Package: shrinkTVP (via r-universe)

October 11, 2026

Type Package

Title Efficient Bayesian Inference for Time-Varying Parameter Models with Shrinkage

Version 3.1.2

Description Efficient Markov chain Monte Carlo (MCMC) algorithms for fully Bayesian estimation of time-varying parameter models with shrinkage priors, both dynamic and static. Details on the algorithms used are provided in Bitto and Frühwirth-Schnatter (2019) <doi:10.1016/j.jeconom.2018.11.006> and Cadonna et al. (2020) <doi:10.3390/econometrics8020020> and Knaus and Frühwirth-Schnatter (2023) <doi:10.48550/arXiv.2312.10487>. For details on the package, please see Knaus et al. (2021) <doi:10.18637/jss.v100.i13>. For the multivariate extension, see the 'shrinkTVPVAR' package.

License GPL (>= 2)

Encoding UTF-8

Depends R (>= 3.3.0)

Imports Rcpp, GIGrvg, stochvol (>= 3.0.3), coda, methods, utils, zoo

LinkingTo Rcpp, RcppArmadillo, GIGrvg, RcppProgress, stochvol, RcppGSL

RoxygenNote 7.3.2

Suggests testthat, knitr, rmarkdown, R.rsp

VignetteBuilder R.rsp

NeedsCompilation yes

Author Peter Knaus [aut, cre] (ORCID: <https://orcid.org/0000-0001-6498-7084>), Angela Bitto-Nemling [aut], Annalisa Cadonna [aut] (ORCID: <https://orcid.org/0000-0003-0360-7628>), Sylvia Frühwirth-Schnatter [aut] (ORCID: <https://orcid.org/0000-0003-0516-5552>), Daniel Winkler [ctb], Kemal Dingic [ctb]

Maintainer Peter Knaus <peter.knaus@wu.ac.at>

Repository https://cran.r-universe.dev

Date/Publication 2026-07-13 16:00:02 UTC

RemoteUrl https://github.com/cran/shrinkTVP

RemoteRef HEAD

RemoteSha 05ebefc1004efd93936eceae76bb86556f937749

Contents

eval_pred_dens . . . 2
fitted.shrinkTVP . . . 3
forecast_shrinkTVP . . . 4
LPDS . . . 5
plot.mcmc.tvp . . . 6
plot.shrinkTVP . . . 9
plot.shrinkTVP_forc . . . 10
predict.shrinkTVP . . . 11
print.shrinkTVP . . . 12
residuals.shrinkTVP . . . 13
shrinkDTVP . . . 14
shrinkTVP . . . 22
simTVP . . . 29
updateTVP . . . 31

Index **38**

`eval_pred_dens` *Evaluate the one-step ahead predictive density of a fitted TVP model*

Description

`eval_pred_dens` evaluates the one-step ahead predictive density of a fitted TVP model resulting from a call to shrinkTVP at the points supplied in `x`. For details on the approximation of the one-step ahead predictive density used, see the vignette.

Usage

```
eval_pred_dens(x, mod, data_test, log = FALSE)
```

Arguments

`x`	a real number or a vector of real numbers, taken to be the points at which the predictive density will be evaluated.
`mod`	an object of class `shrinkTVP`, containing the fitted model for which the predictive density should be evaluated.
`data_test`	a data frame with one row, containing the one-step ahead covariates. The names of the covariates have to match the names of the covariates used during model estimation in the call to `shrinkTVP`.

log — a single logical value determining whether the density should be evaluated on the log scale or not.

Value

The value returned is a vector of length `length(x)`, containing the values of the predictive density evaluated at the points supplied in `x`.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: `LPDS()`, `fitted.shrinkTVP()`, `forecast_shrinkTVP()`, `predict.shrinkTVP()`, `residuals.shrinkTVP()`

Examples

```
# Simulate data
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

# Estimate model
res <- shrinkTVP(y ~ x1 + x2, data = data[1:199, ])

# Create sequence of x values where the density is to be evaluated
x_vals <- seq(0, 12, by = 0.1)

# Evaluate density and plot
dens <- eval_pred_dens(x_vals, res, data[200, ])
plot(x_vals, dens, type = "l")

# Add vertical line where true value of the one-step ahead y lies
abline(v = data$y[200])
```

---

fitted.shrinkTVP *Calculate fitted historical values for an estimated TVP model*

Description

Calculates the fitted values for an estimated TVP model, i.e. $X_t'\beta_t$. Note that in contrast to `predict.shrinkTVP` this does not include the error term.

Usage

```
## S3 method for class 'shrinkTVP'
fitted(object, ...)
```

Arguments

object	A shrinkTVP object
...	Currently ignored.

Value

An object of class shrinkTVP_fitted

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: LPDS(), eval_pred_dens(), forecast_shrinkTVP(), predict.shrinkTVP(), residuals.shrinkTVP()

Examples

```
# Generate synthetic data
sim <- simTVP()

# Estimate a model
res <- shrinkTVP(y ~ x1 + x2, sim$data)

# Calculate fitted values
fitted <- fitted(res)

# Visualize
plot(fitted)
lines(sim$data$y, col = "forestgreen")
```

---

forecast_shrinkTVP *Draw from posterior predictive density of a fitted TVP model*

Description

forecast_shrinkTVP draws from the posterior predictive distribution of a fitted TVP model resulting from a call to shrinkTVP.

