Calibration of Machine Learning Models in glmnetr

Walter K. Kremers, Mayo Clinic, Rochester MN

24 October 2024

The Package

The "An Overview of glmnetr" vignette shows how to run the main package function nested.glmnetr() and how to summarize model performances. If one identifies a well performing model according to the metrics in this summary, e.g. concordance, correlation, deviance ratio, linear calibration, one may want to do further evaluation in terms of calibration. The strongest calibration and validation will involve calibration with new, independent datasets. Frequently one will not have immediate access to such new data sets, or one may want first to do an internal validation before subjecting a model to an external validation. Here we consider an internal validation approach using cross validation or bootstrap re-sampling, similar to how we numerically assessed model performance.

An Example Analysis

To explore calibration we first consider the nested.glmnetr() call from the "An Overview of glmnetr" vignette which fit machine learning models to survival data with family="cox", i.e.

```
set.seed(465783345)
nested.cox.fit = nested.glmnetr(xs, NULL, yt, event, family="cox",
                                dolasso=1, dostep=1, steps_n=40, folds_n=10, track=1)
```

Linear Calibration

Using either print() or summary() on the output object nested.cox.fit one gets, amongst other information, summaries for the linear calibration slopes and intercepts as in

```
summary( nested.cox.fit )
```

```
##   Sample information including number of records, events, number of columns in
##   design (predictor, X) matrix, and df (rank) of design matrix:
##               family                    n               nevent
##                  cox                 1000                  698
##           xs.columns                xs.df      null.dev/nevent
##                  100                   94                12.43
## null.m2LogLik/nevent  sat.m2LogLik/nevent
##                12.43                    0
##
##  For LASSO, and Stepwise regression tuned by df and p, average (Ave) model
```

```
##  performance measures from the 10-fold (NESTED) Cross Validation are given together
##  with naive summaries calculated using all data without cross validation
##
##                       Ave DevRat Ave Slope Ave Concordance Ave Non Zero
## LASSO min                 0.2446    1.0742          0.8705         45.4
## LASSO minR                0.2434    0.9790          0.8701         17.7
## LASSO minR.G0             0.2410    0.9320          0.8691         16.2
## Ridge                     0.2244    1.2850          0.8626         99.0
##                       Naive DevRat Naive Concordance Non Zero
## LASSO min                   0.1696            0.8794       42
## LASSO minR                  0.1720            0.8793       22
## LASSO minR.G0               0.1728            0.8796       20
## Ridge                       0.1718            0.8822       99
##
##                       Ave DevRat Ave Slope Ave Concordance Ave Non Zero
## Stepwise df tuned         0.2501    0.9590          0.8738         14.7
## Stepwise p tuned          0.2527    0.9677          0.8747         14.8
##                       Naive DevRat Naive Concordance Non Zero
## Stepwise df tuned           0.1711            0.8785       15
## Stepwise p tuned            0.1711            0.8785       15
```

Here we see that for many of the models the linear calibration slope term is near 1, the ideal for perfect calibration. For the Cox model any intercept term can be absorbed into the baseline survival function and there is no pertinent intercept term for calibration.

A First Visual

An initial calibration consideration was made in the overview vignette by regressing observed outcomes on the predicteds from the final model based upon the relaxed lasso. This regression was made using splines, in particular the pspline() function from within a coxph() call, as in

```
# Get predicteds from CV relaxed lasso model embedded in nested CV outputs & Plot
xb.hat = predict( object=nested.cox.fit , xs_new=xs, lam=NULL, gam=NULL, comment=FALSE)
# Fit a spline to xb.hat uisng coxph, and plot
#library( survival )                  ## load survival package for Cox model fits
fit1 = coxph(Surv(yt, event) ~ pspline(xb.hat, df=3))
```

```
summary(fit1)
```

```
## Call:
## coxph(formula = Surv(yt, event) ~ pspline(xb.hat, df = 3))
##
##   n= 1000, number of events= 698
##
##                            coef  se(coef) se2     Chisq   DF   p
## pspline(xb.hat, df = 3),  1.068 0.03321  0.03321 1034.15 1.00 6.8e-227
## pspline(xb.hat, df = 3),                           2.78 2.03  2.5e-01
##
##              exp(coef) exp(-coef) lower .95 upper .95
## ps(xb.hat)3  1.056e+01  9.474e-02 1.627e+00 6.847e+01
## ps(xb.hat)4  1.114e+02  8.976e-03 4.088e+00 3.036e+03
```

```
## ps(xb.hat)5  1.175e+03  8.514e-04 1.646e+01 8.384e+04
## ps(xb.hat)6  1.196e+04  8.364e-05 1.125e+02 1.271e+06
## ps(xb.hat)7  1.285e+05  7.779e-06 1.213e+03 1.362e+07
## ps(xb.hat)8  1.737e+06  5.758e-07 1.660e+04 1.818e+08
## ps(xb.hat)9  1.428e+07  7.005e-08 1.347e+05 1.513e+09
## ps(xb.hat)10 1.648e+08  6.068e-09 1.492e+06 1.820e+10
## ps(xb.hat)11 2.320e+09  4.310e-10 1.778e+07 3.029e+11
## ps(xb.hat)12 3.349e+10  2.986e-11 1.633e+08 6.870e+12
##
## Iterations: 4 outer, 14 Newton-Raphson
##      Theta= 0.7960027
## Degrees of freedom for terms= 3
## Concordance= 0.879  (se = 0.005 )
## Likelihood ratio test= 1501  on 3.03 df,   p=<2e-16
```

followed by plotting with

```
termplot(fit1,term=1,se=TRUE, rug=TRUE, lwd.term=2, lwd.se=2, lty.se=1) # , col.term=1, col.se=2
abline(a=0,b=1,lty=3)
title ("Naive calibration curve for relaxed lasso model")
```

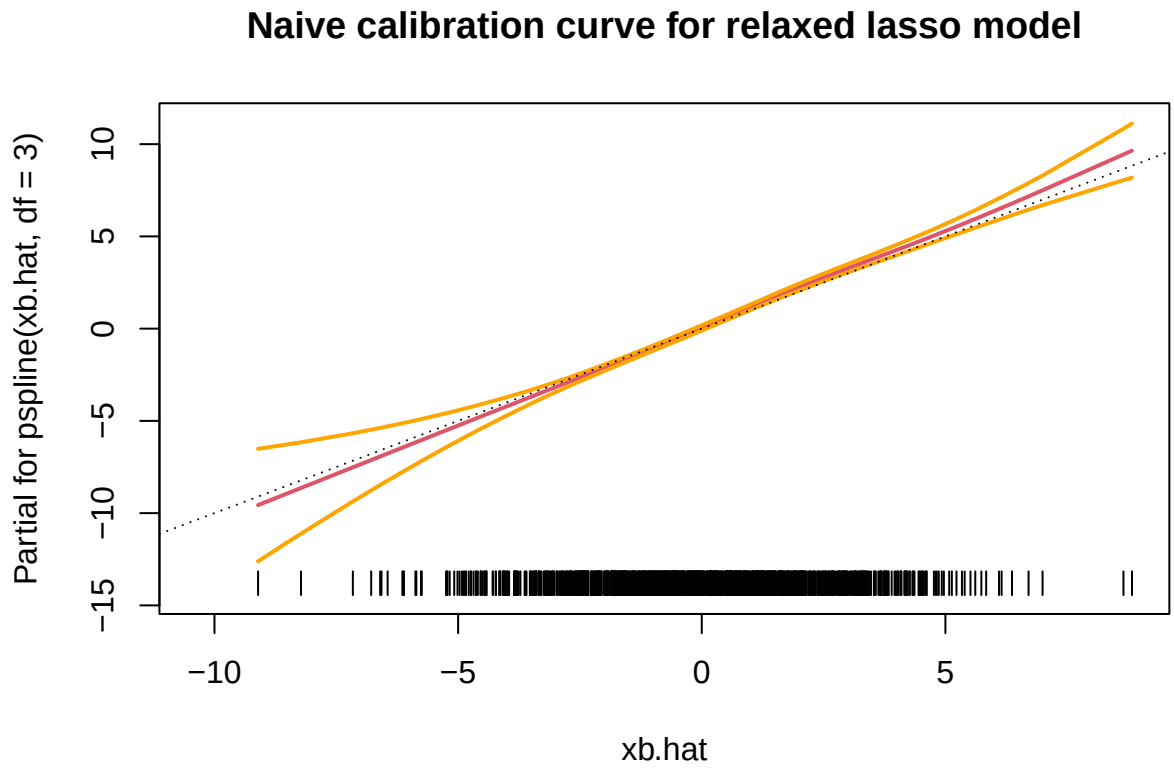

The spline fits may help to understand potential nonlinearities in the model. Here we see, a clibration line which is not far from linear. Still, as noted in the "An Overview of glmnetr" vignette, because the same data are used for model evaluation as well as model derivation, it is hard to put much confidence in such a calibration plot because of potential bias which may suggest a better fit than can be expected for new data.

