
Exercise sheet 6

Exercise 1 Let $X_1, \ldots, X_n$ be a sample from the uniform distribution on $(0, \theta)$ where $\theta > 0$ is unknown. Recall that the maximum likelihood estimator of θ is $X_{(n)} := \max(X_1, \ldots, X_n)$. Consider estimators of θ of the form $\widehat{\theta}_b = bX_{(n)}$, $b > 0$. Find the estimator of this form that has the smallest risk $R(\theta, \widehat{\theta}_b) = \mathrm{E}_\theta \mathcal{L}(\theta, \widehat{\theta}_b)$ for all values of $\theta > 0$ (if such an estimator exists). Do this for the squared error loss function $\mathcal{L}_2(\theta, a) = (a - \theta)^2$ and for the absolute error loss function $\mathcal{L}_1(\theta, a) = |a - \theta|$.

Solution 1 For the squared error loss function the risk $R(\theta, \widehat{\theta}_b) = \mathrm{E}_\theta(\widehat{\theta}_b - \theta)^2$ is

$$\mathrm{var}_\theta(\widehat{\theta}_b) + (bias_\theta(\widehat{\theta}_b))^2 = \tfrac{b^2 n}{(n+1)^2(n+2)}\theta^2 + \left(\tfrac{bn}{n+1} - 1\right)^2 \theta^2,$$

which is minimized at $b^* = (n+2)/(n+1)$.

For the absolute error loss the risk $R(\theta, \widehat{\theta}_b) = \mathrm{E}_\theta|\widehat{\theta}_b - \theta|$ equals

$$\int_0^{\theta/b} (\theta - by)\tfrac{ny^{n-1}}{\theta^n} dy + \int_{\theta/b}^{\theta} (by - \theta)\tfrac{ny^{n-1}}{\theta^n} dy = \tfrac{\theta}{b^n}\left[\tfrac{2}{n+1} + b^n\left(\tfrac{n}{n+1}b - 1\right)\right]$$

for $b > 1$, and

$$\int_0^{\theta} (\theta - by)\tfrac{ny^{n-1}}{\theta^n} dy = \theta\left(1 - b\tfrac{n}{n+1}\right)$$

for $b \leq 1$.

Note that the risk function is therefore defined by parts, thus we have to find a minimizer for each part, and then compare the values of the minima to find the global minimum of the function. The minimizer for the equation $b > 1$ is $b^* = 2^{1/(n+1)}$, and the one for the equation $b \leq 1$ is $b = 1$.

Since both expressions are equal at $b = 1$, and that the value of the risk for $b = b^*$ is strictly smaller than for $b = 1$, the minimum is attained at b^*.

In both cases the optimal value of b is the same independently of the true value of θ, hence the estimators are uniformly best.

Exercise 2 Let $X_1, X_2, \ldots, X_n$ be a sample from a $N(\mu, \sigma^2)$ distribution, where $n > 1$ and both μ and σ^2 are unknown. The MLE of σ^2 can be shown to be

$$S_n^2 = \tfrac{1}{n}\sum_{i=1}^{n}(X_i - \overline{X})^2.$$

(you do not need to prove this, but if you never this result, you should prove it).

(a) Show that S_n^2 is inadmissible for σ^2 for the squared error loss function.

(b) Find the risk of the minimax estimator of σ^2 in the class of estimators of the form $a_n \sum_{i=1}^{n}(X_i - \overline{X})^2$ when $\sigma^2 \in (0, M]$, $M < \infty$.

(c) Why do we consider $\sigma^2 \in (0, M]$ in the previous point? What fails if we consider $\sigma^2 \in (0, \infty)$?

Solution 2

(a) Developing the risk of $a_n \sum_{i=1}^{n}(X_i - \overline{X}_n)^2$ as a function of a_n, shows that it is minimised uniquely when $a = (n+1)^{-1}$ for all $\sigma^2 > 0$. But S_n^2 corresponds to $a_n = 1/n$, and is thus inadmissible.

(b) The best estimator in this class is $T_n^* = \frac{n}{n+1}S_n^2$ with risk $\frac{2\sigma^4}{n+1}$. Therefore, among this class, T_n^* is unique minimax with supremum risk $2M^4/(n+1)$.