Usage

```
forecast_shrinkTVP(mod, newdata, n.ahead)
```

Arguments

mod	an object of class shrinkTVP, containing the fitted model.
newdata	a data frame containing the future covariates. The names of the covariates have to match the names used during model estimation in the call to shrinkTVP.
n.ahead	a single, positive integer indicating the forecasting horizon, i.e. how many time-points into the future the posterior predictive distribution should be sampled from. Can not be larger than the number of rows in newdata.

Value

The value returned is a list object of class shrinkTVP_forc containing the samples from the posterior predictive density.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: LPDS(), eval_pred_dens(), fitted.shrinkTVP(), predict.shrinkTVP(), residuals.shrinkTVP()

Examples

```
# Simulate data
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

# Estimate model
res <- shrinkTVP(y ~ x1 + x2, data = data[1:190, ])

# Forecast
forc <- forecast_shrinkTVP(res, data[191:200, ])

# Plot
plot(forc)
```

LPDS *Calculate the Log Predictive Density Score for a fitted TVP model*

Description

LPDS calculates the one-step ahead Log Predictive Density Score (LPDS) of a fitted TVP model resulting from a call to shrinkTVP For details on the approximation of the one-step ahead predictive density used, see the vignette.

Usage

```
LPDS(mod, data_test)
```

Arguments

mod	an object of class shrinkTVP, containing the fitted model for which the LPDS should be calculated.
data_test	a data frame with one row, containing the one-step ahead covariates and response. The names of the covariates and the response have to match the names used during model estimation in the call to shrinkTVP.

Value

A real number equaling the calculated LPDS.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: eval_pred_dens(), fitted.shrinkTVP(), forecast_shrinkTVP(), predict.shrinkTVP(), residuals.shrinkTVP()

Examples

```
# Simulate data
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

# Estimate model
res <- shrinkTVP(y ~ x1 + x2, data = data[1:199, ])

# Calculate LPDS
LPDS(res, data[200,])
```

---

plot.mcmc.tvp *Graphical summary of posterior distribution for a time-varying parameter*

Description

plot.mcmc.tvp plots empirical posterior quantiles for a time-varying parameter.

Usage

```
## S3 method for class 'mcmc.tvp'
plot(
  x,
  probs = c(0.025, 0.25, 0.75, 0.975),
  shaded = TRUE,
  quantlines = FALSE,
  shadecol = "skyblue",
  shadealpha = 0.5,
  quantlty = 2,
  quantcol = "black",
  quantlwd = 0.5,
  drawzero = TRUE,
  zerolty = 2,
  zerolwd = 1,
  zerocol = "grey",
  ...
)
```

Arguments

x	mcmc.tvp object
probs	vector of boundaries for credible intervals to plot for each point in time, with values in [0,1]. The largest and smallest value form the outermost credible interval, the second smallest and second largest the second outermost and so forth. The default value is c(0.025, 0.25, 0.75, 0.975). Note that there have to be the same number of probs < 0.5 as there are > 0.5.
shaded	single logical value or a vector of logical values, indicating whether or not to shade the area between the pointwise credible intervals. If a vector is given, the first value given is used to determine if the area between the outermost credible interval is shaded, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of quantile pairs. The default value is TRUE.
quantlines	single logical value or a vector of logical values, indicating whether or not to draw borders along the pointwise credible intervals. If a vector is given, the first value given is used to determine whether the outermost credible interval is marked by lines, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of credible intervals. The default value is FALSE.
shadecol	single character string or a vector of character strings. Determines the color of the shaded areas that represent the credible intervals. If a vector is given, the first color given is used for the outermost area, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of shaded areas. Has no effect if shaded = FALSE. The default value is "skyblue".
shadealpha	real number between 0 and 1 or a vector of real numbers between 0 and 1. Determines the level of transparency of the shaded areas that represent the credible

intervals. If a vector is given, the first value given is used for the outermost area, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of shaded areas. Has no effect if `shaded = FALSE`. The default value is `0.5`.

`quantlty` either a single integer in [0,6] or one of the character strings `"blank"`,`"solid"`, `"dashed"`, `"dotted"`, `"dotdash"`, `"longdash"`, `"twodash"` or a vector containing these. Determines the line type of the borders drawn around the shaded areas that represent the credible intervals. Note that if a vector is supplied the elements have to either be all integers or all character strings. If a vector is given, the first value given is used for the outermost area, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of shaded areas. Has no effect if `quantlines = FALSE`. The default value is `2`.

`quantcol` single character string or a vector of character strings. Determines the color of the borders drawn around the shaded areas that represent the credible intervals. If a vector is given, the first color given is used for borders of the outermost area, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of shaded areas. Has no effect if `quantlines = FALSE`. The default value is `"black"`.

`quantlwd` single real, positive number or a vector of real, positive numbers. Determines the line width of the borders drawn around the shaded areas that represent the credible intervals. If a vector is given, the first number given is used for the borders of the outermost area, the second for the second outermost and so forth. Recycled in the usual fashion if the vector is shorter than the number of shaded areas. Has no effect if `quantlines = FALSE`. The default value is `0.5`.

`drawzero` single logical value determining whether to draw a horizontal line at zero or not. The default value is `TRUE`.

`zerolty` single integer in [0,6] or one of the character strings `"blank"`,`"solid"`, `"dashed"`, `"dotted"`, `"dotdash"`, `"longdash"`, `"twodash"`. Determines the line type of the horizontal line at zero. Has no effect if `drawzero = FALSE`. The default value is `2`.

