p_c + p_d = 1, p_a \in [0.45, 0.5], \\ p_b \in [0.15, 0.2], p_c \in [0.3, 0.35], p_d \in [0.05, 0.1]\}$$

To see that the bounds are attained, it suffices to consider the following elements of $\mathcal{M}(\underline{P})$: $(0.45, 0.15, 0.3, 0.1)$, $(0.45, 0.2, 0.3, 0.05)$, $(0.5, 0.15, 0.3, 0.05)$,$(0.45, 0.15, 0.35, 0.05)$.

Sets of desirable gambles

Given a lower prevision $\underline{P}$, we can consider the set of gambles

$$\mathcal{D} := \{f \in \mathcal{K} : \underline{P}(f) \geq 0\},$$

the set of associated desirable gambles. Conversely, given a set of gambles $\mathcal{D}$ we can define

$$\underline{P}(f) := \sup\{\mu : f - \mu \in \mathcal{D}\}$$

Rationality axioms for sets of desirable gambles

If we consider a set of gambles that we find desirable, there are a number of rationality requirements we can consider:

- A gamble that makes us lose money, no matter the outcome, should not be desirable, and a gamble which never makes us lose money should be desirable.
- A change of utility scale should not affect our desirability.
- If two transactions are desirable, so should be their sum.

These ideas define the notion of coherence for sets of gambles.

Coherence of sets of desirable gambles

A set of desirable gambles is coherent if and only if

- If $\sup f < 0$, then $f \notin \mathcal{D}$.
- If $f \geq 0$, then $f \in \mathcal{D}$.
- If $f, g \in \mathcal{D}$, then $f + g \in \mathcal{D}$.
- If $f \in \mathcal{D}, \lambda \geq 0$, then $\lambda f \in \mathcal{D}$.
- If $f + \epsilon \in \mathcal{D}$ for all $\epsilon > 0$, then $f \in \mathcal{D}$.

- If $\mathcal{D}$ is a coherent set of gambles, then the lower prevision it induces is coherent.
- Conversely, a coherent lower prevision $\underline{P}$ determines a coherent set of desirable gambles through the previous formula.

Hence, we have three equivalent representations of our beliefs:

- Coherent lower and upper previsions.
- Closed and convex sets of linear previsions.
- Coherent sets of desirable gambles,

and we can easily go from any of these formulations to the others.

Is coherence too strong?

Some critics to the property of coherence are:

- Descriptive decision theory shows that sometimes beliefs violate the notion of coherence.
- Coherent lower previsions may be difficult to assign for people not familiar with the behavioural theory of imprecise probabilities.
- Other rationality criteria may be also interesting.

Inference: natural extension

Consider the following gambles:

$$f(a) = 5, f(b) = 2, f(c) = -5, f(d) = -10$$
$$g(a) = 2, g(b) = -2, g(c) = 0, g(d) = 5,$$

and assume we make the assessments $\underline{P}(f) = 2, \underline{P}(g) = 0$. Can we deduce anything about how much should we pay for the gamble

$$h(a) = 7, h(b) = 4, h(c) = -5, h(d) = 0?$$

For instance, since $h \geq f + g$, we should be disposed to pay at least $\underline{P}(f) + \underline{P}(g) = 2$. But can we be more specific?

Definition

Consider a coherent lower prevision $\underline{P}$ with domain $\mathcal{K}$, we seek to determine the consequences of the assessments in $\mathcal{K}$ on gambles outside the domain.

The natural extension of $\underline{P}$ to all gambles is given by

$$\underline{E}(f) := \sup\{\mu : \exists f_k \in \mathcal{K}, \lambda_k \geq 0, k = 1, \dots, n : \\ f - \mu \geq \sum_{i=1}^{n} \lambda_k (f_k(\omega) - \underline{P}(f_k))\}$$

$\underline{E}(f)$ is the supremum acceptable buying price for f that can derived from the assessments on the gambles in the domain.

Example

Applying this definition, we obtain that $\underline{E}(h) = 3.4$, by considering

$$h - 3.4 \geq 1.2(f - \underline{P}(f)).$$

Hence, the coherent assessments $\underline{P}(f) = 2$, $\underline{P}(g) = 0$ imply that we should pay at least 3.4 for the gamble h, but not more.

Natural extension: properties

- If $\underline{P}$ does not avoid sure loss, then $\underline{E}(f) = +\infty$ for any gamble f.
- If $\underline{P}$ avoids sure loss, then $\underline{E}$ is the smallest coherent lower prevision on $\mathcal{L}(\mathcal{X})$ that dominates $\underline{P}$ on $\mathcal{K}$.
- $\underline{P}$ is coherent if and only if $\underline{E}$ coincides with $\underline{P}$ on $\mathcal{K}$.
- $\underline{E}$ is then the least-committal extension of $\underline{P}$: if there are other extensions, they reflect stronger assessments than those in $\underline{P}$.