(c) If we consider the same setting but with $\sigma^2 \in (0, \infty)$, then the supremum risk of T_n^2 is infinite, and therefore any decision rule in this class is minimax.

Exercise 3

1. Show that a unique minimax rule is admissible.

2. Show that a Bayes rule with constant risk is minimax.

3. Show that an admissible estimator with constant risk is minimax.

Exercise 4 Assume that $\Theta = \{\theta_1, \ldots, \theta_t\}$ is a finite parameter space and the space of decision rules $\mathcal{D}$ is such that it includes all randomised rules. Define the risk set to be a subset S of $\mathbb{R}^t$ of the form $S = \{(R(\theta_1, d), \ldots, R(\theta_t, d)) : d \in \mathcal{D}\}$.

Show that S is a convex set.

Solution 3

1. Let δ be minimax. If δ' dominates δ, then δ' is also minimax, and therefore δ is not unique minimax.

2. If δ is π-Bayes with constant risk A then for any δ',

$$\sup_\theta R(\theta, \delta') \geq \int_\Theta R(\theta, \delta')\pi(\theta)d\theta \geq \int_\Theta R(\theta, \delta)\pi(\theta)d\theta = \int_\Theta A\pi(\theta)d\theta = A.$$

   Therefore δ is minimax.

3. If δ has constant risk A and is not minimax, then there is a δ' with supremum risk smaller than A. Thus $R(\theta, \delta') < A = R(\theta, \delta)$ for all θ, and δ is inadmissible.

Solution 4 Suppose $x_1 = R(\theta_1, d_1), \ldots, R(\theta_t, d_1))$ and $x_2 = R(\theta_1, d_2), \ldots, R(\theta_t, d_2))$ are two elements of S, and suppose $\lambda \in (0, 1)$. Form a new randomised decision $d = \lambda d_1 + (1 - \lambda)d_2$. Then for every $\theta \in \Theta$, by definition of a randomised decision rule,

$$R(\theta, d) = \lambda R(\theta, d_1) + (1 - \lambda)R(\theta, d_2).$$

Then we see that $\lambda x_1 + (1 - \lambda)x_2$ corresponds to d, and hence is itself a member of S.

Exercise 5 Consider a parameter space with two values $\Theta = \{\theta_1, \theta_2\}$. In each plot in Figure 1, coordinates are $r_1 = R(\theta_1, d)$, $r_2 = R(\theta_2, d)$, the dots and/or thick curves correspond to the values (r_1, r_2) of risk of some non-randomised decision rules, the filled ares are risk sets consisting of points corresponding to all non-randomised and randomised decisions (all convex combinations of the points corresponding to non-randomised decisions). For each risk set:

(a) Draw the set of admissible decisions.

(b) Draw curves corresponding to the decision rules with the same value of the maximal risk $\max(R(\theta_1, d), R(\theta_2, d))$ (i.e., for various values of c, draw "iso-max-risk" curves satisfying $\max(r_1, r_2) = c$).

(c) Use these curves to find the minimax decision(s). Discuss whether it is (they are) unique, randomised, admissible.

(d) Suppose we have prior probabilities $\pi_1 = \pi(\theta_1) \geq 0$, $\pi_2 = \pi(\theta_2) \geq 0$, $\pi_1 + \pi_2 = 1$. Draw curves corresponding to the decisions with the same value of the Bayes risk $\pi_1 R(\theta_1, d) + \pi_2 R(\theta_2, d)$ (i.e., for various values of c, draw "iso-Bayes-risk" curves satisfying $\pi_1 r_1 + \pi_2 r_2 = c$). Do this and the next step for various prior probabilities, for example for $\pi_1 = 0.5, 0.25, 0.75, 0, 1$.

(e) Use these curves to find the Bayes decision(s). Discuss if it is (they are) unique, randomised, admissible.

Solution 5 See Figures 2, 3, 4.