`zerolwd` single real, positive number. Determines the line width of the horizontal line at zero. Has no effect if `drawzero = FALSE`. The default value is `1`.

`zerocol` single character string. Determines the color of the horizontal line at zero. Has no effect if `drawzero = FALSE`. The default value is `"grey"`.

`...` further arguments to be passed to `plot`.

Value

Called for its side effects and returns invisibly.

Author(s)

Peter Knaus `<peter.knaus@wu.ac.at>`

See Also

Other plotting functions: `plot.shrinkTVP()`, `plot.shrinkTVP_forc()`

Examples

```
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

res <- shrinkTVP(y ~ x1 + x2, data)
plot(res$beta$beta_x1)
```

`plot.shrinkTVP` *Graphical summary of posterior distribution*

Description

`plot.shrinkTVP` generates plots visualizing the posterior distribution.

Usage

```
## S3 method for class 'shrinkTVP'
plot(
  x,
  pars = c("beta"),
  nplot = 3,
  h_borders = c(0.05, 0.05),
  w_borders = c(0.02, 0.02),
  ...
)
```

Arguments

`x`	a `shrinkTVP` object.
`pars`	a character vector containing the names of the parameters to be visualized. The names have to coincide with the names of the list elements of the `shrinkTVP` object. Throws an error if any element of `pars` does not fulfill this criterium. The default is `c("beta")`.
`nplot`	positive integer that indicates the number of tvp plots to display on a single page before a new page is generated. The default value is 3.
`h_borders`	single real, positive number smaller than 0.5 or a vector containing two such numbers. Determines the relative amount of space (the total amount summing up to 1) left blank on the left and right of the plot, in that order. The default is `c(0.05, 0.05)`.
`w_borders`	single real, positive number smaller than 0.5 or a vector containing two such numbers. Determines the relative amount of space (the total amount summing up to 1) left blank at the top and bottom of the plot, in that order. The default is `c(0.02, 0.02)`.
`...`	further arguments to be passed to the respective plotting functions.

Value

Called for its side effects and returns invisibly.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other plotting functions: plot.mcmc.tvp(), plot.shrinkTVP_forc()

Examples

```
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

output <- shrinkTVP(y ~ x1 + x2, data)
plot(output)


## Will produce an error because 'hello' is not a parameter in the model
## Not run:
plot(output, pars = c("beta", "hello"))

## End(Not run)
```

---

plot.shrinkTVP_forc *Graphical summary of posterior predictive density*

Description

plot.shrinkTVP_forc generates plots visualizing the posterior predictive density generated by forecast_shrinkTVP.

Usage

```
## S3 method for class 'shrinkTVP_forc'
plot(x, showgap = FALSE, ...)
```

Arguments

x	a shrinkTVP_forc object.
showgap	if showgap = FALSE, the gap between the historical observations and the forecasts is removed. The default value is FALSE.
...	further arguments to be passed to plot.

Value

Called for its side effects and returns invisibly.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other plotting functions: plot.mcmc.tvp(), plot.shrinkTVP()

Examples

```
set.seed(123)
sim <- simTVP()

train <- sim$data[1:190, ]
test <- sim$data[191:200, ]

res <- shrinkTVP(y ~ x1 + x2, train)

forecast <- forecast_shrinkTVP(res, test)
plot(forecast)
lines(sim$data$y, col = "forestgreen")
```

---

predict.shrinkTVP *Calculate predicted historical values for an estimated TVP model*

Description

Calculates the predicted past values for an estimated TVP model, i.e. $X_t'\beta_t + \epsilon_t$. Note that in contrast to fitted.shrinkTVP this includes the error term.

Usage

```
## S3 method for class 'shrinkTVP'
predict(object, ...)
```

Arguments

object	a shrinkTVP object
...	Currently ignored.

Value

An object of class shrinkTVP_pred.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: LPDS(), eval_pred_dens(), fitted.shrinkTVP(), forecast_shrinkTVP(), residuals.shrinkTVP()

Examples

```
# Generate synthetic data
sim <- simTVP(N = 300)

# Estimate a model
res <- shrinkTVP(y ~ x1 + x2, sim$data)

# Calculate predicted values
pred <- predict(res)

# Visualize
plot(pred)
lines(sim$data$y, col = "forestgreen")
```

---

print.shrinkTVP *Nicer printing of shrinkTVP objects*

Description

Nicer printing of shrinkTVP objects

Usage

```
## S3 method for class 'shrinkTVP'
print(x, ...)
```

Arguments

x	a shrinkTVP object.
...	Currently ignored.

Value

Called for its side effects and returns invisibly.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

`residuals.shrinkTVP` *Calculate residuals for an estimated TVP model*

Description

Calculates the residuals for an estimated TVP model, i.e. $y_t - X_t'\beta_t$.

Usage

```
## S3 method for class 'shrinkTVP'
residuals(object, ...)
```

Arguments

`object`	a `shrinkTVP` object.
`...`	Currently ignored.

Value

An object of class `shrinkTVP_resid`

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

See Also

Other prediction functions: `LPDS()`, `eval_pred_dens()`, `fitted.shrinkTVP()`, `forecast_shrinkTVP()`, `predict.shrinkTVP()`

Examples

```
# Generate synthetic data
sim <- simTVP(N = 300)

# Estimate a model
res <- shrinkTVP(y ~ x1 + x2, sim$data)

# Calculate residuals
resids <- residuals(res)

# Visualize
plot(resids)
```

shrinkDTVP *Markov Chain Monte Carlo (MCMC) for time-varying parameter models with dynamic shrinkage*

Description

shrinkTVP samples from the joint posterior distribution of the parameters of a time-varying parameter model with dynamic triple gamma shrinkage, potentially including stochastic volatility (SV), and returns the MCMC draws.

