6.3.A Prediction of mortality from community-acquired pneumonia

Schuetz et al [1] compared 3 previously derived rules for predicting mortality in patients with community-acquired pneumonia. The 3 rules were the Pneumonia Severity Index (PSI), the CURB65, and the CRB65[1]. They used each of these 3 rules to predict risk of death in 373 patients with community–acquired pneumonia seen in the emergency department of a Swiss university hospital, of whom 41 died within 30 days. Their calibration plots are shown on the next page.

For all 3 rules, the predicted and observed 30-day mortality rates differed substantially. The authors therefore re-calibrated the prediction rules.

a) Figure (c) is the calibration plot for the CRB65 rule. The open diamonds (◊) represent the original risk predictions prior to recalibration. Prior to recalibration, did the CRB65 rule overestimate or underestimate mortality risk? Explain briefly.

The CRB65 underestimated mortality because the graph shows that the observed mortality was higher than the predicted mortality.

b) Figure (b) is the calibration plot for C**U**RB65 (note the letter "U"). CURB 65 assigned only 3 patients to its highest risk group. How many of them died?

The observed mortality looks like about 33%, so of the 3 patients in the highest risk group one must have died.

[1] The CRB65 is just the CURB65 without a lab test called the BUN (blood urea nitrogen).

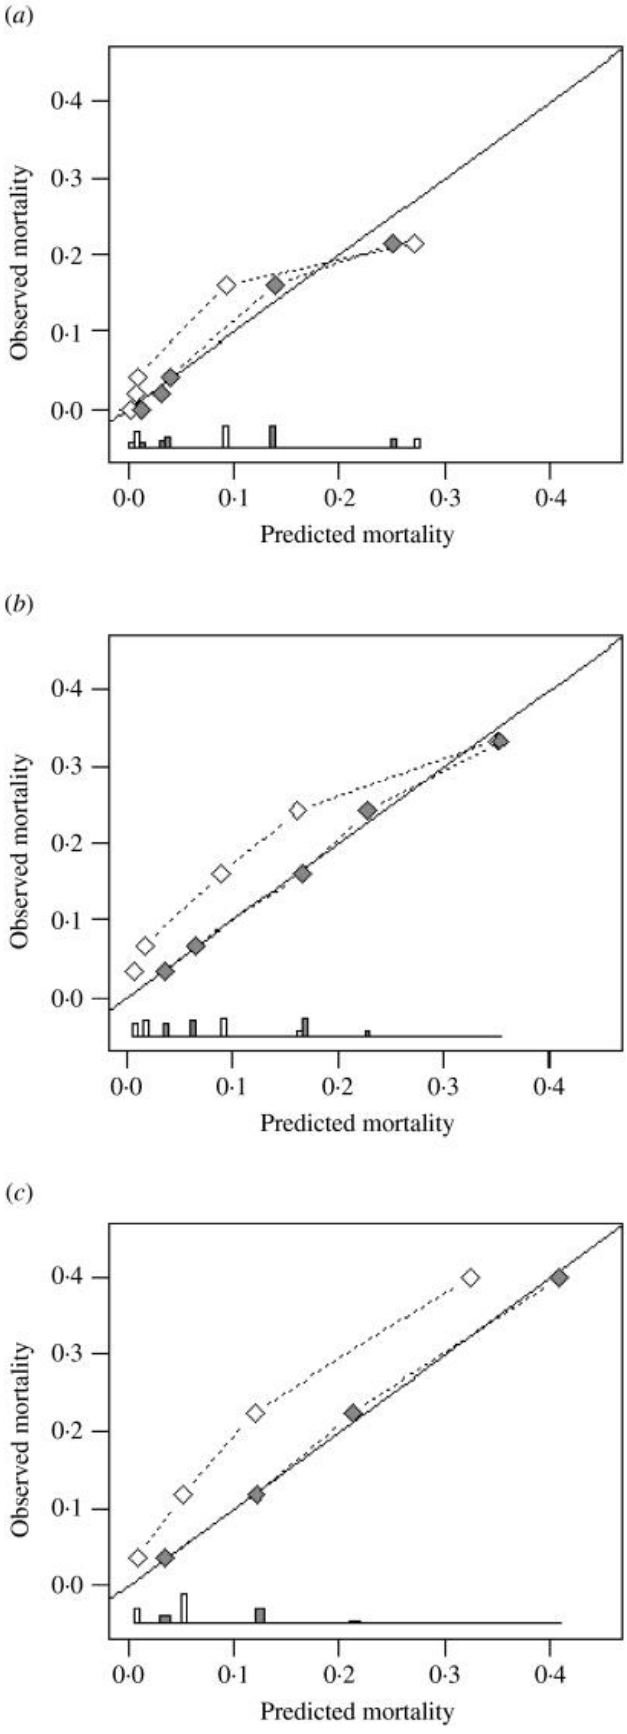

Figure 2: Agreement between predicted and observed 30-day mortality (calibration) for three pneumonia severity prediction rules (*a*) PSI, (*b*) CURB65 and (*c*) CRB65. Observed mortality is plotted according to classes of predicted risk for each prediction rule separately. The solid line of identity represents perfect calibration of predicted risk within new patients. From Schuetz et al[1], used with permission from Cambridge University Press.

◊ Before recalibration
◆ After recalibration

The ROC curves are shown below:

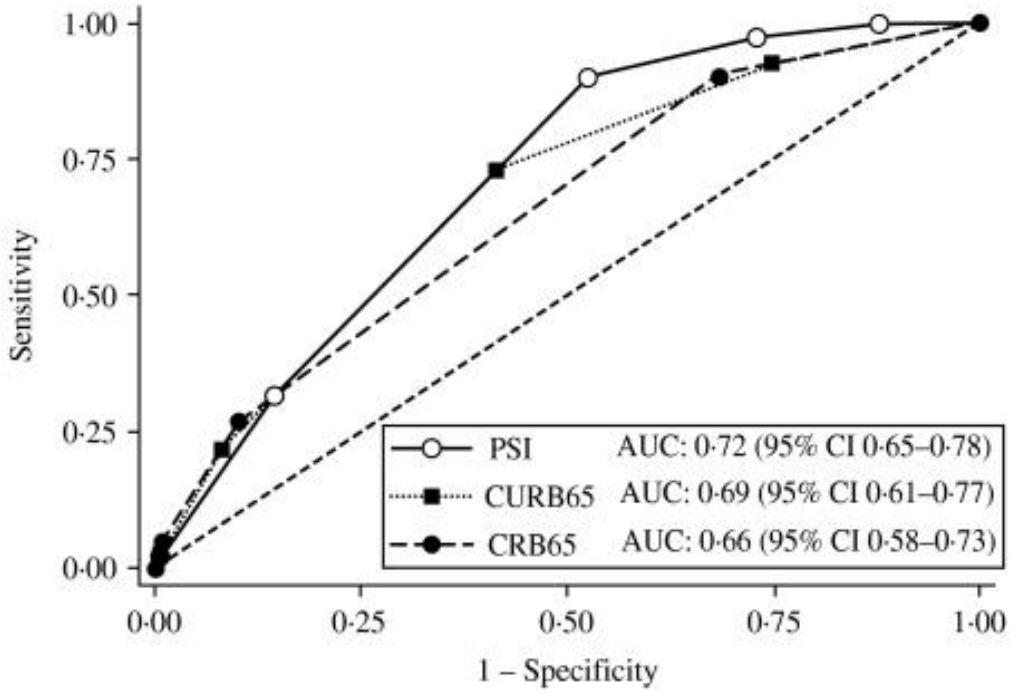

Figure 3. Receiver-operating characteristics analysis for 30-day mortality prediction with three pneumonia severity prediction rules (PSI, CURB65 and CRB65) in 373 patients with community acquired pneumonia. From Schuetz et al[1], used with permission from Cambridge University Press.

c) Do you think the above ROC curves were based on the pre-recalibration or post-recalibration risk predictions? Does it matter? Explain your answer.