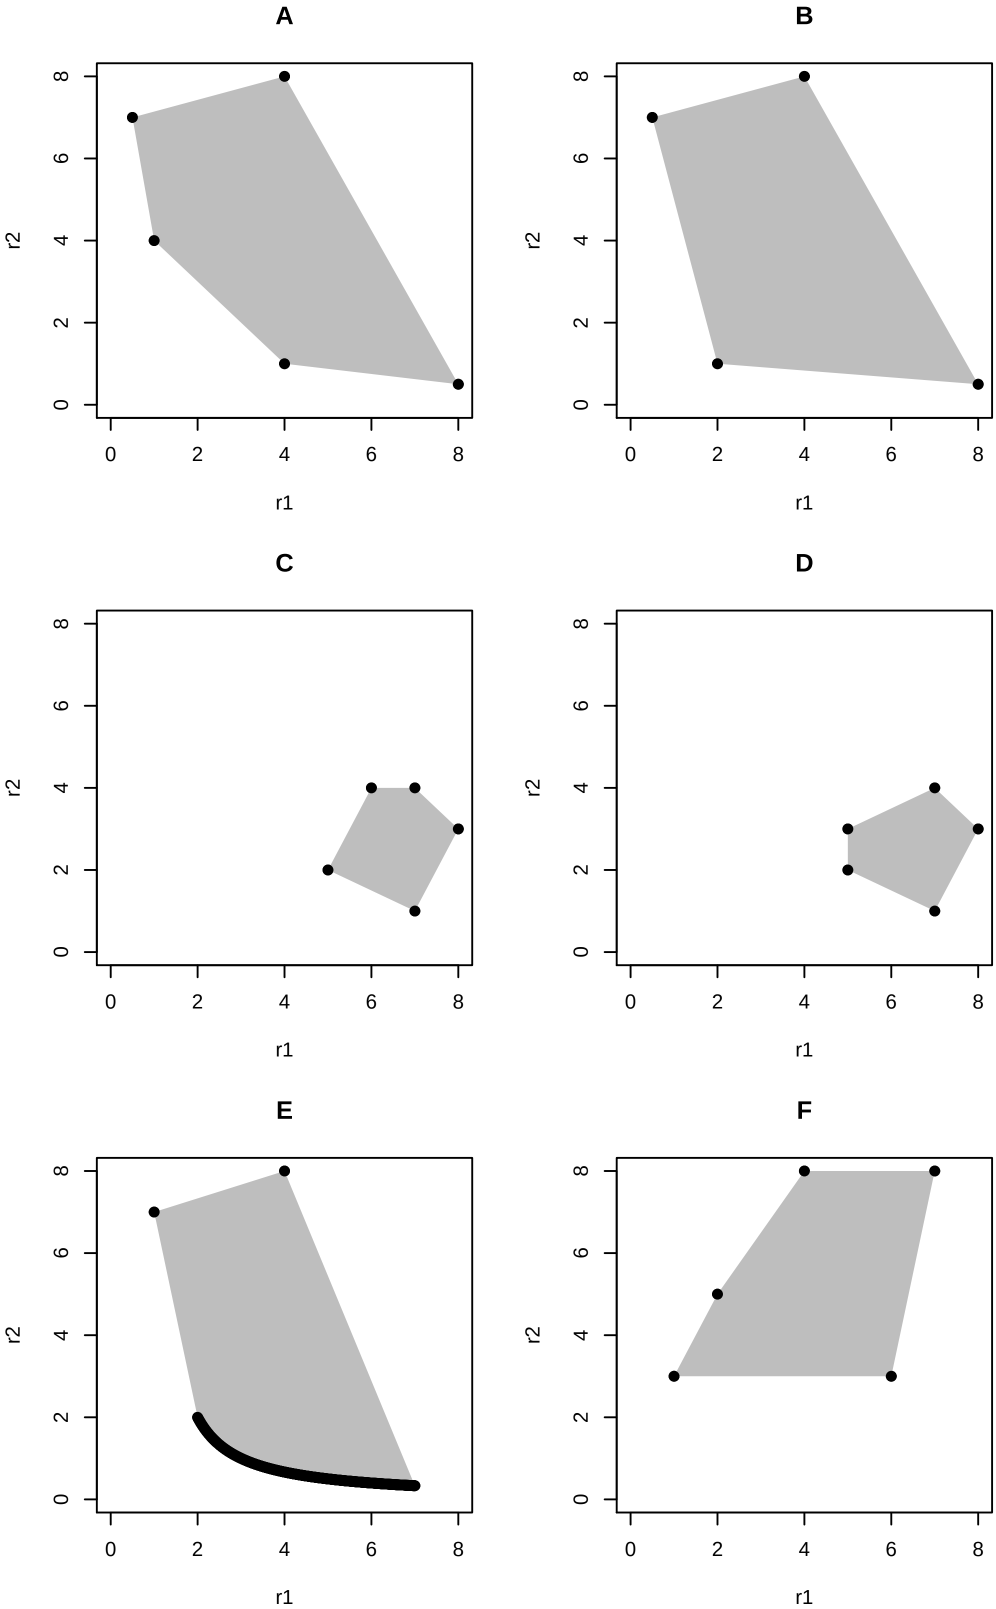

Figure 1: Risk values for non-randomised (black) and randomised (grey) decision rules

