# RevMan-Style Study Markers in SROC Plots

## Overview

Study estimates in a summary receiver operating characteristic (SROC) plot can be represented by symbols whose horizontal and vertical dimensions convey different information. In `dtametaTMB`, the horizontal dimension represents specificity information, whereas the vertical dimension represents sensitivity information.

The package provides several study-symbol sizing methods, including equal sizing, relative sample-size sizing, approximate inverse-standard-error sizing, Fisher-information sizing for applicable models, and two RevMan-style methods:

- `size = "sampsize_revman"`, based on the square roots of the numbers of non-diseased and diseased participants;
- `size = "se_revman"`, based on approximate inverse standard errors of empirical logit specificity and logit sensitivity.

The calculated dimensions are graphical sizing measures. They are not percentage contributions to the precision of the pooled estimate and should not be interpreted as conventional inverse-variance meta-analytic weights. For closed study symbols, `dtametaTMB` applies shape-specific adjustments so that rectangles, ellipses, diamonds, and triangles with the same calculated dimensions have equal enclosed graphical area.

## Notation

For study $i$, let

$$n_{1i} = TP_i + FN_i$$

be the number of diseased participants and

$$n_{0i} = FP_i + TN_i$$

be the number of non-diseased participants. The observed sensitivity and specificity are

$$\widehat{Se}_i = \frac{TP_i}{TP_i + FN_i}$$

and

$$\widehat{Sp}_i = \frac{TN_i}{TN_i + FP_i}.$$

For each sizing method, $q_{Se,i}$ denotes the raw vertical sizing measure and $q_{Sp,i}$ denotes the raw horizontal sizing measure.

## Analysis-specific normalization

Within an analysis, `dtametaTMB` uses one common normalization constant for the horizontal and vertical dimensions and for all studies displayed in the plot.

$$m_i = \max(q_{Se,i}, q_{Sp,i}).$$

For $I$ studies, the analysis-specific constant is

$$c = \frac{9}{\frac{1}{I}\sum_{i=1}^{I} m_i} = \frac{9I}{\sum_{i=1}^{I} \max(q_{Se,i}, q_{Sp,i})}.$$

The unrescaled full horizontal and vertical dimensions are then

$$d_{Sp,i} = 3 + c \cdot q_{Sp,i}$$

and

$$d_{Se,i} = 3 + c \cdot q_{Se,i}.$$

This normalization has two useful properties:

1. the average larger variable-dimension component, $c \cdot m_i$ equals 9;
2. the fixed baseline of 3 keeps a study marker visible when its raw sizing measure is zero in one direction.

### RevMan-style sample-size scaling

For `size = "sampsize_revman"`, the raw measures are

$$q_{Se,i} = \sqrt{n_{1i}}$$

and

$$q_{Sp,i} = \sqrt{n_{0i}}.$$

The analysis-specific sample-size constant is therefore

$$c_{SS} = \frac{9}{\frac{1}{I}\sum_{i=1}^{I} \max\left(\sqrt{n_{1i}}, \sqrt{n_{0i}}\right)}.$$

The corresponding full dimensions are

$$d_{Se,i} = 3 + c_{SS} \cdot \sqrt{n_{1i}}$$

and

$$d_{Sp,i} = 3 + c_{SS} \cdot \sqrt{n_{0i}}.$$

Consequently, a study with more diseased than non-diseased participants is taller than it is wide, whereas a study with more non-diseased than diseased participants is wider than it is tall.

For graphical sample-size scaling, if either $n_{1i}$ or $n_{0i}$ is zero, one is added to both quantities before their square roots are calculated. This follows the RevMan calculation.

### RevMan-style inverse-standard-error scaling

For a binomial proportion $\hat{p}_i$, the large-sample variance of its logit is approximately

$$\mathrm{Var}\left\{\mathrm{logit}(\hat{p}_i)\right\} \approx \frac{1}{n_i \hat{p}_i (1 - \hat{p}_i)}.$$

The corresponding approximate inverse standard error is

$$\frac{1}{SE\left\{\text{logit}(\hat{p}_i)\right\}} \approx \sqrt{n_i \hat{p}_i (1-\hat{p}_i)}.$$