Calibration Using Spline Fits and Resampling

For each of the models fit, nested.glmnetr() saves the X* Beta's from the final model. The nested.glmentr() function also calculates the X* Betas's for the hold out data for each partitioning, i.e. each hold out fold of the outer loop of nested cross validation or the out-of-bag items not selected by the sample whit replacement of the bootstrap sample. In this manner there are multiple subsets, e.g. k from the k-fold nested CV, or calculation of X*Betas based upon independent observations, and each of these subsets can contribute to calibrate the final model. While each of these calibrations will individually have limited information, when combined following the principles of cross validation for boostrap sampling, they will collectively provide a more meaningful evaluation. This is done by the calplot() function as in

```
calplot(nested.cox.fit, wbeta=5)
```

```
##   Range of X*Beta for calibration:
## -9.10632 8.827991
##   Range of calibrated confidence intervals:
## -12.05814 11.24302
```

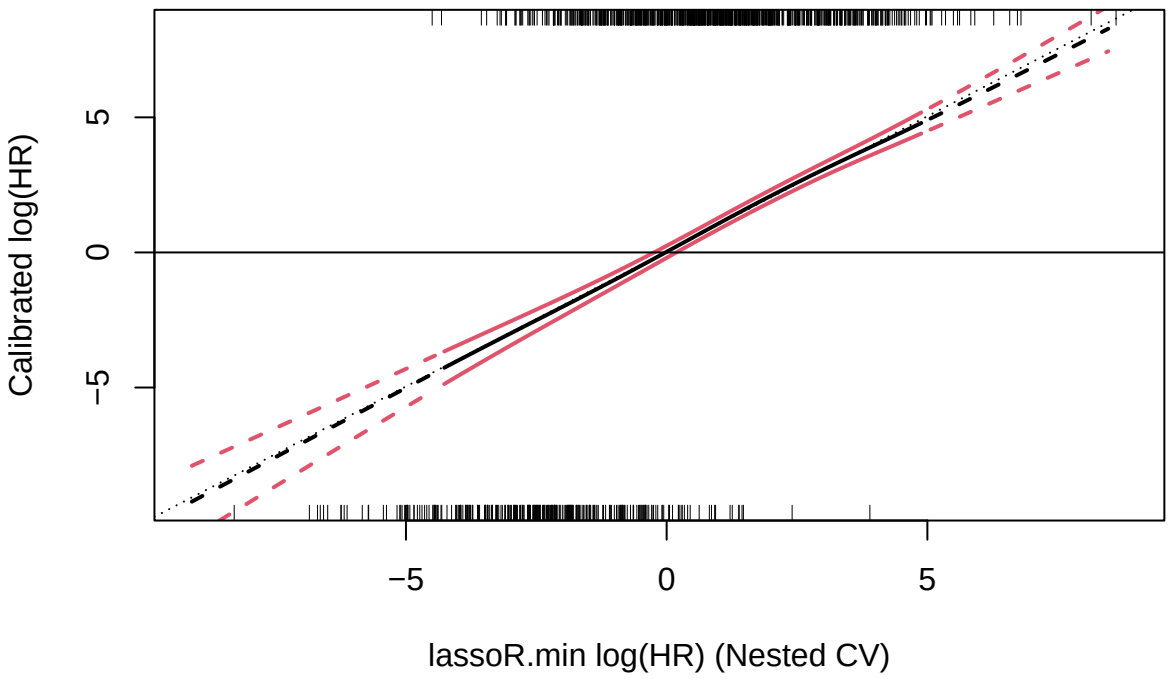


```
##   Due to user specified xlim  1 tick marks are not displayed in rug
##   min:max xb =  -9.32657108298515 8.61924886334891
```

Here we see a smooth, nearly linear predicted log hazard ratio as a function of the model X* Beta from the relaxed lasso model. The bounding lines in red depict the average +/- 2 standard errors (SE) to assist in assessing meaningfulness in any deviation from the ideal identity line, and non linearities. In these curves the

central region with solid lines denotes the region within the range of all the calibration spline fit, i.e. spline fits from all the different leave-out folds of the CV overlap without extrapolation. The dashed lines depict areas out of range for at least one of the leave out folds. Because spline fits can be rather uncertain when extrapolating beyond the data range, one should be more cautious in making strong conclusions in the dashed regions of these plots. Here we see somewhat wider confidence bounds about the overall calibration curve.

A naive calibration curve similar to that shown above can also be easily gotten using the calplot function when specifying to not use the resample for construction as in

```
calplot(nested.cox.fit, wbeta=5, resample=0, xlim=c(-10,9), ylim=c(-15,10))
```

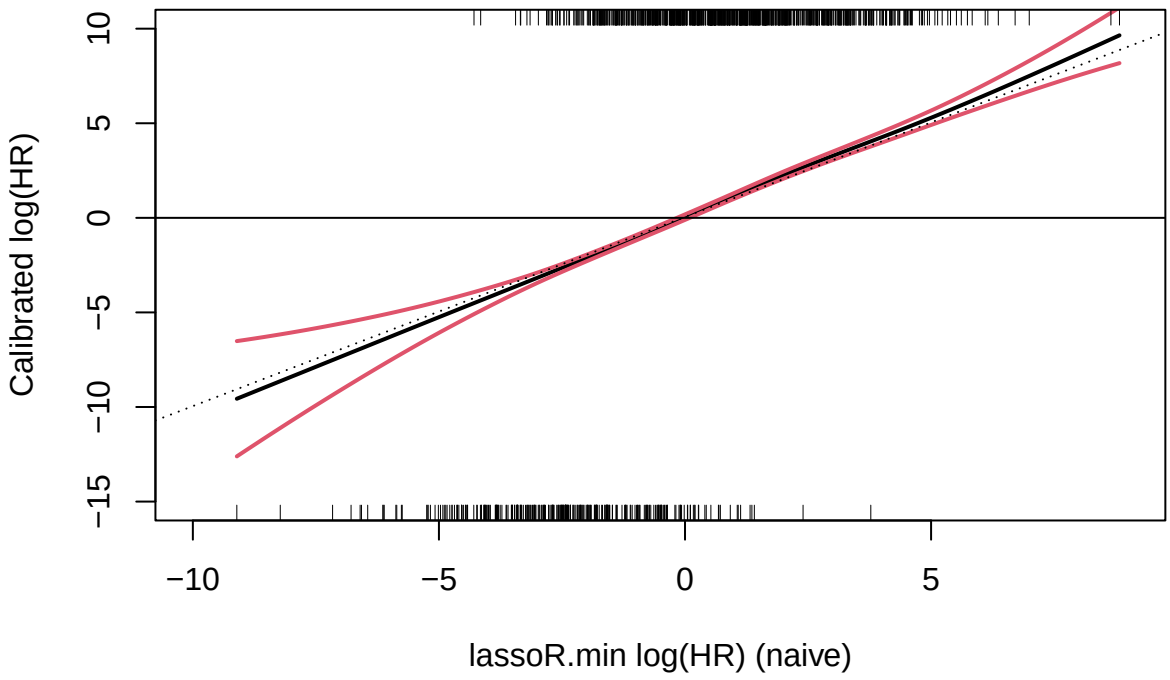

The code above using the termplot() function is provided to show our general approach for derivation of the calibration curves.