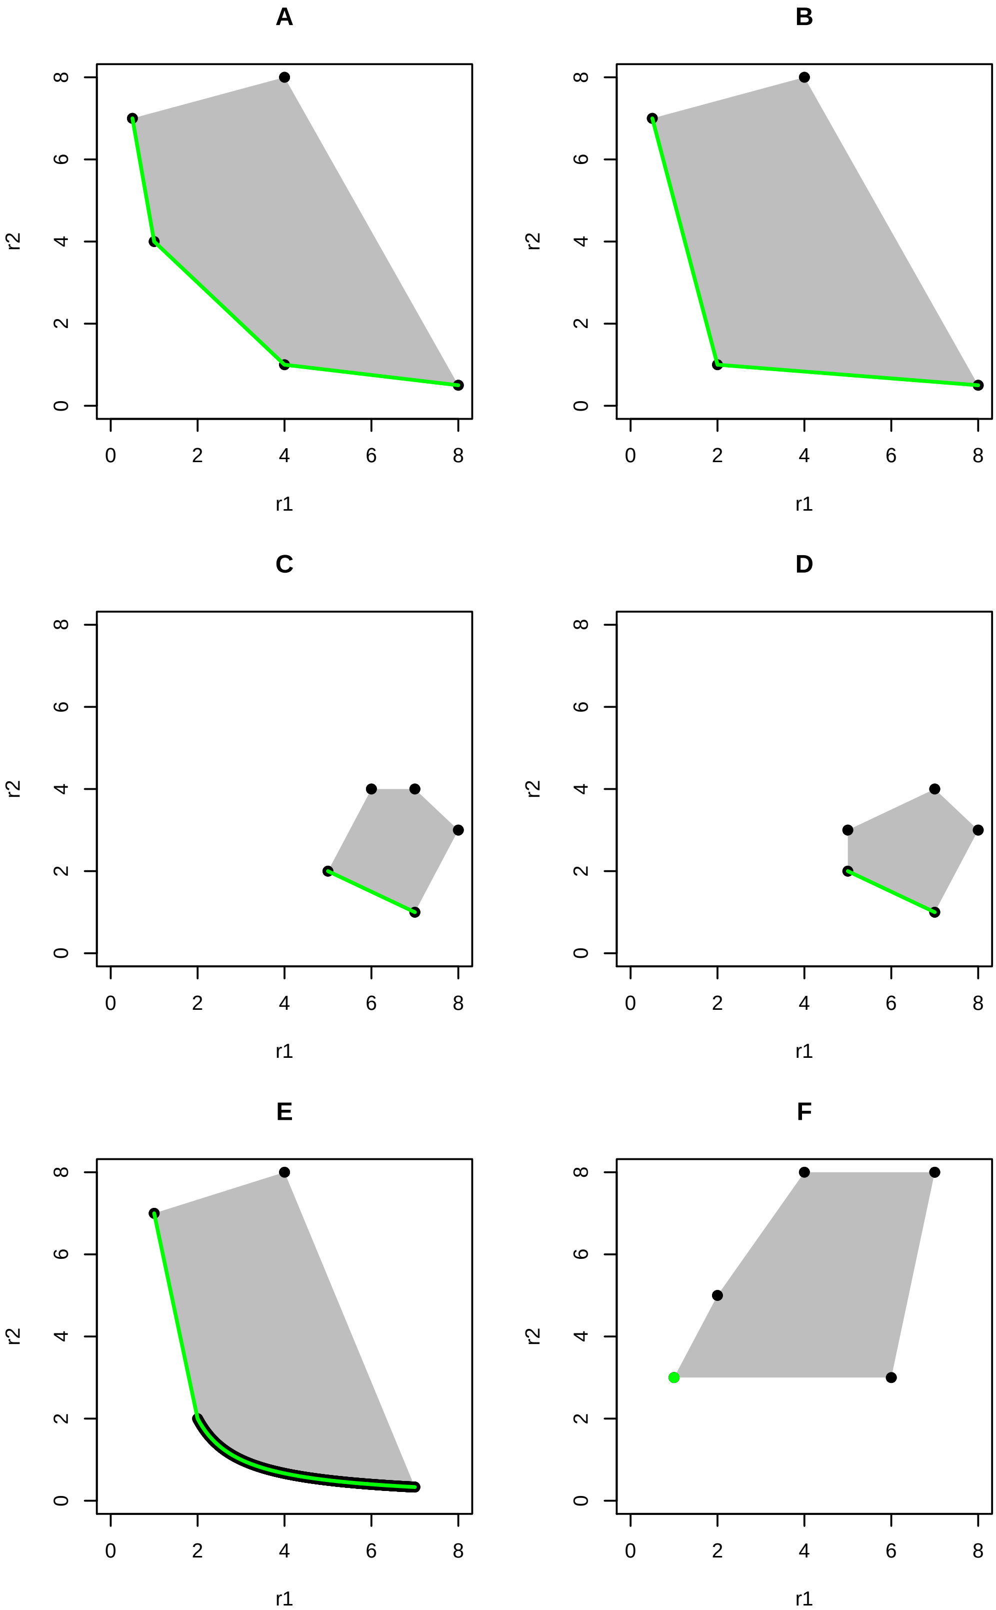

Figure 2: Admissible decisions