For `size = "se_revman"`, the raw measures are therefore

$$q_{Se,i} = \sqrt{n_{1i} \widehat{Se}_i (1 - \widehat{Se}_i)}$$

and

$$q_{Sp,i} = \sqrt{n_{0i} \widehat{Sp}_i (1 - \widehat{Sp}_i)}.$$

When all four cell counts are positive, these raw measures can equivalently be written as

$$q_{Se,i} = \sqrt{\frac{TP_i FN_i}{TP_i + FN_i}} = \frac{1}{\sqrt{\frac{1}{TP_i} + \frac{1}{FN_i}}}$$

and

$$q_{Sp,i} = \sqrt{\frac{TN_i FP_i}{TN_i + FP_i}} = \frac{1}{\sqrt{\frac{1}{TN_i} + \frac{1}{FP_i}}}$$

The analysis-specific constant is

$$c_{SE} = \frac{9}{\frac{1}{I} \sum_{i=1}^{I} \max(q_{Se,i}, q_{Sp,i})}.$$

The full dimensions are

$$d_{Se,i} = 3 + c_{SE} \cdot q_{Se,i}$$

and

$$d_{Sp,i} = 3 + c_{SE} \cdot q_{Sp,i}.$$

For `size = "se_revman"`, if any of the four cell counts in a study is zero, 0.5 is added to all four counts before the raw horizontal and vertical measures are calculated. The correction is used solely for graphical sizing and does not alter the cell counts supplied to the fitted binomial likelihood.

By contrast, `size = "se"` does not apply a continuity correction. A boundary estimate therefore contributes zero to the corresponding unnormalized inverse-standard-error measure.

## Package-specific graphical rescaling

The dimensions above describe the recovered RevMan-style relative geometry. Before the dimensions are passed to the plotting routine, `dtametaTMB` divides them by 10:

$$\tilde{d}_{Sp,i} = \frac{3 + c \cdot q_{Sp,i}}{10}$$

and

$$\tilde{d}_{Se,i} = \frac{3 + c \cdot q_{Se,i}}{10}.$$

The divisor of 10 is a package-specific graphical rescaling constant. It places the symbols in a range compatible with the default graphical scale of 0.02 used by the other sizing methods. It does not alter relative widths, heights, or the ordering of study-symbol dimensions.

The `scale` argument in the plotting methods multiplies all dimensions proportionally. Changing `scale` therefore changes the overall symbol size without changing the relative geometry among studies.

## Symbol geometry and equal-area scaling

The horizontal and vertical dimensions calculated by the selected sizing method define the geometry of a reference rectangle. The horizontal dimension represents specificity-side information, whereas the vertical dimension represents sensitivity-side information.

For ellipses, diamonds, and triangles, `dtametaTMB` applies shape-specific scaling factors so that each closed symbol has the same enclosed area as the corresponding reference rectangle. In particular, the ellipse semiaxes are multiplied by $2/\sqrt{\pi}$, whereas the horizontal and vertical half-extents of diamonds and triangles are multiplied by $\sqrt{2}$.

Because the line-based symbols `"plus"`, `"cross"`, and `"star"` do not enclose a two-dimensional area, they are drawn using the calculated horizontal and vertical extents without an equal-area adjustment.

### A controlled example

The following dataset has sensitivity and specificity equal to 0.8 in every study, while the diseased and non-diseased sample sizes increase across studies.

```
controlled_data <- data.frame(
  study = paste("Study", LETTERS[1:7]),
  TP = c(20, 80, 320, 800, 1600, 4000, 16000),
  FP = c(5, 20, 80, 200, 400, 1000, 4000),
  FN = c(5, 20, 80, 200, 400, 1000, 4000),
  TN = c(20, 80, 320, 800, 1600, 4000, 16000)
)

controlled_data$sens <- with(
  controlled_data,
  TP / (TP + FN)
)

controlled_data$spec <- with(
  controlled_data,
  TN / (TN + FP)
)

controlled_data
#>     study    TP   FP   FN    TN sens spec
#> 1 Study A    20    5    5    20  0.8  0.8
#> 2 Study B    80   20   20    80  0.8  0.8
#> 3 Study C   320   80   80   320  0.8  0.8
#> 4 Study D   800  200  200   800  0.8  0.8
#> 5 Study E  1600  400  400  1600  0.8  0.8
#> 6 Study F  4000 1000 1000  4000  0.8  0.8
#> 7 Study G 16000 4000 4000 16000  0.8  0.8
```

For every study in this example,

$$\widehat{Se}_i = \widehat{Sp}_i = 0.8.$$

The inverse-standard-error measure is therefore

$$\sqrt{n_i(0.8)(0.2)} = 0.4\sqrt{n_i}.$$

The factor 0.4 is common to all raw horizontal and vertical measures. Because the analysis-specific normalization constant absorbs a common multiplicative factor, the sample-size and inverse-standard-error methods produce identical relative marker dimensions in this controlled example.

