6.4.2 *Lower envelope theorem*

Suppose that $\underline{P}$ is coherent, $\underline{P}(B) > 0$, and $\underline{P}(X|B)$ satisfies the GBR. Define $P(X|B)$ for linear previsions P by Bayes' rule, $P(X|B) = P(BX)/P(B)$. Then $\underline{P}(X|B) = \min\{P(X|B): P \in \mathcal{M}(\underline{P})\} = \min\{P(X|B): P \in \text{ext}\,\mathcal{M}(\underline{P})\}$. On the domain where it is determined by the GBR, $\underline{P}(\cdot|B)$ is separately coherent.

Proof. Using the GBR, $0 = \underline{P}(G(X|B)) = \min\{P(G(X|B)): P \in \mathcal{M}(\underline{P})\} = \min\{P(BX) - \underline{P}(X|B)P(B): P \in \mathcal{M}(\underline{P})\} = \min\{P(B)(P(X|B) - \underline{P}(X|B)): P \in \mathcal{M}(\underline{P})\}$. Hence, since $P(B) \geq \underline{P}(B) > 0$ when $P \in \mathcal{M}(\underline{P})$, $\min\{P(X|B): P \in \mathcal{M}(\underline{P})\} = \underline{P}(X|B)$. By Theorem 3.6.2(c), the minimum is achieved by some extreme point of $\mathcal{M}(\underline{P})$. It follows from linearity of P that $P(\cdot|B)$ is a linear prevision and $P(B|B) = 1$. By Corollary 2.8.6, $\underline{P}(\cdot|B)$ is coherent on its domain, and $\underline{P}(B|B) = 1$ by the GBR. ◆

Provided $\underline{P}(B) > 0$, the updated prevision $\underline{P}(\cdot|B)$ is the lower envelope of the updated class $\mathcal{M}^B$, which is obtained by conditioning the linear previsions in $\mathcal{M}(\underline{P})$ on B. The extreme points of $\mathcal{M}^B$ form a subset of the conditioned extreme points of $\mathcal{M}(\underline{P})$. Often, the simplest way to apply the GBR is to apply Bayes' rule to the extreme points of $\mathcal{M}(\underline{P})$, and compute the lower envelope.

To illustrate this, we give two simple examples. In both examples, the conditional probabilities determined by the GBR are highly imprecise. More substantial examples will be given in section 6.6.

6.4.3 *Two tosses of a coin*

Suppose that a fair coin is tossed twice, in such a way that the second toss may depend on the outcome of the first, but You know nothing about the type or degree of dependence. The model $\underline{P}$ suggested in Example 5.13.4 is the lower envelope of all additive probabilities P which assign $P(H_1) = P(H_2) = \frac{1}{2}$ and have an arbitrary degree of dependence between tosses. The two extreme points P_1 and P_2 of $\mathcal{M}(\underline{P})$ are characterized by $P_1(H_1 \cap H_2) = \frac{1}{2}$ and $P_2(H_1 \cap H_2) = 0$.

Now suppose that You observe the outcome of the first toss, represented by the partition $\mathcal{B} = \{H_1, T_1\}$. Because of the symmetry between heads and tails, it suffices to consider the observation H_1. The extreme points P_1 and P_2 can be conditioned by Bayes' rule, giving $P_1(H_2|H_1) = P_1(H_1 \cap H_2)/P_1(H_1) = 1$, and $P_2(H_2|H_1) = P_2(H_1 \cap H_2)/P_2(H_1) = 0$. The updated lower prevision $\underline{P}(\cdot|H_1)$, obtained as the lower envelope of $P_1(\cdot|H_1)$ and $P_2(\cdot|H_1)$, is therefore vacuous. The GBR produces vacuous updated probabilities concerning the second toss, whatever the outcome of the first, even though Your initial probabilities $\underline{P}(H_2) = \overline{P}(H_2) = \frac{1}{2}$ were precise.

It may seem strange that precise probabilities can become vacuous when

You obtain additional information.[3] To see that it is reasonable, think of the extreme points P_1 and P_2 as hypotheses about how the experiment is performed. Suppose that the first toss is made in the usual way, but the second outcome is completely determined by the first: it is either identical (under hypothesis 1) or opposite (hypothesis 2). The model $\underline{P}$ is reasonable if You are completely ignorant about which hypothesis is true. Before observing the first toss, Your probabilities concerning the second are precise because the two hypotheses imply the same probability $P(H_2) = \frac{1}{2}$. After observing the first toss, Your probabilities become vacuous because the two hypotheses predict different outcomes for the second toss. Observing the first toss introduces indeterminacy concerning the second toss, due to disagreement between the hypotheses.

This shows that receiving extra information can sometimes be a bad thing, in the sense that it is certain to produce indeterminacy and indecision. This effect is actually quite common in practice, especially when artificial randomization is involved. To illustrate, suppose that two treatments are compared by allocating them randomly to two different experimental units, and observing which unit responds better. (For example, two medical treatments might be tried on two people.) Consider the event A, that the unit given treatment 1 will respond better than the unit given treatment 2, and assume (as a null hypothesis) that the two treatments have identical effects. Before the random allocation of treatments, the event A has precise probability $\frac{1}{2}$. But after You learn which unit is allocated treatment 1, Your beliefs about A may be quite indeterminate, depending on what You know about the two units.[4] The standard (frequentist) analysis of randomized experiments is based on the precise initial probabilities, and ignores the extra information about the outcome of the randomization.[5]

6.4.4 *The three prisoners*

The same effect is seen in the problem of the three prisoners, described in Example 5.13.10. Prisoner a learns from the governor either that b will be executed (event B), or that c will be executed (event C), so the partition is $\mathcal{B} = \{B, C\}$. The unconditional lower prevision $\underline{P}$, defined in 5.13.10, is the lower envelope of two linear previsions P_1 and P_2, which are uniform distributions on the sets $\{ab, bc, cb\}$ and $\{ac, bc, cb\}$ respectively.

After observing $B = \{ab, cb\}$, P_1 and P_2 are updated to $P_1(\cdot|B)$, which is uniform on $\{ab, cb\}$, and $P_2(\cdot|B)$, which assigns probability one to $\{cb\}$. The event that a will be reprieved is $R = \{ab, ac\}$. Using the GBR, R has updated probabilities $\underline{P}(R|B) = P_2(R|B) = 0$ and $\overline{P}(R|B) = P_1(R|B) = \frac{1}{2}$. Similarly $\underline{P}(R|C) = 0$ and $\overline{P}(R|C) = \frac{1}{2}$. Whatever the governor tells him, prisoner a can conclude only that he will *probably* be executed.[6]