Usage

```
shrinkDTVP(
  formula,
  data,
  mod_type = "double",
  niter = 10000,
  nburn = round(niter/2),
  nthin = 1,
  learn_a_xi = TRUE,
  learn_a_tau = TRUE,
  a_xi = 0.1,
  a_tau = 0.1,
  learn_c_xi = TRUE,
  learn_c_tau = TRUE,
  c_xi = 0.1,
  c_tau = 0.1,
  a_eq_c_xi = FALSE,
  a_eq_c_tau = FALSE,
  learn_kappa2_B = TRUE,
  learn_lambda2_B = TRUE,
  kappa2_B = 20,
  lambda2_B = 20,
  a_psi = 0.5,
  c_psi = 0.5,
  iid = FALSE,
  shrink_inter = TRUE,
  hyperprior_param,
  display_progress = TRUE,
  sv = FALSE,
  sv_param,
  MH_tuning,
  starting_vals
)
```

Arguments

formula	object of class "formula": a symbolic representation of the model, as in the function lm. For details, see formula.
data	*optional* data frame containing the response variable and the covariates. If not found in data, the variables are taken from environment(formula), typically the environment from which shrinkTVP is called. No NAs are allowed in the response variable and the covariates.
mod_type	character string that reads either "triple", "double" or "ridge". Determines whether the triple gamma, double gamma or ridge prior are used for theta_sr and beta_mean. The default is "double".
niter	positive integer, indicating the number of MCMC iterations to perform, including the burn-in. Has to be larger than or equal to nburn + 2. The default value is 10000.
nburn	non-negative integer, indicating the number of iterations discarded as burn-in. Has to be smaller than or equal to niter - 2. The default value is round(niter / 2).
nthin	positive integer, indicating the degree of thinning to be performed. Every nthin draw is kept and returned. The default value is 1, implying that every draw is kept.
learn_a_xi	logical value indicating whether to learn a_xi, the spike parameter of the state variances. Ignored if mod_type is set to "ridge". The default value is TRUE.
learn_a_tau	logical value indicating whether to learn a_tau, the spike parameter of the mean of the initial values of the states. Ignored if mod_type is set to "ridge". The default value is TRUE.
a_xi	positive, real number, indicating the (fixed) value for a_xi. Ignored if learn_a_xi is TRUE or mod_type is set to "ridge". The default value is 0.1.
a_tau	positive, real number, indicating the (fixed) value for a_tau. Ignored if learn_a_tau is TRUE or mod_type is set to "ridge". The default value is 0.1.
learn_c_xi	logical value indicating whether to learn c_xi, the tail parameter of the state variances. Ignored if mod_type is not set to "triple" or a_eq_c_xi is set to TRUE. The default value is TRUE.
learn_c_tau	logical value indicating whether to learn c_tau, the tail parameter of the mean of the initial values of the states. Ignored if mod_type is not set to "triple" or a_eq_c_tau is set to TRUE. The default value is TRUE.
c_xi	positive, real number, indicating the (fixed) value for c_xi. Ignored if learn_c_xi is TRUE, mod_type is not set to "triple" or a_eq_c_xi is set to TRUE. The default value is 0.1.
c_tau	positive, real number, indicating the (fixed) value for c_tau. Ignored if learn_c_xi is TRUE, mod_type is not set to "triple" or a_eq_c_tau is set to TRUE. The default value is 0.1.
a_eq_c_xi	logical value indicating whether to force a_xi and c_xi to be equal. If set to TRUE, beta_a_xi and alpha_a_xi are used as the hyperparameters and beta_c_xi and alpha_c_xi are ignored. Ignored if mod_type is not set to "triple". The default value is FALSE.

a_eq_c_tau logical value indicating whether to force a_tau and c_tau to be equal. If set to TRUE, beta_a_tau and alpha_a_tau are used as the hyperparameters and beta_c_tau and alpha_c_tau are ignored. Ignored if mod_type is not set to "triple". The default value is FALSE.

learn_kappa2_B logical value indicating whether to learn kappa2_B, the global level of shrinkage for the state variances. The default value is TRUE.

learn_lambda2_B
logical value indicating whether to learn the lambda2_B parameter, the global level of shrinkage for the mean of the initial values of the states. The default value is TRUE.

kappa2_B positive, real number, indicating the (fixed) value for kappa2_B. Ignored if learn_kappa2_B is TRUE. The default value is 20.

lambda2_B positive, real number, indicating the (fixed) value for lambda2_B. Ignored if learn_lambda2_B is TRUE. The default value is 20.

a_psi positive, real number, or a vector of length equal to the number of covariates containing positive, real numbers. Indicates the value for a_psi, which is the pole parameter of the dynamic triple gamma. The default value is 0.5.

c_psi positive, real number, or a vector of length equal to the number of covariates containing positive, real numbers. Indicates the value for c_psi, which is the tail parameter of the dynamic triple gamma. The default value is 0.5.

iid logical value indicating whether the innovations are assumed to be independent and identically distributed. If set to TRUE, the innovations are assumed to be a priori iid triple gamma. If set to FALSE, the prior on the innovations is the dynamic triple gamma specification of Knaus and Frühwirth-Schnatter (2023). The default value is FALSE.