The following code reproduces the calculations used by the plotting methods.

```
n1 <- with(controlled_data, TP + FN)
n0 <- with(controlled_data, FP + TN)

# Sample-size measures
q_ss_se <- sqrt(n1)
q_ss_sp <- sqrt(n0)
c_ss <- 9 / mean(pmax(q_ss_se, q_ss_sp))

# Inverse-standard-error measures
q_ise_se <- with(
  controlled_data,
  sqrt(n1 * sens * (1 - sens))
)

q_ise_sp <- with(
  controlled_data,
  sqrt(n0 * spec * (1 - spec))
)

c_se <- 9 / mean(pmax(q_ise_se, q_ise_sp))

comparison <- data.frame(
  study = controlled_data$study,
  sample_size_width = (3 + c_ss * q_ss_sp) / 10,
  sample_size_height = (3 + c_ss * q_ss_se) / 10,
  inverse_se_width = (3 + c_se * q_ise_sp) / 10,
  inverse_se_height = (3 + c_se * q_ise_se) / 10
)

comparison
#>     study sample_size_width sample_size_height inverse_se_width
#> 1 Study A         0.3973797          0.3973797        0.3973797
#> 2 Study B         0.4947593          0.4947593        0.4947593
#> 3 Study C         0.6895186          0.6895186        0.6895186
#> 4 Study D         0.9158831          0.9158831        0.9158831
#> 5 Study E         1.1709902          1.1709902        1.1709902
```

```
#> 6 Study F         1.6771564          1.6771564        1.6771564
#> 7 Study G         3.0543128          3.0543128        3.0543128
#>   inverse_se_height
#> 1         0.3973797
#> 2         0.4947593
#> 3         0.6895186
#> 4         0.9158831
#> 5         1.1709902
#> 6         1.6771564
#> 7         3.0543128
```

The two pairs of dimensions are equal in this example, apart from numerical rounding. These dimensions define the reference rectangular geometry. If an ellipse, diamond, or triangle is selected instead, the shape-specific equal-area adjustment is applied during rendering. The underlying study-specific dimensions and their ordering remain unchanged.

## Using the sizing methods

The RevMan-style methods are selected through the `size` argument of the applicable SROC plotting methods.

```
fit <- fitRutterGatsonis(
  data = controlled_data,
  TP = TP,
  FP = FP,
  FN = FN,
  TN = TN,
  study = study
)
#> Warning in fitRutterGatsonis(data = controlled_data, TP = TP, FP = FP, FN = FN,
#> : The Hessian matrix is not positive definite. Parameter estimates and standard
#> errors may be unreliable. Consider checking model specification, data quality,
#> or potential identifiability issues.
fit
#>
#> Rutter & Gatsonis Model
#> -----------------------
#>
#> Number of studies : 7
#> Optimizer         : Converged
#> Hessian           : Not positive definite
```

```
#> Max |grad|         : 0.002677994
#> -2 log likelihood : 93.659 ( df = 5 )
#> AIC               : 103.659
#> BIC               : 103.388
#>
#> Lambda            : 2.775
#> Theta             : 0.053
#> beta              : 0.077
#>
#> Use summary() for parameter estimates.
```

Use square-root sample-size scaling with:

```
plot(fit,
     symbol = "rectangle",
     size = "sampsize_revman",
     scale = 0.02,
     HSROC = FALSE)
```