In terms of sets of linear previsions

Given a lower prevision $\underline{P}$ and its set of dominating linear prevision $\mathcal{M}(\underline{P})$, the natural extension $\underline{E}$ of $\underline{P}$ is the lower envelope of $\underline{P}$. This provides the natural extension with a sensitivity analysis interpretation.

$\underline{P}$ is coherent if and only if $\mathcal{M}(\underline{E}) = \mathcal{M}(\underline{P})$.

In terms of sets of gambles

Consider a coherent set of desirable gambles $\mathcal{D}$. Its natural extension $\mathcal{E}$ is the set of gambles

$$\mathcal{E} := \{g \in \mathcal{L}(\mathcal{X}) : (\forall \delta > 0)(\exists n \geq 0, \lambda_k \in \mathbb{R}^+, f_k \in \mathcal{D})$$
$$g \geq \sum_{k=1}^{n} \lambda_k f_k - \delta\}.$$

It is the smallest coherent set of desirable gambles that contains $\mathcal{D}$. It is the smallest closed convex cone including $\mathcal{D}$ and all non-negative gambles.

All these procedures of natural extension agree with one another: if we consider a coherent lower prevision $\underline{P}$, its set of desirable gambles $\mathcal{D}_{\underline{P}}$, the natural extension of this set $\mathcal{E}_{\mathcal{D}_{\underline{P}}}$ and then the coherent lower prevision associated to this set, we obtain the natural extension of $\underline{P}$.

Hence, we have three equivalent ways of representing our behavioural dispositions:

- Coherent lower previsions.
- Sets of linear previsions.
- Sets of desirable gambles.

Particular cases

The natural extension coincides with some familiar extension procedures for some particular cases of coherent lower previsions:

- Lebesgue integration of a probability measure.
- Choquet integration of 2-monotone lower probabilities.
- Bayes' rule for probability measures.
- Robust Bayesian analysis.
- Logical deduction.

Conditional lower previsions

- Definition.
- Consistency requirements.
- Natural extension.

Updating information

So far, we have assumed that all we know about the outcome of the experiment modelled by is that it belongs to a set $\mathcal{X}$.

But we may have some additional information about this outcome, for instance that it belongs to a set B.

We need to update then our assessment by means of a conditional lower prevision

Example (cont.)

We are in the semifinals of the Euro Cup of basketball, and the remaining teams are Spain, Lithuania, France, and Serbia.

For the gamble f on $\{a, b, c, d\} = \{Greece, Spain, Lithuania, Other\}$ given by $f(a) = 5, f(b) = 2, f(c) = -5, f(d) = -10$, I had given the supremum buying price $\underline{P}(f) = 2$.

But now I should probably lower this supremum buying price, unless I am certain that Spain will be the winner!

Definition

Consider a subset B of $\mathcal{X}$, and a gamble f on $\mathcal{X}$.

$\underline{P}(f|B)$ represents our supremum acceptable buying price for f, after we come to know that the outcome of the experiment belongs to B.

If we consider a partition $\mathcal{B}$ of $\mathcal{X}$, we define $\underline{P}(f|\mathcal{B})$ as the gamble that takes the value $\underline{P}(f|B)$ on the elements of B. It is called a conditional lower prevision.

Separate coherence

A first consistency requirement is that the updated assessments are separately coherent. This means that:

- $\underline{P}(B|B) = 1$ for any $B \in \mathcal{B}$.
- $\underline{P}(\cdot|B)$ is a coherent lower prevision.

- Consequence: $\underline{P}(\cdot|B)$ is determined by its values on B: for any $B \in \mathcal{B}$,

$$I_B h = I_B h' \Rightarrow \underline{P}(h|B) = \underline{P}(h'|B).$$

Example(cont)

The assessments $\underline{P}(a) = 0.45, \underline{P}(b) = 0.15, \underline{P}(c) = 0.3, \underline{P}(d) = 0.05$ that I gave before the championships started, are not coherent anymore: separate coherence implies that

$$\underline{P}(a|Spain, Lithuania, France, Serbia) = 0.$$

I should have $\underline{P}(b, c, d|Spain, Lithuania, France, Serbia) = 1$.

Separate coherence: equivalent formulation

If the domain $\mathcal{K}$ of $\underline{P}(f|\mathcal{B})$ is a linear space, this holds if and only if for any $\lambda \geq 0, f, g \in \mathcal{K}$ and $B \in \mathcal{B}$,

- $\underline{P}(f|B) \geq \inf_{x \in B} f(x)$
- $\underline{P}(f+g|B) \geq \underline{P}(f|B) + \underline{P}(g|B)$
- $\underline{P}(\lambda f|B) = \lambda \underline{P}(f|B)$

Consistency with the initial assessments

Not only our updated lower previsions have to be coherent, but we need them to be coherent with the initial assessments.

For instance, if we consider a gamble f on $\{a, b, c, d\}$ given by $f(a) = -1, f(b) = 0, f(c) = 1, f(d) = 2$ and we make $\underline{P}(f) = 1.5$, it does not make sense that if we learn that the outcome of the experiment is either c or d then we make $\underline{P}(f|\{c, d\}) = 1$.