In the Nested CV figure we see two rugs, one below and one above the plotted region. The rug below depicts the model X* Beta's which are not associated with an event and the rug above depicts X* Beta's which are associated with events. When there are lots of data points it can be hard to read these rugs. One can use the vref option in calplot to draw two vertical lines where the first separates the smaller vref% of the X* Beta's form the rest, and a second which separates the larger vref% of the data. To depict the hazard ratios (HR) instead of the X*Beta for the Cox model one can use the option plothr, where one assigns a numerical value for the product between tick marks, e.g. exp(1) or 10. Combining these two options we have the example

```
calplot(nested.cox.fit, wbeta=5, vref=1, plothr=100)
```

```
##   Range of X*Beta for calibration:
## -9.10632 8.827991
```

```
##   Range of calibrated confidence intervals:
## -12.05814 11.24302
```

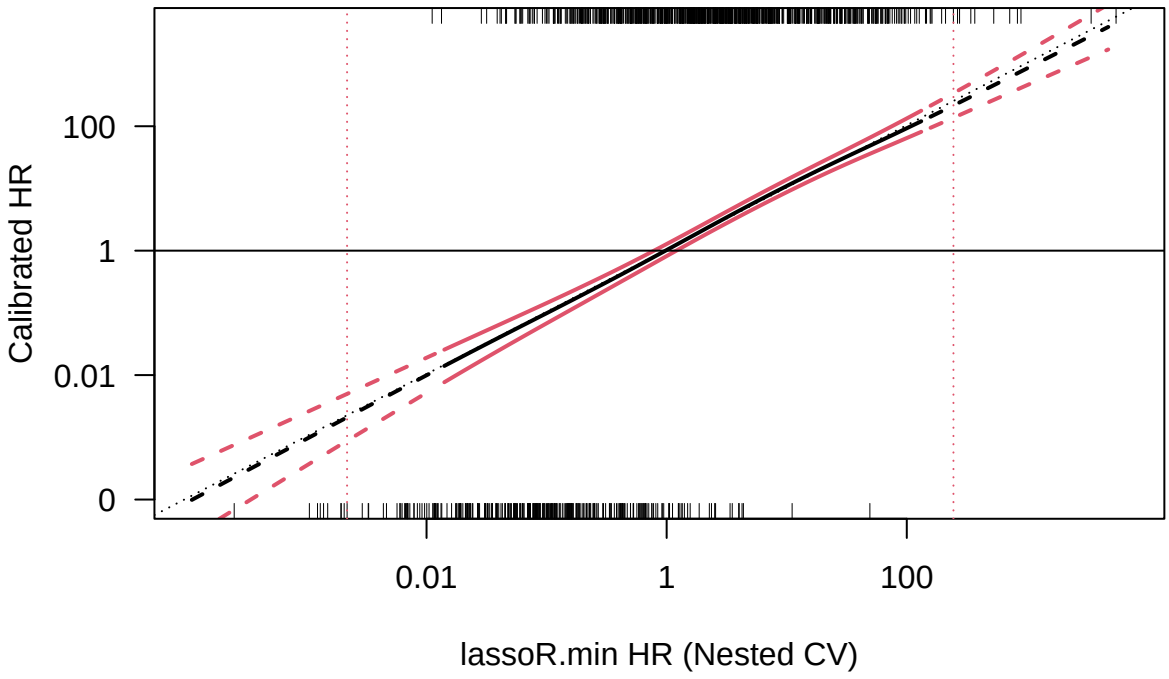


```
##   Due to user specified xlim  1 tick marks are not displayed in rug
##   min:max xb =  -9.32657108298515 8.61924886334891
```

The user can also use different colors for the lines with the options col.term, col.se. One can also specify xlim and ylim in case a few data points cause an excessive amount of white space or odd aspect ratio in the plots.

To view the calibration plots form the individual leave out cross validation folds, one may specify foldplot= 1. In that this generates many figures, we omit in this vignette actually producing plots using this option specification, and instead assign plotfold=1 which overlays the individual calibration curves, albeit without the +/- 2 SE limits for the individual CV folds. The overall calibration (average of the individual CV fold calibrations) and overall +/- 2 SE limits though are maintained.

```
calplot(nested.cox.fit, 5, plotfold=1, vref=1, plothr=100)
```

```
##   Range of X*Beta for calibration:
## -9.10632 8.827991
##   Range of calibrated confidence intervals:
## -12.05814 11.24302
```

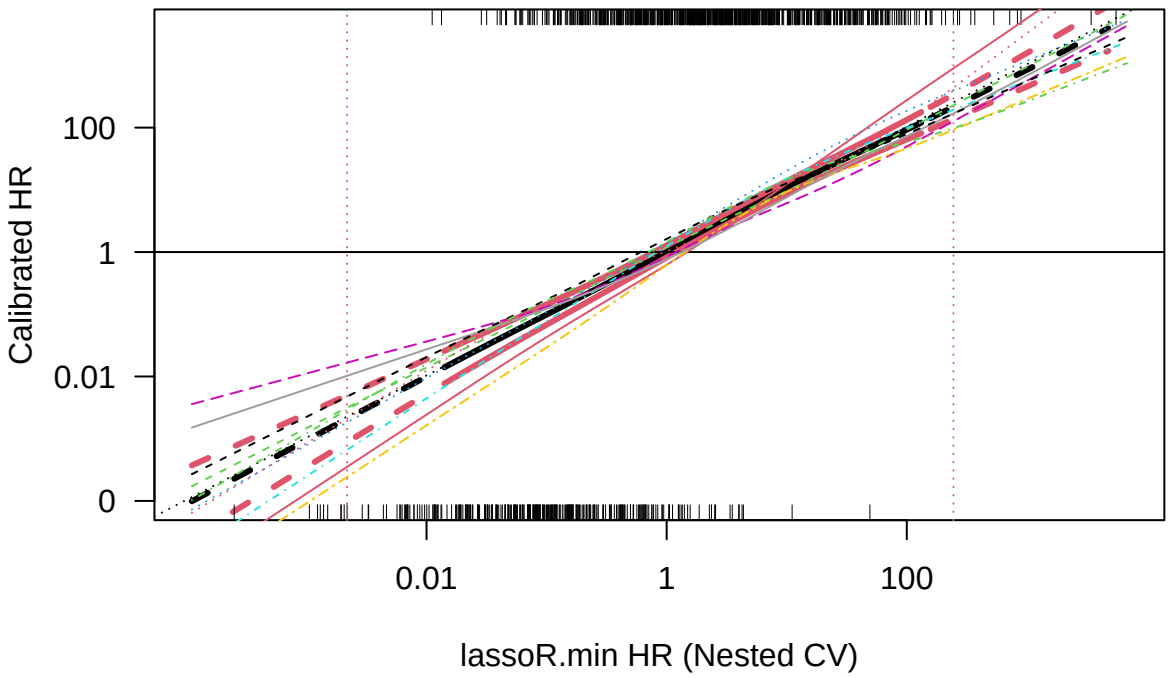


```
##   Due to user specified xlim  1 tick marks are not displayed in rug
##   min:max xb =  -9.32657108298515 8.61924886334891
```