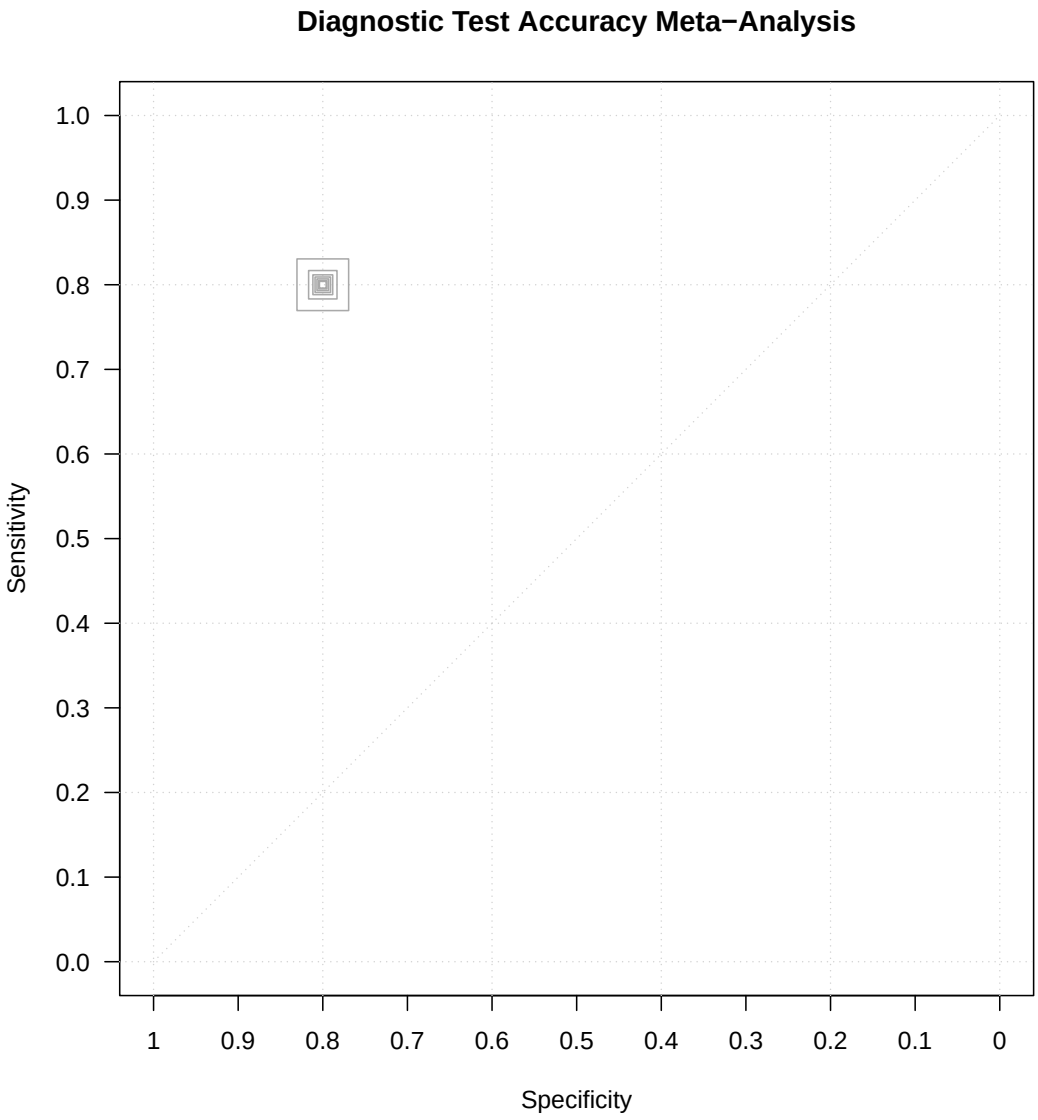


Use inverse-standard-error scaling with:

```
plot(fit,
     symbol = "rectangle",
     size = "se_revman",
     scale = 0.02,
     HSROC = FALSE)
```