shrink_inter logical value indicating whether to dynamically shrink the intercept. Note that shrinkage is still applied to the theta_sr and beta_mean associated with the intercept. The intercept column is automatically determined by the function and does not have to be included in the formula. This is done by finding the column that contains only 1s. The default value is TRUE.

hyperprior_param
optional named list containing hyperparameter values. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. All hyperparameter values have to be positive, real numbers. The following hyperparameters can be supplied:

- c0: The default value is 2.5.
- g0: The default value is 5.
- G0: The default value is 5 / (2.5 - 1).
- e1: The default value is 0.001.
- e2: The default value is 0.001.
- d1: The default value is 0.001.
- d2: The default value is 0.001.
- alpha_a_xi: The default value is 5.
- alpha_a_tau: The default value is 5.

- beta_a_xi: The default value is 10.
- beta_a_tau: The default value is 10.
- alpha_c_xi: The default value is 5.
- alpha_c_tau: The default value is 5.
- beta_c_xi: The default value is 2.
- beta_c_tau: The default value is 2.
- a_rho: The default value is 2.
- b_rho: The default value is 0.95.
- alpha_rho: The default value is 0.5.
- beta_rho: The default value is 3.

display_progress

logical value indicating whether the progress bar and other informative output should be displayed. The default value is TRUE.

sv

logical value indicating whether to use stochastic volatility for the error of the observation equation. For details please see stochvol, in particular svsample. The default value is FALSE.

sv_param

optional named list containing hyperparameter values for the stochastic volatility parameters. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. Ignored if sv is FALSE. The following elements can be supplied:

- Bsigma_sv: positive, real number. The default value is 1.
- a0_sv: positive, real number. The default value is 5.
- b0_sv: positive, real number. The default value is 1.5.
- bmu: real number. The default value is 0.
- Bmu: real number. larger than 0. The default value is 1.

MH_tuning

optional named list containing values used to tune the MH steps for a_xi, a_tau, c_xi, and c_tau. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. The arguments for a_xi(a_tau) are only used if learn_a_xi(learn_a_tau) is set to TRUE and mod_type is not equal to "ridge". The arguments for c_xi(c_tau) are only used if learn_c_xi(learn_c_tau) is set to TRUE and mod_type is equal to "triple". Arguments ending in "adaptive" are logical values indicating whether or not to make the MH step for the respective parameter adaptive. Arguments ending in "tuning_par" serve two different purposes. If the respective MH step is not set to be adaptive, it acts as the standard deviation of the proposal distribution. If the respective MH step is set to be adaptive, it acts as the initial standard deviation. Arguments ending in "target_rate" define the acceptance rate the algorithm aims to achieve. Arguments ending in "max_adapt" set the maximum value by which the logarithm of the standard deviation of the proposal distribution is adjusted. Finally, arguments ending in "batch_size" set the batch size after which the standard deviation of the proposal distribution is adjusted. The following elements can be supplied:

- a_xi_adaptive: logical value. The default is TRUE.
- a_xi_tuning_par: positive, real number. The default value is 1.

- a_xi_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- a_xi_max_adapt: positive, real number. The default value is 0.01.
- a_xi_batch_size: positive integer. The default value is 50.
- a_tau_adaptive: logical value. The default is TRUE.
- a_tau_tuning_par: positive, real number. The default value is 1.
- a_tau_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- a_tau_max_adapt: positive, real number. The default value is 0.01.
- a_tau_batch_size: positive integer. The default value is 50.
- c_xi_adaptive: logical value. The default is TRUE.
- c_xi_tuning_par: positive, real number. The default value is 1.
- c_xi_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- c_xi_max_adapt: positive, real number. The default value is 0.01.
- c_xi_batch_size: positive integer. The default value is 50.
- c_tau_adaptive: logical value. The default is TRUE.
- c_tau_tuning_par: positive, real number. The default value is 1.
- c_tau_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- c_tau_max_adapt: positive, real number. The default value is 0.01.
- c_tau_batch_size: positive integer. The default value is 50.
- rho_adaptive: logical value. The default is TRUE.
- rho_tuning_par: positive, real number. The default value is 1.
- rho_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- rho_max_adapt: positive, real number. The default value is 0.01.
- rho_batch_size: positive integer. The default value is 50.

starting_vals *optional* named list containing the values at which the MCMC algorithm will be initialized. In the following d refers to the number of covariates, including the intercept and expanded factors. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. The following elements can be supplied:

- beta_mean_st: vector of length d containing single numbers. The default is rep(0, d).
- theta_sr_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- tau2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- xi2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- kappa2_st: vector of length d containing single, positive numbers. The default is rep(1, d).

- `lambda2_st`: vector of length `d` containing single, positive numbers. The default is `rep(1, d)`.
- `kappa2_B_st`: positive, real number. The default value is 20.
- `lambda2_B_st`: positive, real number. The default value is 20.
- `a_xi_st`: positive, real number. The default value is 0.1.
- `a_tau_st`: positive, real number. The default value is 0.1.
- `c_xi_st`: positive, real number. The default value is 0.1. Note that the prior for `c_xi` is restricted to (0, 0.5).
- `c_tau_st`: positive, real number. The default value is 0.1. Note that the prior for `c_tau` is restricted to (0, 0.5).
- `sv_mu_st`: real number. The default value is -10.
- `sv_phi_st`: positive, real number between -1 and 1. The default value is 0.5.
- `sv_sigma2_st` : positive, real number. The default value is 1.
- `C0_st`: positive, real number. The default value is 1.
- `sigma2_st`: positive, real number if `sv` is `FALSE`, otherwise a vector of positive, real numbers of length `N`. The default value is 1 or a vector thereof.
- `h0_st`: real number. The default value is 0.
- `lambda_0_st` vector of length `d` containing positive, real numbers. The default value is `rep(1, d)`.
- `rho_st`: vector of length `d` containing real numbers between 0 and `b_rho`. The default value is `rep(max(0.1, hyperprior_param$b_rho - 0.1), d)`.