As we see from the above calls the first term in the calplot() function call is an output object form a nested.glmnetr() call. The second term, wbeta, specifies "which beta" or model is to be used for deriving the model X*Beta's. Here, as we see in the figure x-axis label, the 5 determines the relaxed lasso model. Instead of making a hard to remember key the user can leave this term unspecified and a key will be directed to the R console. The actual numbers for the different models will depend on which models are fit and so this key is dynamic.

```
calplot(nested.cox.fit)
```

```
##  specify num for wbeta =
##                       Var num
## null             0.000000   1
## Lasso.1se        3.924673   2
## lasso.min        5.284657   3
## lassoR.1se       4.752238   4
## lassoR.min       6.308573   5
## lassoR0.1se      5.399642   6
## lassoR0.min      6.939080   7
## ridge            3.316922   8
## Lasso.1se cal    6.476655   9
## lasso.min cal    7.333155  10
## lassoR.1se cal   6.365883  11
```

```
## lassoR.min cal  6.960332  12
## lassoR0.1se cal 5.399642  13
## lassoR0.min cal 6.939080  14
## ridge cal       7.550848  15
## step.df         7.267784  16
## step.p          7.273054  17
```

From this key we read of the numers corresponding to the respective models. The variance in X*beta for the "null" model is 0 as the intercept for the Cox model is arbitrarily assinged the value of 0 for each resample model fit. From this key we see we can produce a calibration plot for the ridge regression model by setting wbeta = 8 (wbeta for "which beta"), as in

```
calplot(nested.cox.fit, 8)
```

```
##   Range of X*Beta for calibration:
## -8.186478 7.32873
##   Range of calibrated confidence intervals:
## -16.24789 14.06614
```

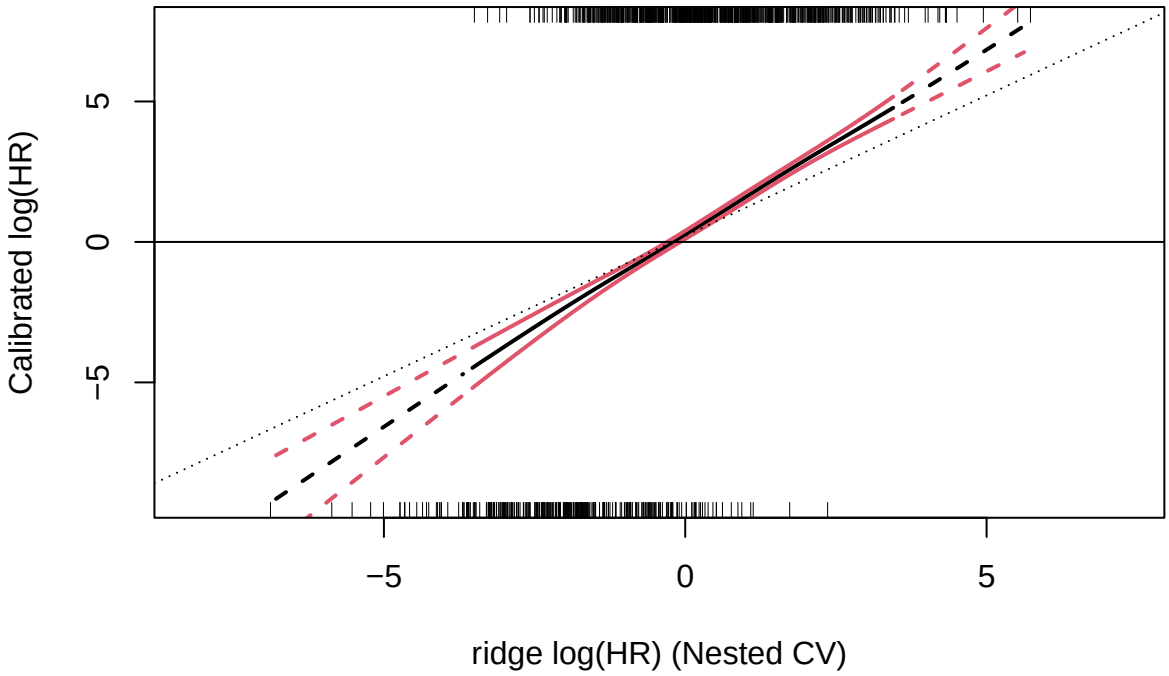

Here we see the model is not ideally calibrated as the calibration curve largely does not include the identity line, and it requires a correction to achieve an un (less) biased estimation of the hazard ratio. Inspecting the calibration curve for a step wise regression model

```
calplot(nested.cox.fit, 16)
```

```
##   Range of X*Beta for calibration:
```

```
## -3.62831 6.633604
##   Range of calibrated confidence intervals:
## -5.710064 7.688828
```

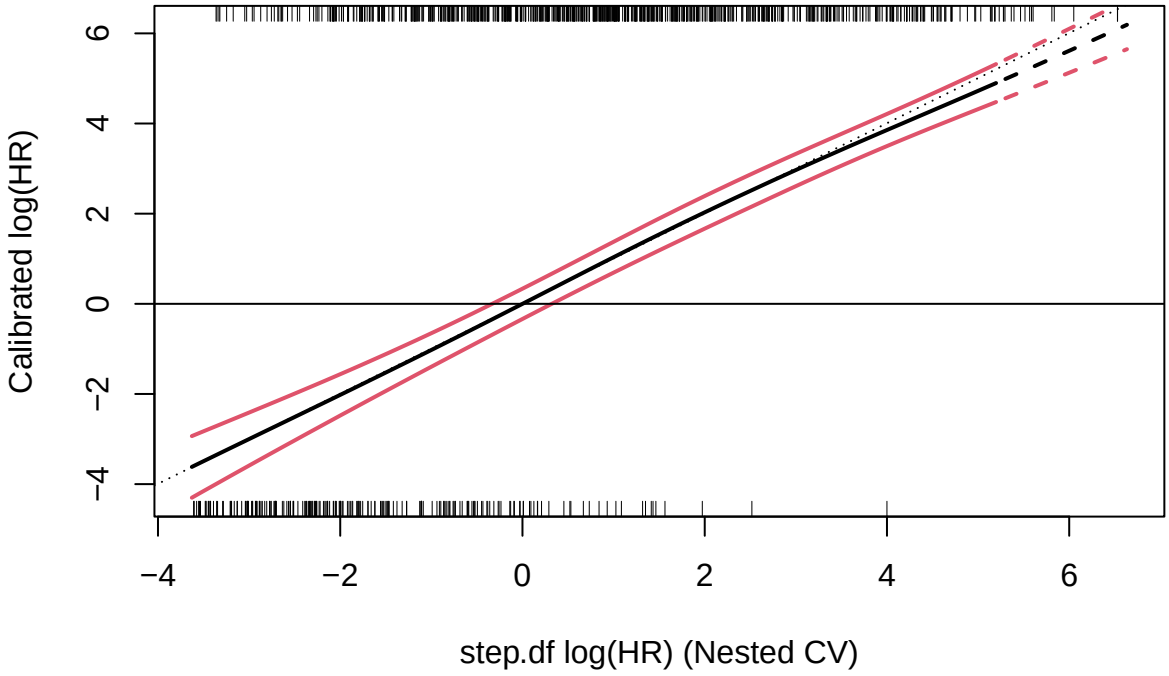


```
##   Due to user specified xlim  93 tick marks are not displayed in rug
##   min:max xb =  -9.82210618760391 9.96096261915667
```

we see that the response is roughly linear but numerically at least there seems to be some correction for over fitting.

To obtain the numerical values used to construct these calibration plots one may specify plot=0 (or plot=2 to plot and obtain the numerical data) in list format as in

```
tmp = calplot(nested.cox.fit, 5, plot=0)
```

```
##   Range of X*Beta for calibration:
## -9.10632 8.827991
##   Range of calibrated confidence intervals:
## -12.05814 11.24302
```

```
str(tmp)
```

```
## List of 5
##  $ estimates     : num [1:101, 1:5] -9.11 -8.93 -8.75 -8.57 -8.39 ...
```

```
##   ..- attr(*, "dimnames")=List of 2
##   .. ..$ : chr [1:101] "1" "2" "3" "4" ...
##   .. ..$ : chr [1:5] "plotxb" "est" "se" "lower" ...
##  $ est.resample  : num [1:10, 1:101] -11.73 -8.67 -9.53 -11.06 -5.63 ...
##  $ se.resample   : num [1:10, 1:101] 4.91 4.08 5.06 4.2 3.53 ...
##  $ lower.resample: num [1:10, 1:101] -21.5 -16.8 -19.7 -19.4 -12.7 ...
##  $ upper.resample: num [1:10, 1:101] -1.909 -0.506 0.601 -2.665 1.429 ...
```

These data may be further processed by the user.