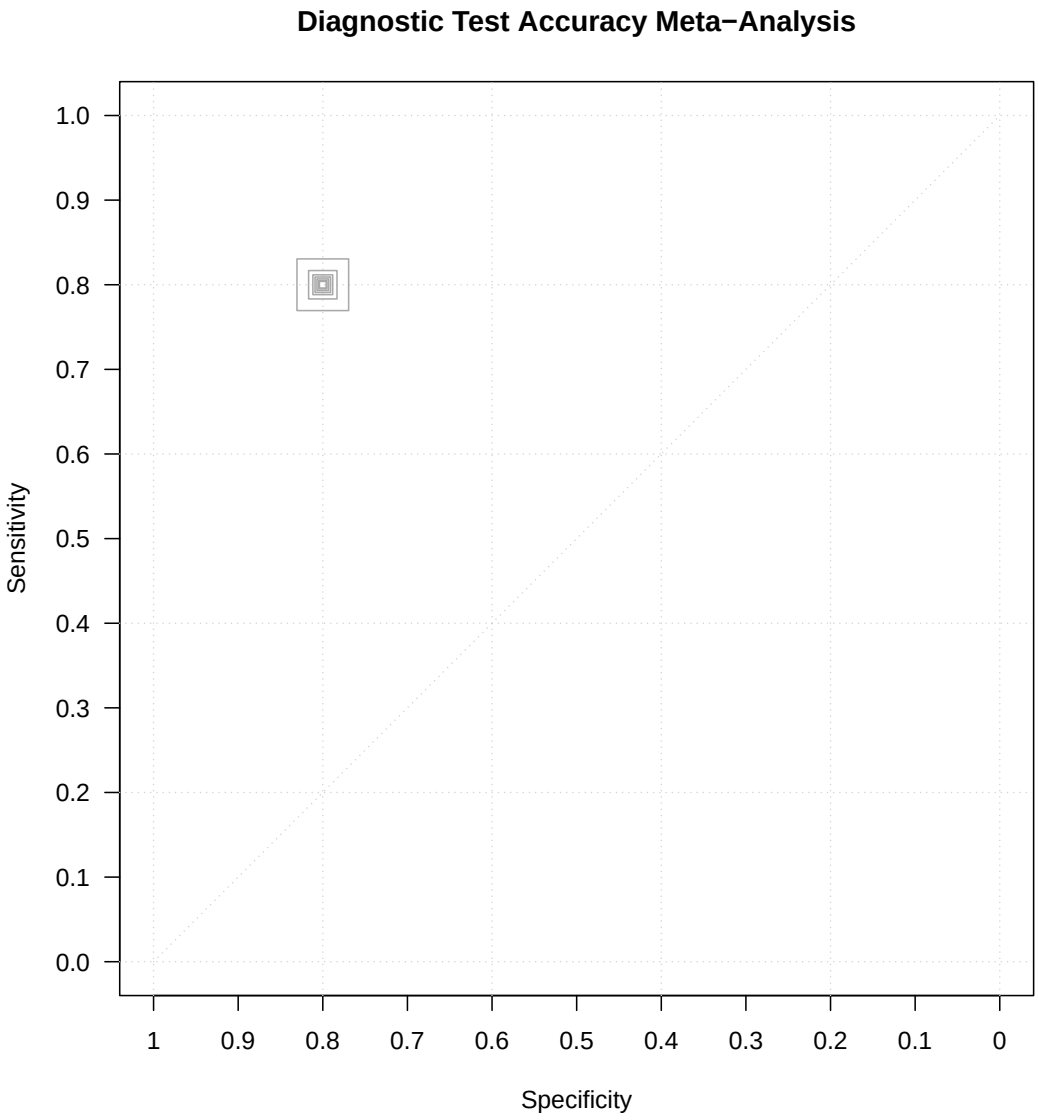


The overall marker size can be adjusted without changing the relative dimensions:

```
plot(fit,
     symbol = "rectangle",
     size = "se_revman",
     scale = 0.02*1.5, # 50% larger than the default
     HSROC = FALSE)
```

