CoTiMA 1.0.4: APA-style tables

29 October 2026

1 APA-style parameter tables (`ctmaTables`)

The function exports conventional CoTiMA and native ctsem matrix-parameter summaries to editable Word (the default), legacy Word XML, LaTeX, or PDF. An in-memory preview is available with `format="none"`. No new model is fitted. Unsupported summary structures, including arbitrary hierarchical summaries or raw random-effect correlation arrays, produce an explicit error rather than being silently relabelled as transformed estimates.

```
# 'fit' must already contain the fitted model.
preview <- ctmaTables(
  fit,
  parameters = c("drift", "asymCINT"),
  statistics = c("estimate", "2.5", "median", "97.5"),
  format = "none"
)
preview$data        # numeric source values
preview$display     # rounded presentation
preview$notes

# Word output needs these optional packages once:
# install.packages(c("officer", "flextable"))
ctmaTables(fit,
  parameters = c("drift[2,1]", "drift[1,2]", "asymCINT"),
  statistics = c("estimate", "sd", "2.5", "median", "97.5"),
  defaultOrder = FALSE,
  activeDirectory = activeDirectory,
  file = "model_parameters",
  format = c("docx", "tex")
)
```

Matrix families may also be supplied without quotation marks: `parameters=c(drift, diffusion)` or `parameters=c(DRIFT, DiFFusion)`. These forms include every entry in the selected matrices. Bare family names are reserved even if objects with those names exist in the calling environment. Quoted selectors and character-vector variables remain supported for programming.

Parameter names are matched exactly. Matrix families and matrix-coordinate selectors are case-insensitive; row denotes outcome and column denotes predictor. Family names such as `drift` select all available entries. `defaultOrder=TRUE` places measurement and initial-state components before mean components, drift cross-effects before auto-effects, then diffusion and other quantities. FALSE preserves the supplied selector order. If parameters or statistics are omitted, an interactive session asks yes/no for every available item. Noninteractive and multiple-model calls require explicit selections.

Source statistics retain their definitions. Posterior SD is not an analytic SE; legacy `n_eff` is not bulk ESS. ESS and Rhat are available only when the selected summary actually sup-

plies them for the selected quantities. Missing requested statistics cause an error by default; `missing="note"` leaves blank cells and explains why. Numeric results retain source precision; display rounding is controlled by `digits`, which defaults to 4. A named `labels` vector can provide publication-friendly row labels. Fixed parameters are not counted by counting rows: the table uses the explicit fit parameter count.

An optional `summary` must agree with the fit's stored or regenerated full parameter table and explicit fit statistics. This rejects mismatched values or likelihoods. It cannot establish the historical origin of an untagged legacy summary. A fingerprint is attached to the returned summary, but a supplied fingerprint does not override value comparisons. Regenerated summaries must use matching options; otherwise omit the supplied summary.

1.1 All primary studies from an initial fit

An object returned by `ctmaInit()` is now handled directly. Estimates are stacked by primary study, with the original study labels retained. Individual ctsem fit objects do not need to have been saved, and no models are refitted.

```
ctmaTables(initFit,
  parameters = c(drift, diffusion),
  statistics = c("estimate", "sd", "2.5", "97.5"),
  activeDirectory = activeDirectory,
  file = "primary_study_estimates",
  format = "docx")
```

The aggregate summary supplies drift estimates and uncertainty. Other selected quantities use retained summary entries or saved coefficient lists where available. Diffusion values are reconciled to fitted time units using the recorded multiplier; rounding of the stored original-time summary can limit recovered uncertainty precision. Study citations are retained in `data$Source` and included in the notes when provided. Optional `modelNames` replaces the primary-study display labels.

Legacy initial summaries do not retain medians, ESS or Rhat; these are not invented. If the native fits are saved in `studyFitList`, missing quantities can instead be obtained by summarizing those fitted objects. `missing="note"` leaves unavailable statistics blank, while the default raises an error. Uncertainty columns historically labelled SE by `ctmaInit()` originate from sd/Sd summaries and are reported as SD. Some retained T0 covariance coefficients lack a saved SD. The reported $-2LL$ and parameter counts are totals across the separately fitted primary studies, not statistics for one pooled model. For missing quantities, the default `recomputeSummary="ask"` asks once before computing summaries for all saved primary fits. No model is refitted. Use `recomputeSummary="yes"` in batch scripts or `"no"` to keep only existing results. Legacy initial objects without saved native fits still require external joint draws for DT uncertainty.

1.2 Discrete-time effects and model comparisons

```
# Lags are expressed in the fitted model's current time units.
# Native Julia transformed drift draws are recognized automatically.
ctmaTables(fit, parameters = "drift",
  statistics = c("estimate", "sd", "2.5", "median", "97.5"),
  timeIntervals = c(1, 6, 12),
  recomputeSummary = "yes", # Permit summaries of saved fits; never
     refit
  timeUnit = "months", format = "none")

ctmaTables(list(Reference = fit1, Alternative = fit2),
```

```
  parameters = "drift", statistics = c("estimate", "2.5", "97.5"),
  format = "docx", activeDirectory = activeDirectory,
  file = "model_comparison")
```

