Exercise: Optimal threshold on classification probability

Consider a case where we have learned a conditional probability distribution $P(y|\mathbf{x})$. Suppose there are only two classes, and let $p_0 = P(Y = 0|\mathbf{x})$ and $p_1 = P(Y = 1|\mathbf{x})$. Consider the loss matrix below:

predicted label $\hat{y}$	true label y	
	0	1
0	0	λ_{01}
1	λ_{10}	0

1. Show that the decision $\hat{y}$ that minimizes the expected loss is equivalent to setting a probability threshold θ and predicting $\hat{y} = 0$ if $p_1 < \theta$ and $\hat{y} = 1$ if $p_1 \geq \theta$. What is θ as a function of λ_{01} and λ_{10}? (Show your work.)

2. Show a loss matrix where the threshold is 0.1. (Show your work.)