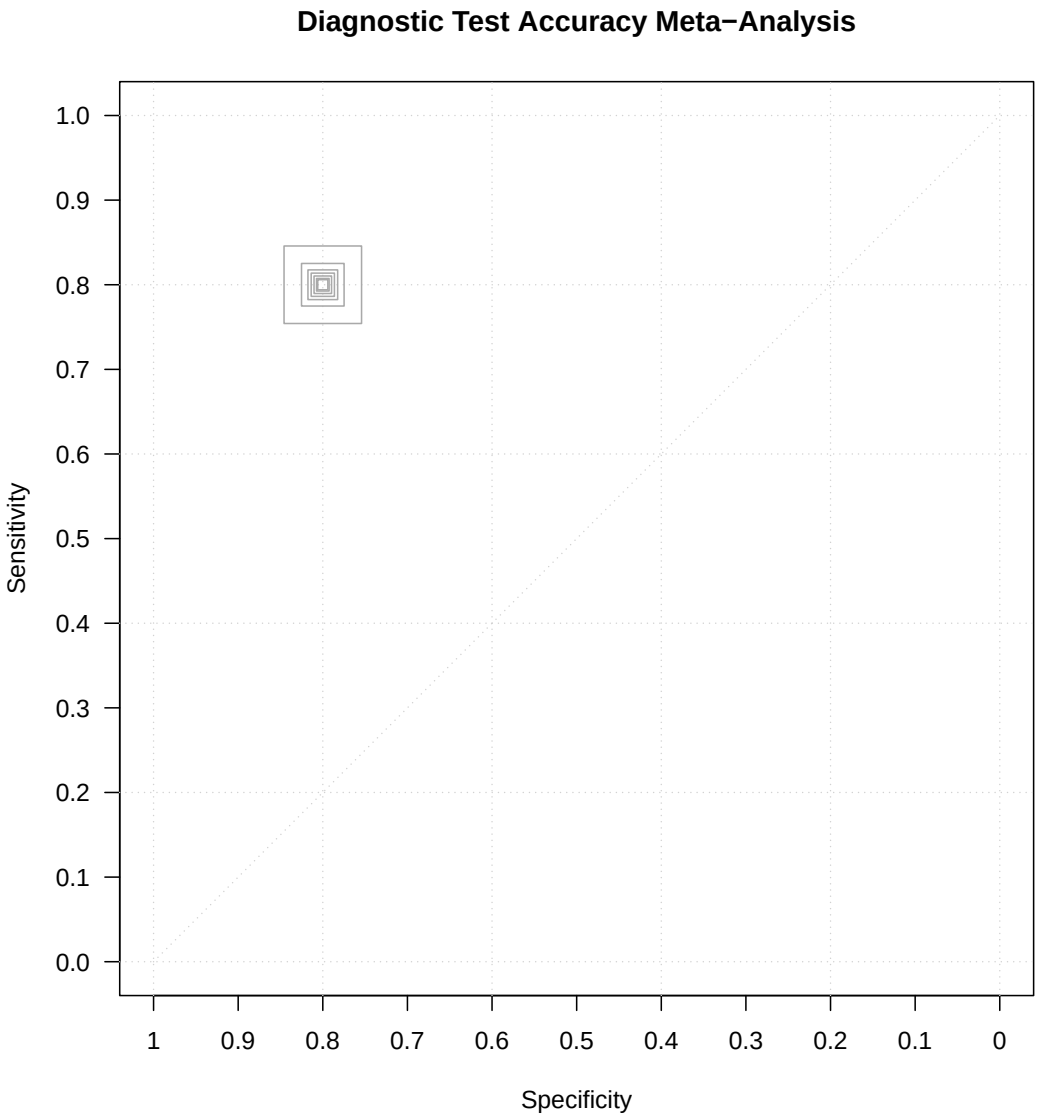


```
plot(fit,
     symbol = "triangle",
     size = "se_revman",
     scale = 0.02*1.5, # 50% larger than the default
     HSROC = FALSE)
```