Details

For details concerning the algorithms please refer to the papers by Bitto and Frühwirth-Schnatter (2019), Cadonna et al. (2020) and Knaus and Frühwirth-Schnatter (2023). For more details on the package and the usage of the functions, see Knaus et al. (2021).

Value

The value returned is a list object of class `shrinkTVP` containing

`beta`	`list` object containing an `mcmc.tvp` object for the parameter draws from the posterior distribution of the centered states, one for each covariate. In the case that there is only one covariate, this becomes just a single `mcmc.tvp` object.
`beta_mean`	`mcmc` object containing the parameter draws from the posterior distribution of beta_mean.
`theta_sr`	`mcmc` object containing the parameter draws from the posterior distribution of the square root of theta.
`tau2`	`mcmc` object containing the parameter draws from the posterior distribution of tau2.
`xi2`	`mcmc` object containing the parameter draws from the posterior distribution of xi2.

psi — list object containing an mcmc.tvp object for the parameter draws from the posterior distribution of psi, one for each covariate. In the case that there is only one covariate, this becomes just a single mcmc.tvp object.

lambda_p — list object containing an mcmc.tvp object for the parameter draws from the posterior distribution of lambda_p, one for each covariate. In the case that there is only one covariate, this becomes just a single mcmc.tvp object.

kappa_p — *(optional)* list object containing an mcmc.tvp object for the parameter draws from the posterior distribution of kappa_p, one for each covariate. In the case that there is only one covariate, this becomes just a single mcmc.tvp object. Not returned if iid is not TRUE.

lambda2 — *(optional)* mcmc object containing the parameter draws from the posterior distribution of lambda2. Not returned if mod_type is not "triple".

kappa2 — *(optional)* mcmc object containing the parameter draws from the posterior distribution of kappa2. Not returned if mod_type is not "triple".

a_xi — *(optional)* mcmc object containing the parameter draws from the posterior distribution of a_xi. Not returned if learn_a_xi is FALSE or mod_type is "ridge".

a_tau — *(optional)* mcmc object containing the parameter draws from the posterior distribution of a_tau. Not returned if learn_a_tau is FALSE or mod_type is "ridge".

c_xi — *(optional)* mcmc object containing the parameter draws from the posterior distribution of c_xi. Not returned if learn_c_xi is FALSE or mod_type is not "triple".

c_tau — *(optional)* mcmc object containing the parameter draws from the posterior distribution of c_tau. Not returned if learn_c_tau is FALSE or mod_type is not "triple".

lambda2_B — *(optional)* mcmc object containing the parameter draws from the posterior distribution of lambda2_B. Not returned if learn_lambda2_B is FALSE or mod_type is "ridge".

kappa2_B — *(optional)* mcmc object containing the parameter draws from the posterior distribution of kappa2_B. Not returned if learn_kappa2_B is FALSE or mod_type is "ridge".

rho — *(optional)* mcmc object containing the parameter draws from the posterior distribution of rho. Not returned if iid is not TRUE.

sigma2 — mcmc object containing the parameter draws from the posterior distribution of sigma2. If sv is TRUE, sigma2 is additionally an mcmc.tvp object.

C0 — *(optional)* mcmc object containing the parameter draws from the posterior distribution of C0. Not returned if sv is TRUE.

sv_mu — *(optional)* mcmc object containing the parameter draws from the posterior distribution of the mu parameter for the stochastic volatility model on the errors. Not returned if sv is FALSE.

sv_phi — *(optional)* mcmc object containing the parameter draws from the posterior distribution of the phi parameter for the stochastic volatility model on the errors. Not returned if sv is FALSE.

sv_sigma2 — *(optional)* mcmc object containing the parameter draws from the posterior distribution of the sigma2 parameter for the stochastic volatility model on the errors. Not returned if sv is FALSE.

MH_diag	*(optional)* named list containing statistics for assessing MH performance. Not returned if no MH steps are required or none of them are specified to be adaptive.
internals	list object containing two arrays that are required for calculating the LPDS.
priorvals	list object containing hyperparameter values of the prior distributions, as specified by the user.
model	list object containing the model matrix, model response and formula used.
summaries	list object containing a collection of summary statistics of the posterior draws.