It should not matter, because the recalibration would not change the rank order of the predictions, which is what determines the ROC curve.

d) Below is an enlarged version of the calibration plot for the PSI with 3 risk classes circled. Draw a circle around the part of the ROC curve that corresponds to these 3 risk classes.

(*a*)

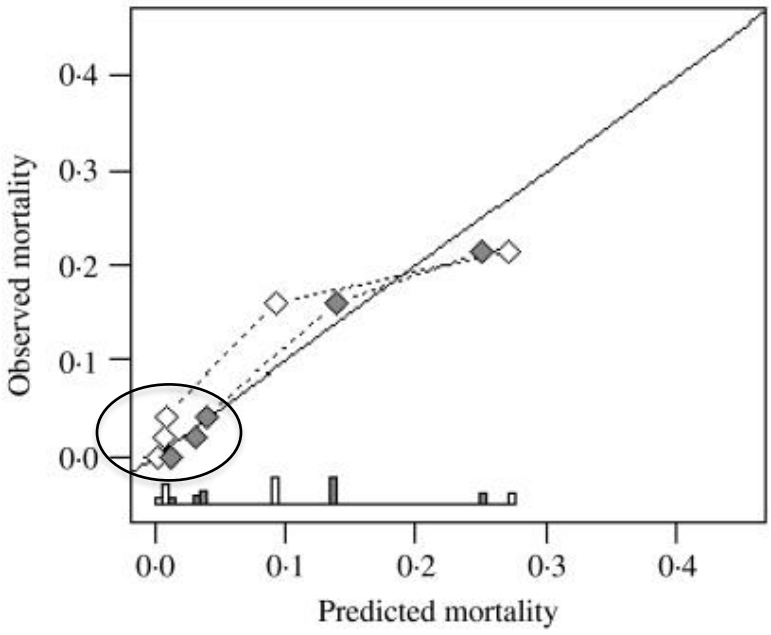
0·4
0·3
0·2
0·1
0·0
Observed mortality
0·0
0·1
0·2
0·3
0·4
Predicted mortality

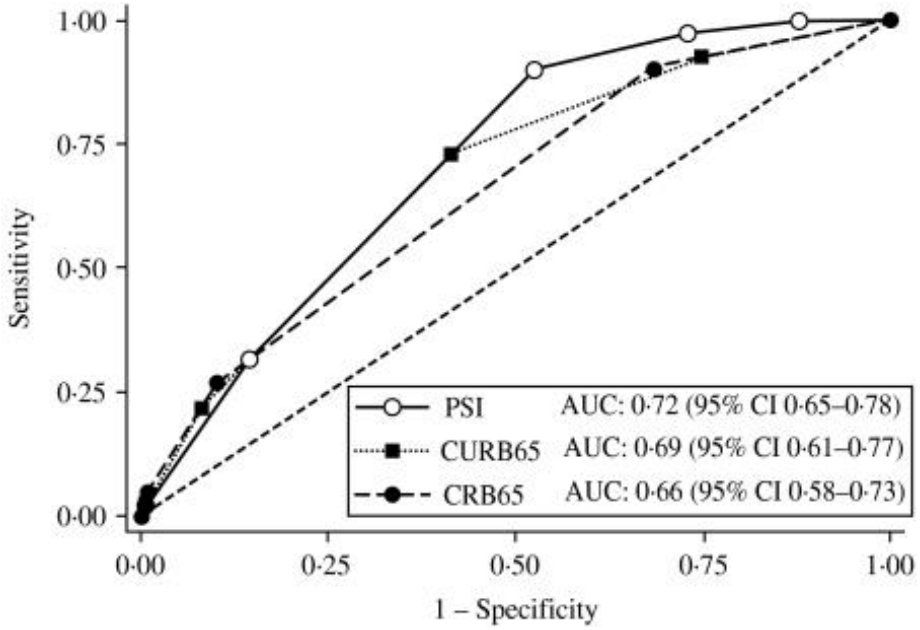
1·00
0·75
0·50
0·25
0·00
Sensitivity
PSI AUC: 0·72 (95% CI 0·65–0·78)
CURB65 AUC: 0·69 (95% CI 0·61–0·77)
CRB65 AUC: 0·66 (95% CI 0·58–0·73)
0·00
0·25
0·50
0·75
1·00
1 – Specificity

This is a hard one. The points in the lower left of the calibration plot are those where the predicted mortality is lowest. So those would be the most reassuring results, i.e., those with the lowest likelihood ratios (LR). The lowest LR are those at the upper right of the ROC curve, where the slopes are closest to zero.