A Binomial Model

For nested.glmnetr() analyses with family = "binomial" with the call

```
yb = simdata$yb
nested.bin.fit = nested.glmnetr(xs,NULL,yb,NULL,family="binomial",
      dolasso=1, doxgb=list(nrounds=250), dorf=1, doorf=1, doann=1,
      folds_n=10, seed=219301029, track=1)
```

an example calibration plot is

```
calplot(nested.bin.fit, 5, plotfold=0)
```

```
##   Range of X*Beta for calibration:
## -8.437476 9.551432
##   Range of calibrated confidence intervals:
## -27.53518 50.62764
```

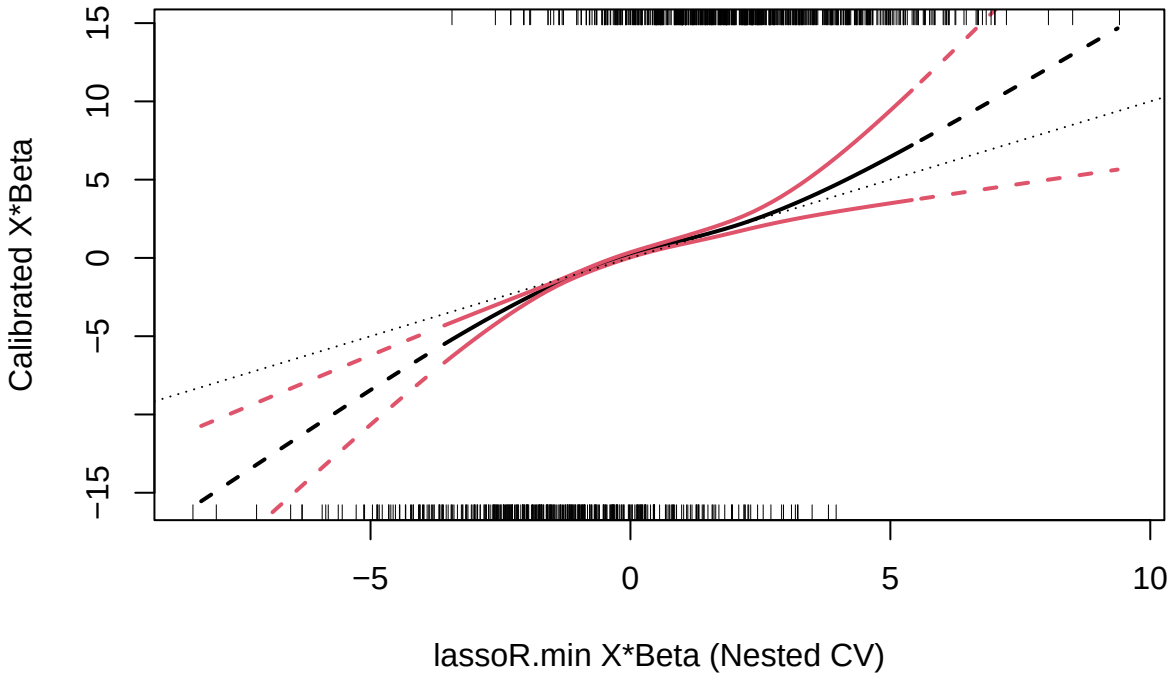


```
##   Due to user specified xlim  1 tick marks are not displayed in rug
##   min:max xb =  -8.41621868029082 10.0536159656975
```

Since these data were generated with probabilities exp(X*beta)/(1+exp(X*beta)) we want that the medal would calibrate linearly. Next we look at the "fully" relaxed model where an unpenalized model is fit based upon the non-zero terms in the fully penalized lasso model.

```
calplot(nested.bin.fit, 7, plotfold=0)
```

```
##   Range of X*Beta for calibration:
## -9.622804 10.61131
##   Range of calibrated confidence intervals:
## -29.67889 49.2557
```

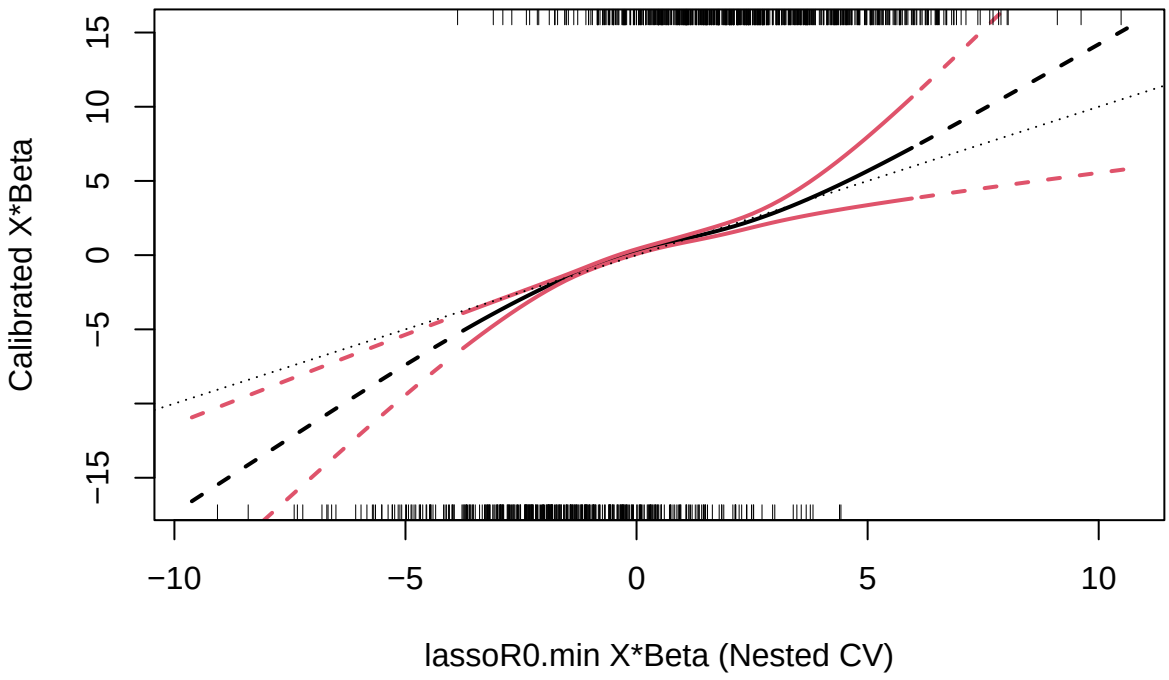


```
##   Due to user specified xlim  2 tick marks are not displayed in rug
##   min:max xb =  -9.75294699449418 11.0684055394026
```

This may calibrate slightly better but is not as linear as we might expect.