To display the output, use plot and summary. The summary method displays the specified prior values stored in priorvals and the posterior summaries stored in summaries, while the plot method calls coda's plot.mcmc or the plot.mcmc.tvp method. Furthermore, all functions that can be applied to coda::mcmc objects (e.g. coda::acfplot) can be applied to all output elements that are coda compatible.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

References

Bitto, A., & Frühwirth-Schnatter, S. (2019). "Achieving shrinkage in a time-varying parameter model framework." *Journal of Econometrics*, 210(1), 75-97. <doi:10.1016/j.jeconom.2018.11.006>

Cadonna, A., Frühwirth-Schnatter, S., & Knaus, P. (2020). "Triple the Gamma—A Unifying Shrinkage Prior for Variance and Variable Selection in Sparse State Space and TVP Models." *Econometrics*, 8(2), 20. <doi:10.3390/econometrics8020020>

Knaus, P., Bitto-Nemling, A., Cadonna, A., & Frühwirth-Schnatter, S. (2021) "Shrinkage in the Time-Varying Parameter Model Framework Using the R Package shrinkTVP." *Journal of Statistical Software* 100(13), 1–32. <doi:10.18637/jss.v100.i13>

Knaus, P., & Frühwirth-Schnatter, S. (2023). "The Dynamic Triple Gamma Prior as a Shrinkage Process Prior for Time-Varying Parameter Models." arXiv preprint arXiv:2312.10487. <doi:10.48550/arXiv.2312.10487>

See Also

plot.shrinkTVP, plot.mcmc.tvp

Examples

```
set.seed(123)
sim <- simTVP(DTG = TRUE, theta = c(0, 1, 0), beta_mean = c(1, 1, 0), rho = 0.95, c_psi = 2)
data <- sim$data

## Example 1, match the true underlying process
res <- shrinkDTVP(y ~ x1 + x2, data = data, c_psi = 2)
# summarize output
summary(res)

## Example 2, dynamic horseshoe
```

```
res <- shrinkDTVP(y ~ x1 + x2, data = data)


## Example 3, modify hyperparameters
res <- shrinkDTVP(y ~ x1 + x2, data = data,
                  hyperprior_param = list(a_rho = 1,
                                          alpha_rho = 0.5,
                                          beta_rho = 0.5))
```

---

shrinkTVP *Markov Chain Monte Carlo (MCMC) for time-varying parameter models with shrinkage*

Description

shrinkTVP samples from the joint posterior distribution of the parameters of a time-varying parameter model with shrinkage, potentially including stochastic volatility (SV), and returns the MCMC draws.

Usage

```
shrinkTVP(
  formula,
  data,
  mod_type = "double",
  niter = 10000,
  nburn = round(niter/2),
  nthin = 1,
  learn_a_xi = TRUE,
  learn_a_tau = TRUE,
  a_xi = 0.1,
  a_tau = 0.1,
  learn_c_xi = TRUE,
  learn_c_tau = TRUE,
  c_xi = 0.1,
  c_tau = 0.1,
  a_eq_c_xi = FALSE,
  a_eq_c_tau = FALSE,
  learn_kappa2_B = TRUE,
  learn_lambda2_B = TRUE,
  kappa2_B = 20,
  lambda2_B = 20,
  hyperprior_param,
  display_progress = TRUE,
  sv = FALSE,
```

```
  sv_param,
  MH_tuning,
  starting_vals
)
```

Arguments

formula	object of class "formula": a symbolic representation of the model, as in the function lm. For details, see formula.
data	*optional* data frame containing the response variable and the covariates. If not found in data, the variables are taken from environment(formula), typically the environment from which shrinkTVP is called. No NAs are allowed in the response variable and the covariates.
mod_type	character string that reads either "triple", "double" or "ridge". Determines whether the triple gamma, double gamma or ridge prior are used for theta_sr and beta_mean. The default is "double".
niter	positive integer, indicating the number of MCMC iterations to perform, including the burn-in. Has to be larger than or equal to nburn + 2. The default value is 10000.
nburn	non-negative integer, indicating the number of iterations discarded as burn-in. Has to be smaller than or equal to niter - 2. The default value is round(niter / 2).
nthin	positive integer, indicating the degree of thinning to be performed. Every nthin draw is kept and returned. The default value is 1, implying that every draw is kept.
learn_a_xi	logical value indicating whether to learn a_xi, the spike parameter of the state variances. Ignored if mod_type is set to "ridge". The default value is TRUE.
learn_a_tau	logical value indicating whether to learn a_tau, the spike parameter of the mean of the initial values of the states. Ignored if mod_type is set to "ridge". The default value is TRUE.
a_xi	positive, real number, indicating the (fixed) value for a_xi. Ignored if learn_a_xi is TRUE or mod_type is set to "ridge". The default value is 0.1.
a_tau	positive, real number, indicating the (fixed) value for a_tau. Ignored if learn_a_tau is TRUE or mod_type is set to "ridge". The default value is 0.1.
learn_c_xi	logical value indicating whether to learn c_xi, the tail parameter of the state variances. Ignored if mod_type is not set to "triple" or a_eq_c_xi is set to TRUE. The default value is TRUE.
learn_c_tau	logical value indicating whether to learn c_tau, the tail parameter of the mean of the initial values of the states. Ignored if mod_type is not set to "triple" or a_eq_c_tau is set to TRUE. The default value is TRUE.
c_xi	positive, real number, indicating the (fixed) value for c_xi. Ignored if learn_c_xi is TRUE, mod_type is not set to "triple" or a_eq_c_xi is set to TRUE. The default value is 0.1.
c_tau	positive, real number, indicating the (fixed) value for c_tau. Ignored if learn_c_xi is TRUE, mod_type is not set to "triple" or a_eq_c_tau is set to TRUE. The default value is 0.1.

a_eq_c_xi — logical value indicating whether to force a_xi and c_xi to be equal. If set to TRUE, beta_a_xi and alpha_a_xi are used as the hyperparameters and beta_c_xi and alpha_c_xi are ignored. Ignored if mod_type is not set to "triple". The default value is FALSE.