DT uncertainty is obtained by exponentiating each complete joint drift draw, then summarizing the transformed distribution. CT interval endpoints are never exponentiated separately. Missing DT summaries can be obtained through `summary(saved_fit, timeinterval=lag)` for each requested lag. This delegates raw-to-model parameter transformations to ctsem, using the retained joint draws and model specification. Typically `stanfit$rawest` is a single raw point estimate; `stanfit$rawposterior` contains joint raw draws. A point estimate alone cannot supply DT uncertainty. The same summary fallback applies to conventional CoTiMA and native ctsem fits. Alternatively supply `driftDraws` with dimensions draws by outcomes by predictors; for multiple models supply one array per model. These external arrays must correspond to the fitted matrix. If joint draws are unavailable, `dtPointOnly=TRUE` explicitly permits a plug-in point estimate and leaves unavailable DT uncertainty blank. State-dependent drift requires model-specific calculations. DT summaries concern the chosen drift representation and are not averages over individual or study-specific trajectories. Time labels do not rescale time automatically.

Comparisons align matrix coordinates. Users must ensure that these coordinates refer to the same substantive variables and time units in every model. Different parameter names can be handled by using matrix-coordinate selectors. Blank cells identify entries absent from a model; they are not replaced by zero.

1.3 Regular Stan fits and retained draws

The table function supports regular ctsem Stan fits, both directly and inside CoTiMA fit objects. Julia and the experimental ctsem branch are not needed for this path. A pooled CoTiMA model stores a single native fit directly in `studyFitList`; an initial fit stores a list of primary fits. Both layouts are supported.

```
CoTiMA::ctmaTables(fit,
  parameters = drift,
  statistics = c("mean", "sd", "2.5", "median", "97.5"),
  timeIntervals = c(1, 5),
  recomputeSummary = "yes", format = "none")
```

Saved `pop_DRIFT` arrays from optimized Stan fits and sampled Stan fits are recognized as joint transformed draws. If an optimized fit has lost its transformation cache, retained `stanfit$rawposterior` draws can be transformed through that fit's saved Stan program and standata. This is native parameter transformation, with Kalman filtering disabled; it does not refit the model or replace its point likelihood. The stored `rawest` point vector alone cannot supply uncertainty. Missing raw draws or the original model/data setup must be restored rather than replaced by invented intervals.

Additional summary calculations follow `recomputeSummary`. A native fit without a retained base parameter summary must first be summarized to construct the requested table. Newly generated native summaries use enough precision for the requested display, while a caller-supplied rounded summary retains its own precision. Direct native Stan fits are labelled ML, MAP or Bayesian from their recorded estimation controls. Bayesian and optimization-based uncertainty retain their different interpretations. ESS and Rhat remain selectable only if the source supplies them for the requested transformed quantities; diagnostics on raw parameters are not silently relabelled as diagnostics on transformed effects.

Table exports use `fontSize = 10` points by default. Set, for example, `fontSize = 12` to change titles, headers, cells and notes in Word, XML, TeX and PDF output. This does not change the numeric estimates.

1.4 Page orientation, ordering and compact rows

Page orientation and row arrangement are independent. The default `layout="long"` uses portrait pages; `layout="wide"` uses landscape. This applies to Word, legacy XML, PDF and TeX output. The default remains one parameter and its selected statistics per row (`rowLayout="parameter"`). For initial fits, `sortBy="study"` groups estimates by study, while `sortBy="parameter"` presents each parameter across studies before the next parameter. This ordering option is ignored for other fit types. Parameter ordering within these groups follows `defaultOrder`; use FALSE to preserve selector order.

```
# Portrait table, grouped by parameter across primary studies
ctmaTables(initFit, parameters = drift,
  statistics = c("mean", "2.5", "97.5"),
  sortBy = "parameter", layout = "long", format = "none")

# Landscape table, all selected drift entries in one row per study
ctmaTables(initFit, parameters = drift,
  statistics = c("mean", "2.5", "97.5"),
  rowLayout = "drift", layout = "wide", format = "none")

# All selected parameters in one logical row per study
ctmaTables(initFit, parameters = c(drift, diffusion),
  statistics = c("mean", "sd"),
  rowLayout = "all", layout = "wide", format = "none")
```

With `rowLayout="drift"`, other parameter families retain separate rows; DT effects are grouped separately for each lag. With `rowLayout="all"`, all selected parameters share a row per study, or a single row for a single model or model comparison. Compact columns carry grouped parameter headings and separate statistic headings. These options do not alter numeric values: `data` retains its long schema, and `display` contains the requested presentation. If `rowLayout="all"` leaves only one row per study, `sortBy` has no additional row-ordering effect.

Tables that would otherwise be too wide continue across labelled column panels in the chosen page orientation. Parameter/statistic groups are kept together where possible; no selected column or study is discarded. The returned `panels` component contains these export panels. A very large selection therefore need not fit on one physical page even with one logical row per study.

1.5 Export and table notes

Word output uses optional `officer` and `flextable` packages; PDF requires XeLaTeX, for example from a TeX distribution. XML is explicitly Word 2003 WordprocessingML, not spreadsheet XML. TeX text is escaped, and requested outputs are staged before publication. Existing files are preserved unless `overwrite=TRUE`. Table notes distinguish ML, MAP and posterior sampling when metadata permit; report explicit parameter count and $-2LL$; and describe uncertainty and time units. A log-posterior is never substituted for a log-likelihood. Additional notes can be supplied with `tableNote`.