A Binomial Model Calibrated Using Bootstrap

Considering this we next fit models using bootstrap, that is fit models based upon random samples from the original sample (with replacement) of size the same as the original sample. Then we fit calibration curves for the out-of-bag sample units for each bootstrap sample, that is the elements of the original sample that are not selected by the bootstrap sample. This is done by specifying the number of bootstrap samples for calculation with the bootstrap option in the calplot() call. First we perform the bootstrap model generation and Out Of Bag (OOB) performance calculations using the nested.glmnetr() function with the *boostrap* option,

```
yb = simdata$yb
nested.bin.boot.fit = nested.glmnetr(xs,NULL,yb,NULL,family="binomial",
      dolasso=1, dorf=1,
      folds_n=10, seed=219301029, track=1, bootstrap=20)
```

with resulting plot

```
calplot(nested.bin.boot.fit, 5, plotfold=1)
```

```
##   Range of X*Beta for calibration:
## -8.437476 9.551432
##   Range of calibrated confidence intervals:
## -18.03857 10.64065
```

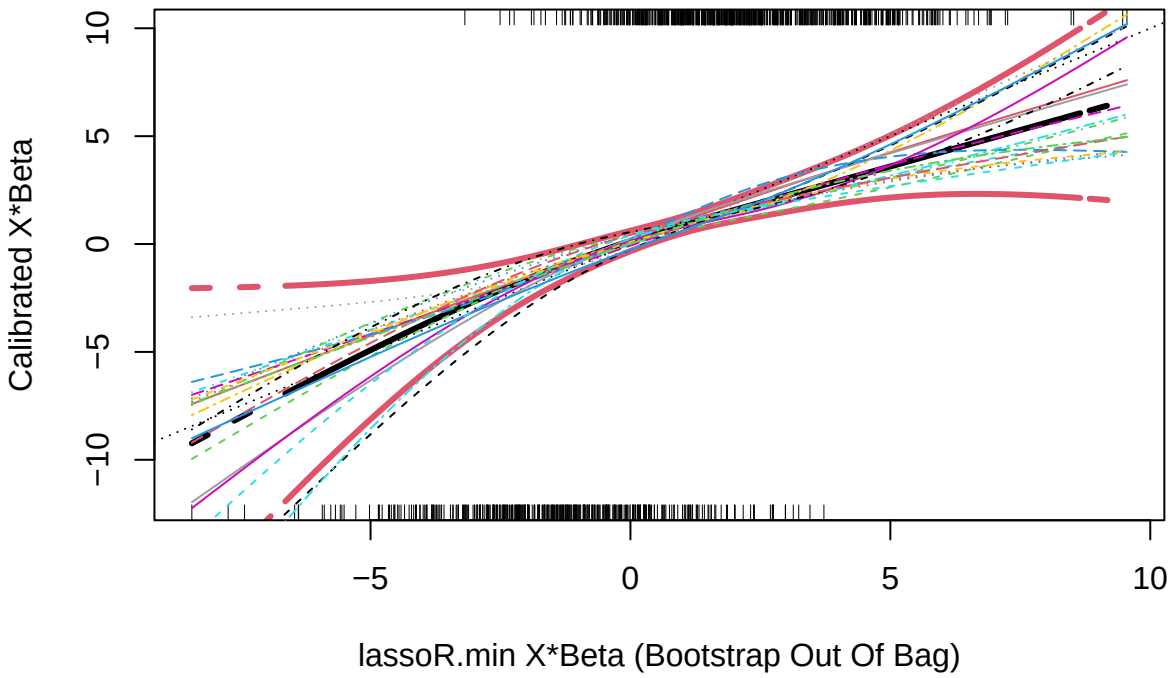

Here we see a similar yet slightly different characteristics between the (nested) cross validation and bootstrap calibration plots. While technically the bootstrap might seem more accurate due to the confidence interval containing the ideal line (which we know applies due ot the way we simulated the data), this may as well be due to the random nature of the re-sample selection process. For the bootstrap one can run a larger number of re-samples to minimize this effect. For the cross validation one may repeat the whole process many times (see the glmnet package vignettes) and then average, but we have not done this here.

Variability in bootstrap in-bag calibrations

In earlier works Austin PC, Steyerberg EW. Graphical assessment of internal and external calibration of logistic regression models by using loess smoothers. Stat Med. 2014; 33(3):517-35. doi: 10.1002/sim.5941. as well as Riley RD, Pate A, Dhiman P, Archer L, Martin GP, Collins GS. Clinical prediction models and the multiverse of madness. BMC Med. 2023; 21(1):502. doi: 10.1186/s12916-023-03212-y fit models based upon bootstrap samples and then fit calibration curves based upon the In Bag data points and model XBeta (predicteds). This follows the usual procedure for bootstrap analysis.

```
calplot(nested.bin.boot.fit, wbeta=5, oob=0, plotfold=1)
```

```
##   Range of X*Beta for calibration:
## -8.437476 9.551432
##   Range of calibrated confidence intervals:
## -39.06028 13.9715
```

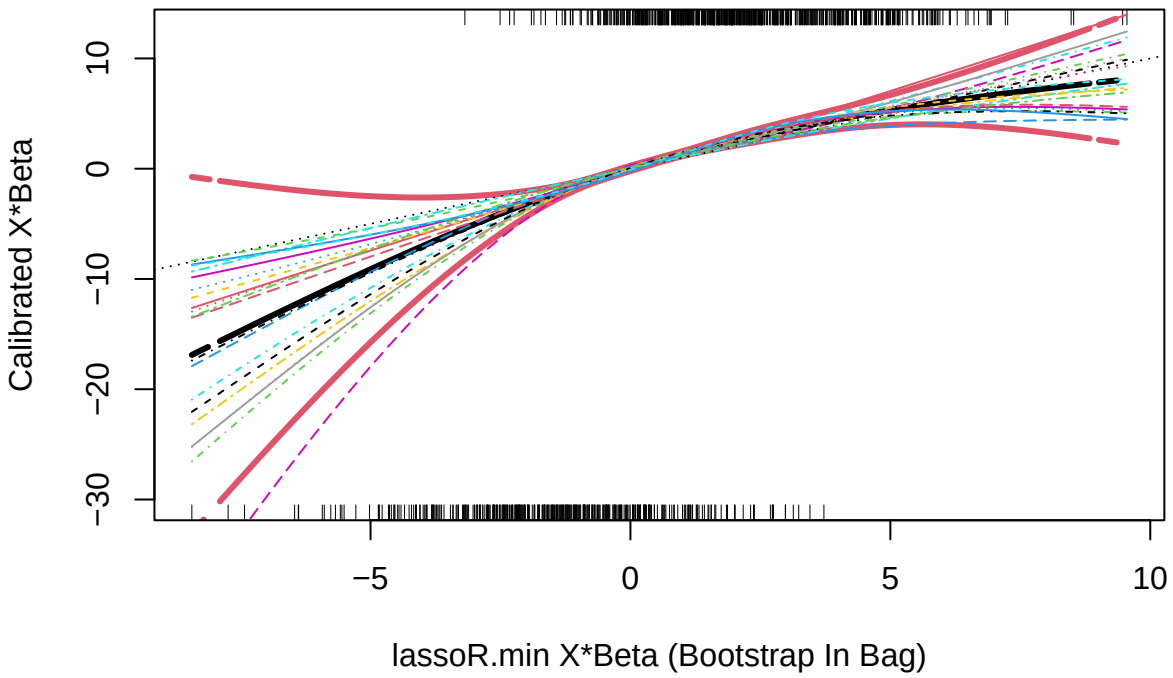

This figure shows the variability expected in the construction of calibration curves as described by the above authors. These curves however make no adjustment for bias correction possible using the bootstrap. Making the usual bias adjustment we have (XBeta_full) - (XBeta_i - XBeta_full) or 2*XBeta_full - XBeta_i. These can be produced with the call

```
calplot(nested.bin.boot.fit, wbeta=5, bootci=1, plotfold=1)
```

```
##   For (bootci == 1) & (oob == 1), oob is set to 0
##   Range of X*Beta for calibration:
## -8.437476 9.551432
##   Range of calibrated confidence intervals:
```