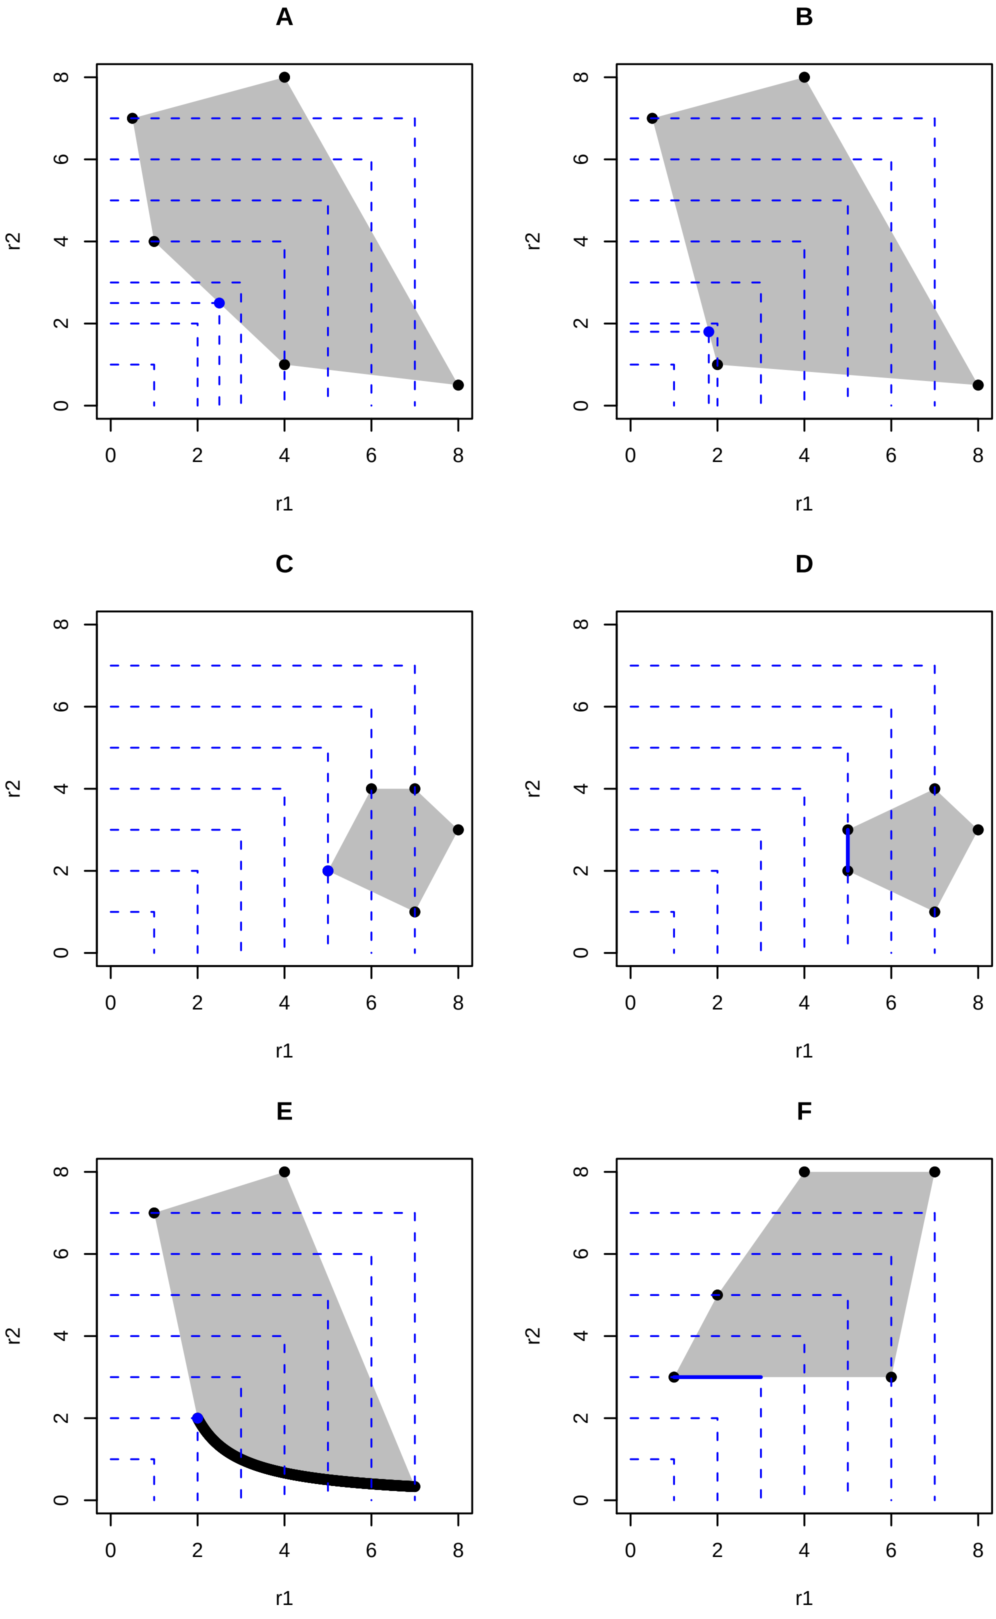

Figure 3: “Iso-max-risk” curves