The connection between unconditional and conditional lower previsions follows from the updating principle:

- We are disposed to accept a gamble f after observing B if and only if the gamble $I_B f$ is desirable.

Coherence of conditional and unconditional previsions

Consider an unconditional lower prevision $\underline{P}$ on a linear space of gambles $\mathcal{K}$ and a conditional lower prevision $\underline{P}(\cdot|\mathcal{B})$ with linear domain $\mathcal{H}$. They are coherent if and only if for any $f_1, f_2 \in \mathcal{K}$, $g_1, g_2 \in \mathcal{H}$ and $B \in \mathcal{B}$,

- $\sup_x[f_1 - \underline{P}(f_1) + g_1 - \underline{P}(g_1|\mathcal{B}) - (f_2 - \underline{P}(f_2))](x) \geq 0.$
- $\sup_x[f_1 - \underline{P}(f_1) + g_1 - \underline{P}(g_1|\mathcal{B}) - I_B(g_2 - \underline{P}(g_2|B))](x) \geq 0.$

Interpretation

The interpretation of this condition is that the supremum acceptable buying price for a gamble (either conditional or unconditional) should not be raised by considering other (conditional or unconditional) gambles.

A similar condition can be given for non-linear domains.

Equivalent conditions

It follows that if $\underline{P}$ and $\underline{P}(\cdot|\mathcal{B})$ are coherent and the domains are rich enough, then

$$\underline{P}(I_B(f - \underline{P}(f|B))) = 0$$

for any $B \in \mathcal{B}$. This is called the Generalised Bayes Rule.

- If $\underline{P}(B) > 0$, then $\underline{P}(f|B)$ is the unique value that satisfies the Generalised Bayes Rule.
- In that case, $\underline{P}(f|B)$ can be calculated as the lower envelope of the values $P(f|B)$, where $P \geq \underline{P}$ and $P(f|B)$ is calculated using Bayes' rule.

Example (cont.)

Given our initial coherent assessments

$$\overline{P}(a) = 0.5, \overline{P}(b) = 0.2, \overline{P}(c) = 0.35, \overline{P}(d) = 0.1$$
$$\underline{P}(a) = 0.45, \underline{P}(b) = 0.15, \underline{P}(c) = 0.30, \underline{P}(d) = 0.05,$$

and if we know that the outcome will belong to $\{b, c, d\}$, we can update them using the envelope theorem, obtaining

$$\underline{P}(b|\{b, c, d\}) = 3/11, \underline{P}(c|\{b, c, d\}) = 6/11, \underline{P}(d|\{b, c, d\}) = 1/11.$$

GBR for linear previsions

When P and $P(\cdot|\mathcal{B})$ are linear and the partition $\mathcal{B}$ is finite, the GBR becomes

$$P(f|B) = \frac{P(fI_B)}{P(B)} \text{ if } P(B) > 0.$$

More generally, if $\mathcal{B}$ is infinite, coherence is equivalent to $P(f) = P(P(f|\mathcal{B}))$ for any gamble f, which is stronger than the GBR.

Natural extension

As in the unconditional case, we can study the behavioural consequences of the assessments given by coherent conditional and unconditional previsions. Given coherent $\underline{P}$ on $\mathcal{K}$ and $\underline{P}(\cdot|\mathcal{B})$ on $\mathcal{H}$, their natural extension is

$$\underline{E}(f) = \sup\{\mu : \exists f_1 \in \mathcal{K}, f_2 \in \mathcal{H}, f - \mu \geq f_1 - \underline{P}(f_1) + f_2 - \underline{P}(f_2|\mathcal{B})\},$$

and

$$\underline{E}(f|B) = \begin{cases} \max\{\beta : \underline{E}(I_B(f-\beta)) \geq 0\} & \text{if } \underline{E}(B) > 0 \\ \sup\{\beta : I_B(f-\beta) \geq I_B(g - \underline{P}(g|\beta)) \text{ for some } g \in \mathcal{H}\} & \text{otherwise} \end{cases}$$

for all $f \in \mathcal{L}(\mathcal{X})$.

Problems with the natural extension

The definition of the natural extension does not always have the properties of the natural extension from the unconditional case:

- In some cases there are no coherent extensions.
- If there are, the natural extension may be only a lower bound of the smallest coherent extensions.
- It provides the smallest coherent extensions when the partition $\mathcal{B}$ is finite.

Other types of extensions

- Regular extension: we compute the lower envelope of the conditional linear previsions dominating $\underline{P}, \underline{P}(\cdot|\mathcal{B})$.
- Marginal extension: we have a conditional lower prevision $\underline{P}(\cdot|\mathcal{B})$ on $\mathcal{L}(\mathcal{X})$ and a lower prevision $\underline{P}$ on the gambles that are constant in the elements of $\mathcal{B}$.