Statistical Learning - Vienna, Fall2017

Homework problem 4

The effect of identical predictors in Ridge and Lasso (3.28,3.29 in ESL 2nd ed.)

Assume we have a univariate model with one x variable and no intercept. We fit constrained ridge regression and lasso with a given constraint s on the norm (ℓ_2 norm squared for ridge, ℓ_1 norm for lasso). Now we add a second identical variable $x^* = x$ and refit the models with the same constraint. What happens to the coefficients $\hat{\beta}$ of both models? How does the two-dimensional solution to the new problem relate to the one-dimensional solution to the old one in each case? Is it unique? Assume the constraint is much smaller than the norm of the least squares solution, so it is tight.

Hint: The behavior of ridge and lasso under this scenario is quite different. Since both predictors x, x^* are identical, a coefficient can be divided between them in different ways which give the same fit. Consider what different divisions do to the norm of the coefficient vector in each case, and use that to infer the optimal solution. You can also simulate to gain intuition.