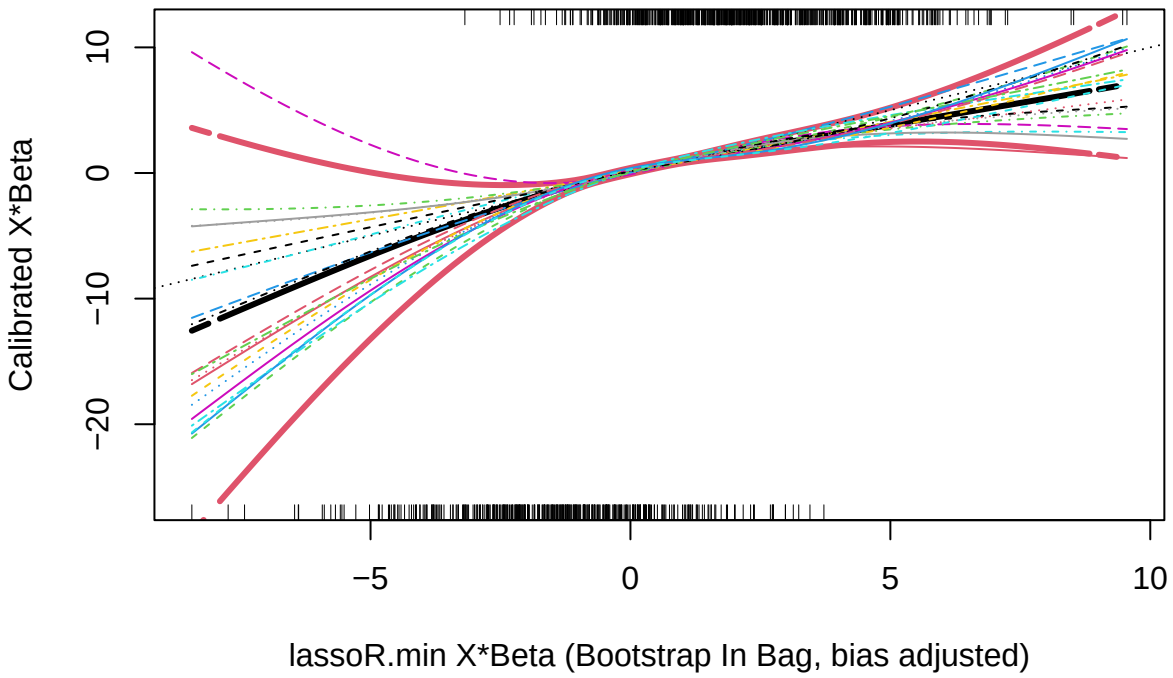

This figure differs from the simple in-bag calibrations. These two figures differ meaningfully from the out-of-bag calibration curves described above.

These two plots were for the relaxed lasso model fit. The corresponding out-of-bag bootstrap calibration plot above for the random forest model shows large variability but confidence limits containing the identity line.

Bootstrap In-Bag Calibration for a Random Forest

The variability in the in-bag bootstrap calibratin curves for the random forest model is depicted in the graph

```
calplot(nested.bin.boot.fit, wbeta=16, bootci=0, oob=0, plotfold=1)
```

```
##   Range of X*Beta for calibration:
## -4.119037 5.517453
##   Range of calibrated confidence intervals:
## -94.68819 115.7647
```

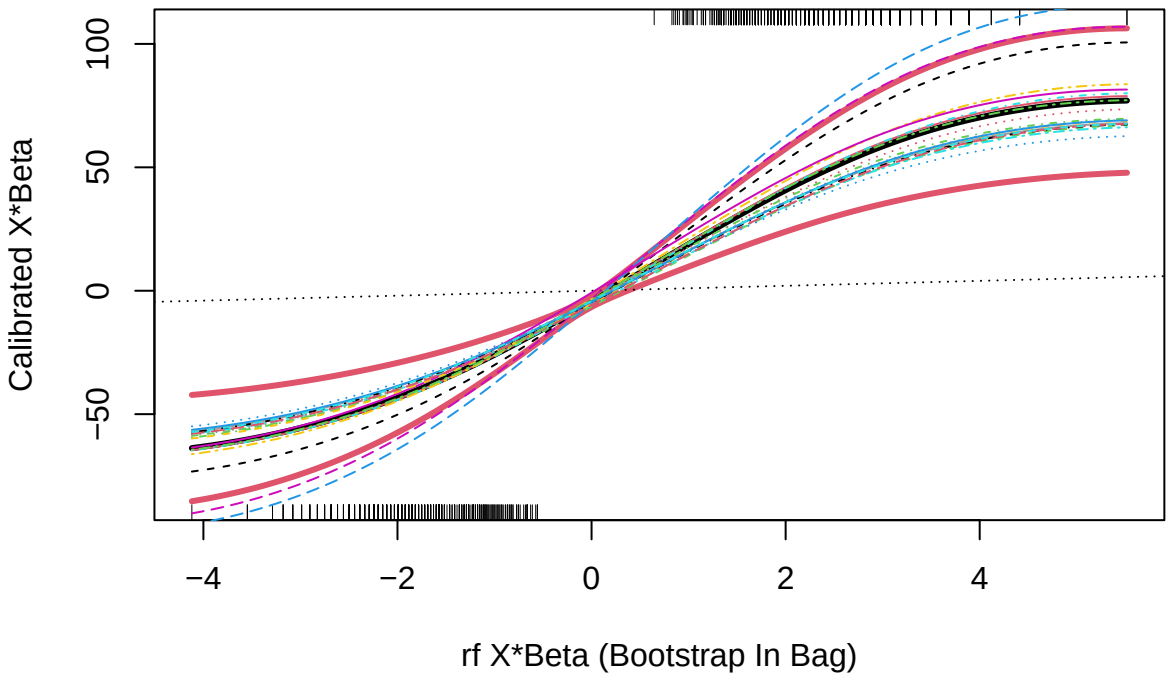

Here we see a strong deviation from the identity line. The bias corrected bootstrap calibration curves

```
calplot(nested.bin.boot.fit, wbeta=16, bootci=1, plotfold=1)
```

```
##   For (bootci == 1) & (oob == 1), oob is set to 0
##   Range of X*Beta for calibration:
## -4.119037 5.517453
##   Range of calibrated confidence intervals:
```

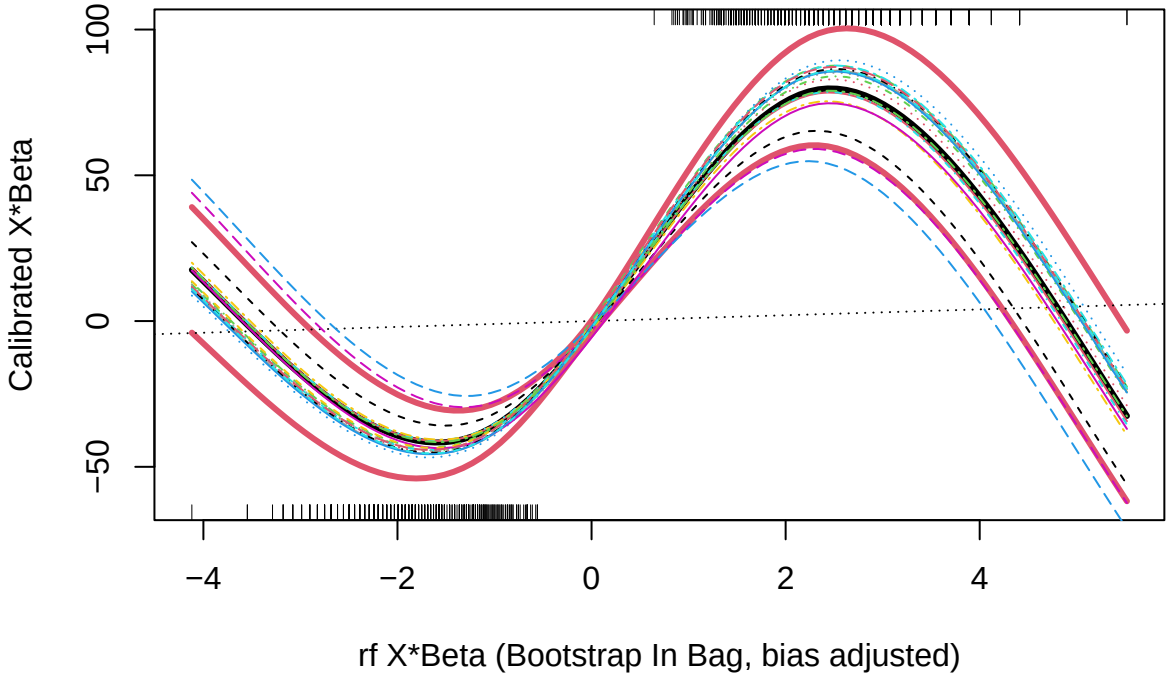

have a more unusual and expected deviation form the ideal calibration curves of the identity line. Returning to the out-of-bag calibration bootstrap plots

```
calplot(nested.bin.boot.fit, wbeta=16, plotfold=1)
```

```
##   Range of X*Beta for calibration:
## -4.119037 5.517453
##   Range of calibrated confidence intervals:
## -53.99674 59.86033
```