and minimax decisions

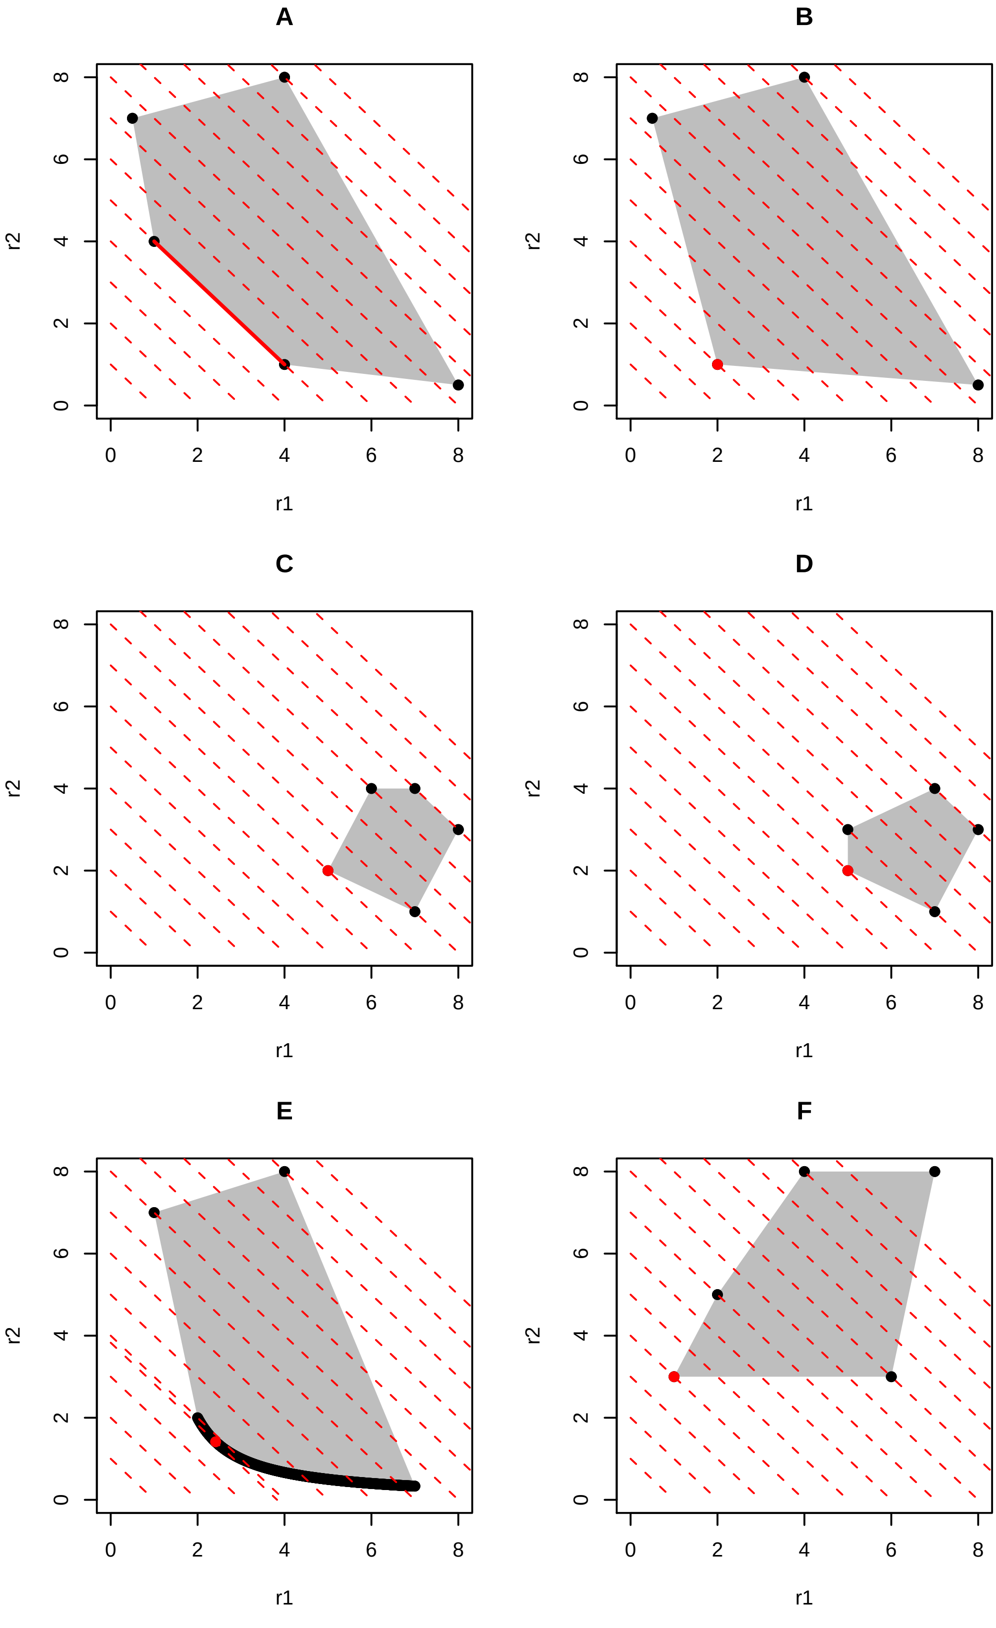

Figure 4: "Iso-Bayes-risk" curves and Bayes decisions for $\pi_1 = \pi_2 = 0.5$