Additional notes, not needed for credit:

Note that each point on the calibration plot corresponds to a group of patients, rather than a cutoff, so points on calibration plots correspond to a segments on the ROC curve.

You can't tell from the calibration plot how many subjects are covered by each point, but you can get a sense of that from the ROC curve: longer line segments means more people. By definition, the vertical distance or "rise" is the proportion of the D+ group in the interval and the horizontal distance or "run" is the proportion of the D- group in the interval. To get the proportion of the entire population in the interval, you have to know the proportion of D+ patients in the sample, P(D+). Then the overall proportion is P(D+)*rise + (1-P(D+))*run.

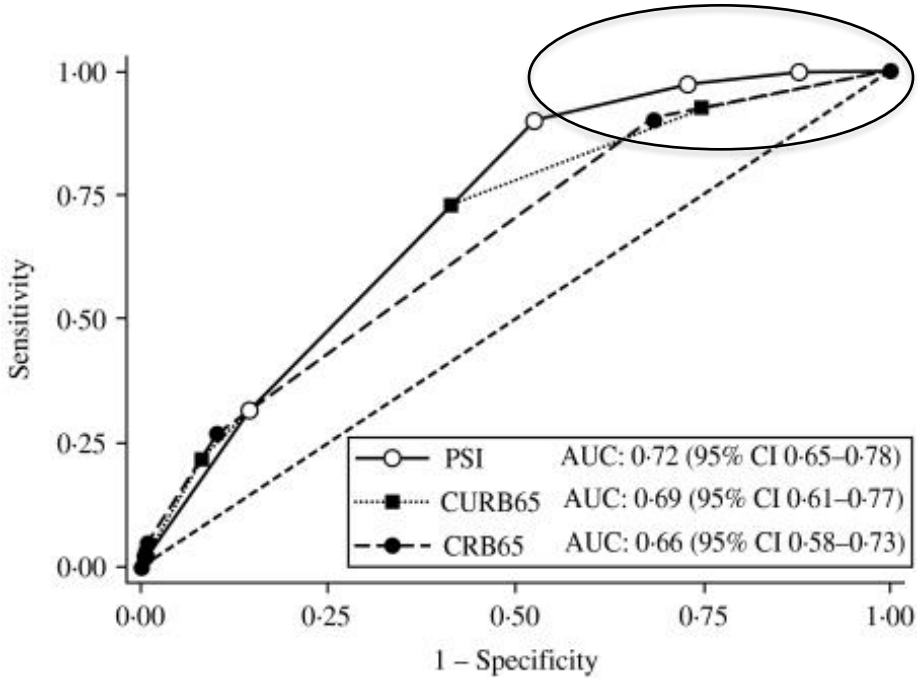

f) The authors were interested in a rule that could identify pneumonia patients at such low risk of death that they could be safely discharged from the emergency department. Even after re-calibration, only one of the 3 rules could identify patients at low enough risk to send home. Which of the 3 rules was it? Explain how you know.

It would have to be the one that was able to achieve the highest sensitivity, the PSI. You can tell this either from the ROC curve (it reaches sensitivity of 100%) or from the calibration plots: only the PSI has a point with zero observed mortality.

1. Schuetz P, Koller M, Christ-Crain M, Steyerberg E, Stolz D, Muller C, et al. Predicting mortality with pneumonia severity scores: importance of model recalibration to local settings. Epidemiol Infect. 2008;136(12):1628-37.