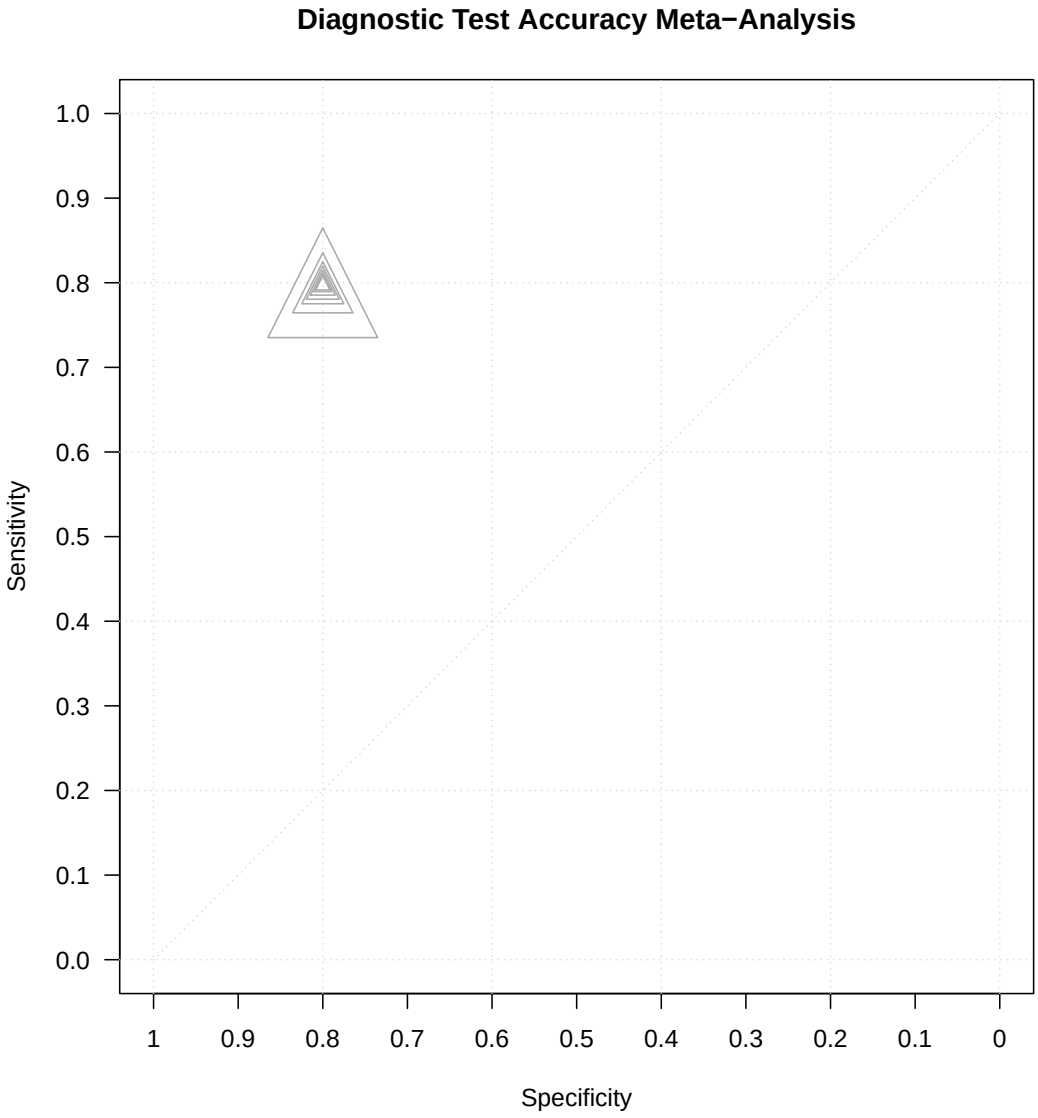


## Relation to the other sizing methods

The original `dtametaTMB` methods answer related but distinct graphical questions.

- `size = "equal"` displays all studies with equal dimensions.
- `size = "sampsize"` uses normalized relative diseased and non-diseased sample sizes directly. For study $i$, the horizontal and vertical sizing measures are

$$w_{Sp,i} = 100 \frac{n_{0i}}{\sum_{j=1}^{I} n_{0j}}$$

and

$$w_{Se,i} = 100 \frac{n_{1i}}{\sum_{j=1}^{I} n_{1j}},$$

where $n_{0i} = FP_i + TN_i$ and $n_{1i} = TP_i + FN_i$. Thus, the horizontal and vertical measures each sum to 100 across studies.

- `size = "se"` uses normalized approximate inverse standard errors on the logit scale. The unnormalized specificity-side and sensitivity-side measures are

$$q_{Sp,i} = \sqrt{n_{0i} \widehat{Sp}_i (1 - \widehat{Sp}_i)}$$

and

$$q_{Se,i} = \sqrt{n_{1i} \widehat{Se}_i (1 - \widehat{Se}_i)}.$$

The displayed horizontal and vertical sizing measures are then

$$w_{Sp,i} = 100 \frac{q_{Sp,i}}{\sum_{j=1}^{I} q_{Sp,j}}$$

and

$$w_{Se,i} = 100 \frac{q_{Se,i}}{\sum_{j=1}^{I} q_{Se,j}}.$$

These measures follow from the large-sample approximation

$$\text{Var}\left\{\text{logit}(\hat{p}_i)\right\} \approx \frac{1}{n_i \hat{p}_i (1 - \hat{p}_i)}.$$

- `size = "fisher"`, where available, uses parameter-specific percentage study contributions derived from a decomposition of the pseudo-linearized Fisher information matrix. Horizontal and vertical dimensions represent contributions toward pooled logit specificity and pooled logit sensitivity, respectively. In subgroup models, these contributions are calculated separately for each subgroup-specific summary parameter and therefore sum to 100 within each subgroup, unless specificity and/or sensitivity are constrained to be common in the subgroups.
- `size = "fisher_revman"`, where available, uses the same parameter-specific percentage Fisher-information contributions as `size = "fisher"`, but applies the fixed-baseline RevMan-style graphical normalization as used in `size = "sampsize_revman"` and `size = "se_revman"`.
- `size = "sampsize_revman"` and `size = "se_revman"` apply the fixed-baseline, analysis-normalized RevMan-style transformations described in this vignette.

The choice of method affects the graphical interpretation of marker dimensions. None of these options changes the fitted meta-analytic model.

## Latent-class extension

In latent-class models, true disease status is unobserved. Separate observed diseased and non-diseased sample sizes are therefore unavailable.