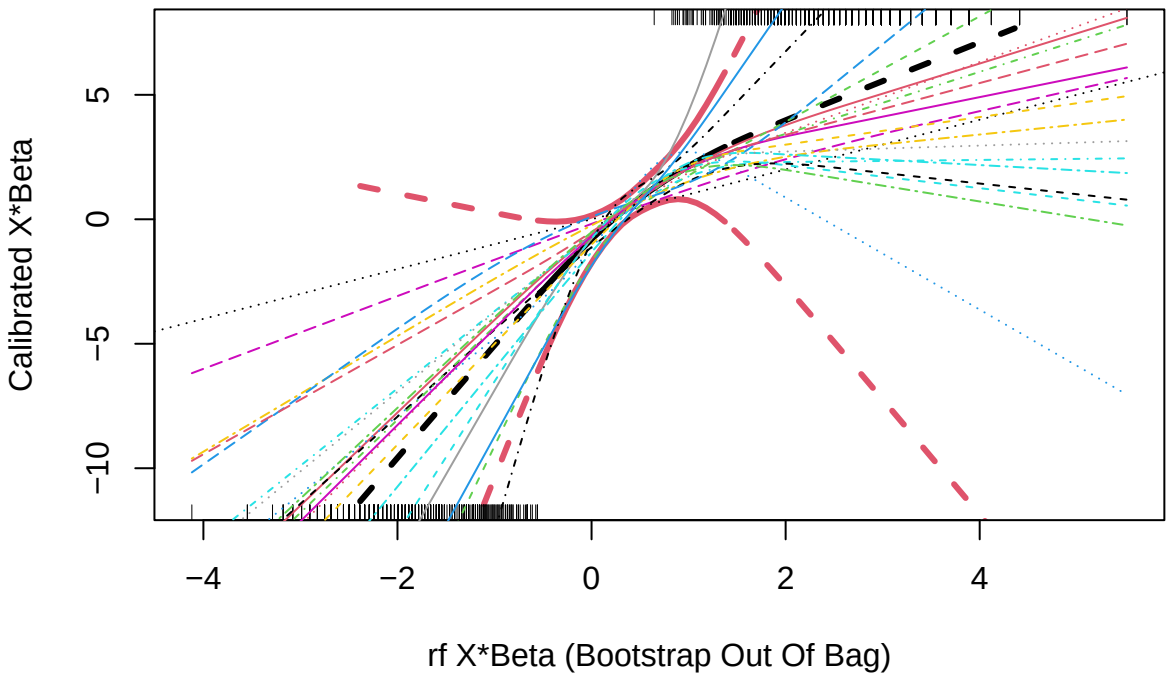

we see that these too are highly varible but more consistent with an appropriately fitting model than suggested by the in-bag plots.

While these two in-bag calibration curve sets strongly depart from the ideal of the identity line, the out-of-bag curves while possibly highly variable were consistent as a group with the ideal of the identity line. The bootstrap out-of-bag calibration curves have a clearer basis and depict more reasonable findings. In a sense this is not unexpected. When fitting machine learning models one typically evaluates model fit based upon the out-of-bag data points and not the in-bag data points. A difficulty here could be that the calibration curves, calculated at particular values for XBeta, do not describe a well defined parameter in terms of the disturbution function. Potentially as well the fitting process for machine learning models might not lend it self to the assumptions typically used to support bootstrap inferences. In particular as we see here when fitting random forest models, with the multiple repeat observations in the bootstrap resamples the model refits are overly optimistic, more strongly overfit and give more extreme XBeta as evidenced in these figures. We present the in-bag and biased adjusted bootstrap calibration curves not to promote their usage but to 1) show the variability in model fit and 2) how they may give poor inferences for a simple dataset. The user may want to generate other datasets of different known form and see how the bootstrap performs for their data.

A Normal (Gaussian Errors) Model

First calculating the numerical summaries of prediction performance

```
nested.gau.fit = nested.glmnetr(xs,NULL,y_,NULL,family="gaussian",
      dolasso=1, doxgb=list(nrounds=250), dorf=1, doorf=1, doann=list(bestof=10),
      folds_n=10, seed=219301029, track=1)
```

and then plotting

```
calplot(nested.gau.fit, wbeta=7)
```

```
##   Range of X*Beta for calibration:
## -8.643322 10.09823
##   Range of calibrated confidence intervals:
## -10.25017 13.60574
```

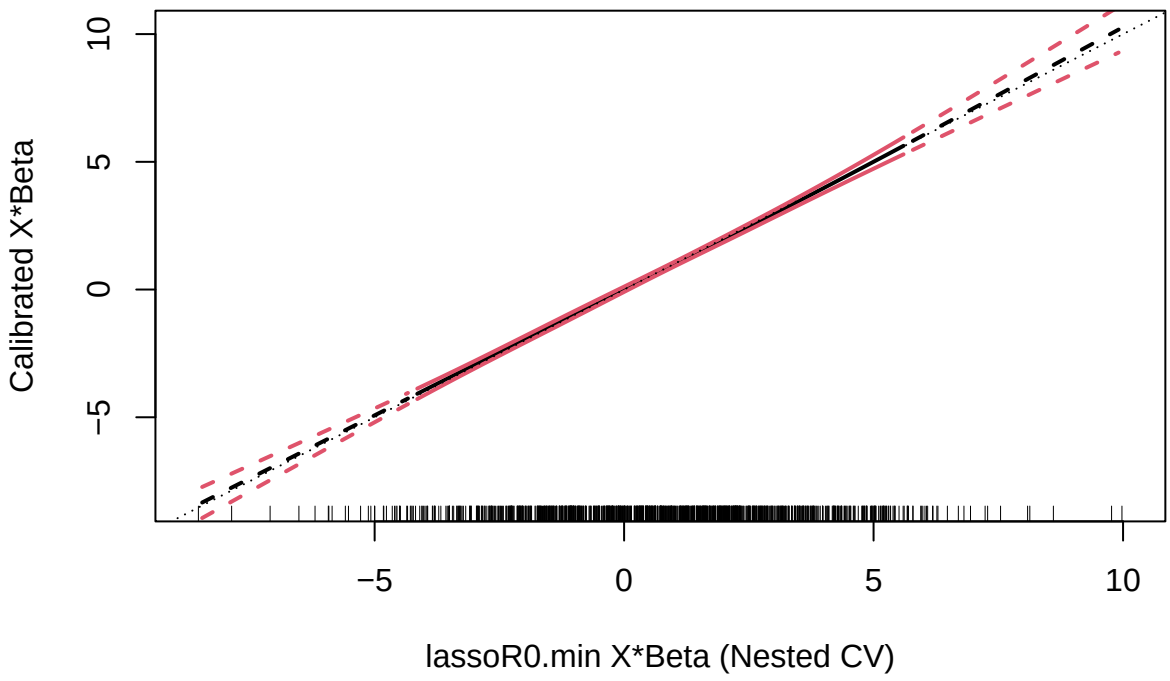

we see a small but probably not significant deviation from the ideal calibration line, the identity function.

Perspective

The summary and calibration plot functions used here do not address all needed for model validation and calibration but do allow a meaningful and un (or minimally) biased summary of model fits. The original outcome variable and X*betas are stored as vectors or matrices in the output with names y_, xbetas (full model) and xbetas.cv (cross-validation) and xbetas.boot.oob and betas.boot.inb (bootstrap in and out of bag) allowing the user to further inspect model fits and to perform other the calibrations. For cross-validation analyses the fold information is contained in the output object in the vector foldid. For bootstrap analyses the first column of xbetas.boot.oob and betas.boot.inb describe the replication of the bootstrap sampling process and the second columns the sequential index for the data point in the input dataset.