### Total sample-size scaling

For `size = "sampsize_revman"`, the same total-sample-size measure is used in both directions:

$$q_i = \sqrt{n_i}.$$

Because the horizontal and vertical measures are identical, the normalization constant reduces to

$$c_n = \frac{9}{\frac{1}{I}\sum_{i=1}^{I}\sqrt{n_i}}.$$

The resulting horizontal and vertical dimensions are equal

$$\tilde{d}_i = \frac{3 + c_n \cdot \sqrt{n_i}}{10}.$$

If the total sample size $n_i$ is zero, it is replaced by one for graphical sizing before taking the square root. This adjustment affects marker dimensions only.

### Empirical Bayes precision scaling

For `size = "eb_revman"`, let $V_{Se,i}$ and $V_{Sp,i}$ denote the conditional posterior variances of the study-specific empirical Bayes logit-sensitivity and logit-specificity estimates. The raw measures are

$$q_{Se,i} = \frac{1}{\sqrt{V_{Se,i}}}$$

and

$$q_{Sp,i} = \frac{1}{\sqrt{V_{Sp,i}}}.$$

The analysis-specific constant is

$$c_{EB} = \frac{9}{\frac{1}{I}\sum_{i=1}^{I} \max(q_{Se,i}, q_{Sp,i})}.$$

The package-rescaled dimensions are

$$\tilde{d}_{Se,i} = \frac{3 + c_{EB} \cdot q_{Se,i}}{10}$$

and

$$\tilde{d}_{Sp,i} = \frac{3 + c_{EB} \cdot q_{Sp,i}}{10}.$$

The empirical Bayes option applies the same graphical normalization principle to latent-class posterior precision.

## Interpretation and limitations

The final graphical normalization was mainly reconstructed from exported RevMan output rather than from a complete official rendering specification. The reconstructed implementation includes the following features:

- separate horizontal and vertical sizing measures;
- square-root sample-size scaling;
- approximate inverse-logit-standard-error scaling;
- a fixed baseline full dimension of 3;
- one analysis-specific constant shared by both directions;
- normalization based on the mean study-level maximum raw measure;
- equal-area rendering of rectangles, ellipses, diamonds, and triangles through shape-specific graphical adjustments;
- proportional scaling of all marker dimensions through the plotting `scale` argument.

For `size = "se_revman"`, `dtametaTMB` follows the raw inverse-standard-error calculation by adding 0.5 to all four cells when any cell is zero. This adjustment affects marker dimensions only and does not alter the binomial likelihood or any fitted-model result.

For `size = "se"`, no continuity correction is applied. Boundary sensitivity or specificity estimates therefore receive zero unnormalized inverse-standard-error information in the corresponding direction.

For subgroup analyses, `dtametaTMB` calculates one common normalization constant across all studies displayed in the plot. Global normalization preserves comparison across subgroups, so equal marker dimensions correspond to equal raw transformed sample-size or inverse-standard-error measures throughout the figure. This differs from the current RevMan implementation, which calculates the normalization separately within each subgroup.

Physical marker sizes can still differ across plotting devices, export formats, panel dimensions, line widths, and rendering systems. The implementation aims to reproduce the relative marker geometry and normalization principle, not identical device pixels under every graphical configuration.

The formulas should therefore be understood as transparent RevMan-style graphical transformations. They do not alter parameter estimates, confidence regions, prediction regions, or the fitted SROC model.

## Summary

The RevMan-style methods use separate horizontal and vertical sizing measures for specificity and sensitivity. For each analysis, `dtametaTMB` applies

$$c = \frac{9}{\text{mean}_i \left[\max(q_{Se,i}, q_{Sp,i})\right]},$$

followed by the package-rescaled dimensions

$$\tilde{d}_{Se,i} = \frac{3 + cq_{Se,i}}{10} \qquad \text{and} \qquad \tilde{d}_{Sp,i} = \frac{3 + cq_{Sp,i}}{10}.$$

For `size = "sampsize_revman"`, the raw measures are the square roots of the diseased and non-diseased sample sizes. For `size = "se_revman"`, they are approximate inverse standard errors of empirical logit sensitivity and specificity. Graphical adjustments for zero dimensions and zero cells are described above.

The calculated dimensions define a reference rectangle. Ellipses, diamonds, and triangles are rescaled to preserve the same enclosed area, while line-based symbols use the calculated horizontal and vertical extents directly.

The latent-class extensions apply the same normalization principle to total sample size or empirical Bayes posterior precision. All calculations described in this vignette affect symbol rendering only and do not alter the fitted meta-analytic model.