a_eq_c_tau — logical value indicating whether to force a_tau and c_tau to be equal. If set to TRUE, beta_a_tau and alpha_a_tau are used as the hyperparameters and beta_c_tau and alpha_c_tau are ignored. Ignored if mod_type is not set to "triple". The default value is FALSE.

learn_kappa2_B — logical value indicating whether to learn kappa2_B, the global level of shrinkage for the state variances. The default value is TRUE.

learn_lambda2_B — logical value indicating whether to learn the lambda2_B parameter, the global level of shrinkage for the mean of the initial values of the states. The default value is TRUE.

kappa2_B — positive, real number, indicating the (fixed) value for kappa2_B. Ignored if learn_kappa2_B is TRUE. The default value is 20.

lambda2_B — positive, real number, indicating the (fixed) value for lambda2_B. Ignored if learn_lambda2_B is TRUE. The default value is 20.

hyperprior_param — *optional* named list containing hyperparameter values. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. All hyperparameter values have to be positive, real numbers. The following hyperparameters can be supplied:

- c0: The default value is 2.5.
- g0: The default value is 5.
- G0: The default value is 5 / (2.5 - 1).
- e1: The default value is 0.001.
- e2: The default value is 0.001.
- d1: The default value is 0.001.
- d2: The default value is 0.001.
- alpha_a_xi: The default value is 5.
- alpha_a_tau: The default value is 5.
- beta_a_xi: The default value is 10.
- beta_a_tau: The default value is 10.
- alpha_c_xi: The default value is 5.
- alpha_c_tau: The default value is 5.
- beta_c_xi: The default value is 2.
- beta_c_tau: The default value is 2.

display_progress — logical value indicating whether the progress bar and other informative output should be displayed. The default value is TRUE.

sv — logical value indicating whether to use stochastic volatility for the error of the observation equation. For details please see stochvol, in particular svsample. The default value is FALSE.

`sv_param` *optional* named list containing hyperparameter values for the stochastic volatility parameters. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. Ignored if `sv` is `FALSE`. The following elements can be supplied:

- `Bsigma_sv`: positive, real number. The default value is 1.
- `a0_sv`: positive, real number. The default value is 5.
- `b0_sv`: positive, real number. The default value is 1.5.
- `bmu`: real number. The default value is 0.
- `Bmu`: real number. larger than 0. The default value is 1.

`MH_tuning` *optional* named list containing values used to tune the MH steps for `a_xi`, `a_tau`, `c_xi`, and `c_tau`. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. The arguments for `a_xi`(`a_tau`) are only used if `learn_a_xi`(`learn_a_tau`) is set to `TRUE` and `mod_type` is not equal to `"ridge"`. The arguments for `c_xi`(`c_tau`) are only used if `learn_c_xi`(`learn_c_tau`) is set to `TRUE` and `mod_type` is equal to `"triple"`. Arguments ending in "adaptive" are logical values indicating whether or not to make the MH step for the respective parameter adaptive. Arguments ending in "tuning_par" serve two different purposes. If the respective MH step is not set to be adaptive, it acts as the standard deviation of the proposal distribution. If the respective MH step is set to be adaptive, it acts as the initial standard deviation. Arguments ending in "target_rate" define the acceptance rate the algorithm aims to achieve. Arguments ending in "max_adapt" set the maximum value by which the logarithm of the standard deviation of the proposal distribution is adjusted. Finally, arguments ending in "batch_size" set the batch size after which the standard deviation of the proposal distribution is adjusted. The following elements can be supplied:

- `a_xi_adaptive`: logical value. The default is `TRUE`.
- `a_xi_tuning_par`: positive, real number. The default value is 1.
- `a_xi_target_rate`: positive, real number, between 0 and 1. The default value is 0.44.
- `a_xi_max_adapt`: positive, real number. The default value is 0.01.
- `a_xi_batch_size`: positive integer. The default value is 50.
- `a_tau_adaptive`: logical value. The default is `TRUE`.
- `a_tau_tuning_par`: positive, real number. The default value is 1.
- `a_tau_target_rate`: positive, real number, between 0 and 1. The default value is 0.44.
- `a_tau_max_adapt`: positive, real number. The default value is 0.01.
- `a_tau_batch_size`: positive integer. The default value is 50.
- `c_xi_adaptive`: logical value. The default is `TRUE`.
- `c_xi_tuning_par`: positive, real number. The default value is 1.
- `c_xi_target_rate`: positive, real number, between 0 and 1. The default value is 0.44.
- `c_xi_max_adapt`: positive, real number. The default value is 0.01.
- `c_xi_batch_size`: positive integer. The default value is 50.

- c_tau_adaptive: logical value. The default is TRUE.
- c_tau_tuning_par: positive, real number. The default value is 1.
- c_tau_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- c_tau_max_adapt: positive, real number. The default value is 0.01.
- c_tau_batch_size: positive integer. The default value is 50.

starting_vals *optional* named list containing the values at which the MCMC algorithm will be initialized. In the following d refers to the number of covariates, including the intercept and expanded factors. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. The following elements can be supplied:

- beta_mean_st: vector of length d containing single numbers. The default is rep(0, d).
- theta_sr_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- tau2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- xi2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- kappa2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- lambda2_st: vector of length d containing single, positive numbers. The default is rep(1, d).
- kappa2_B_st: positive, real number. The default value is 20.
- lambda2_B_st: positive, real number. The default value is 20.
- a_xi_st: positive, real number. The default value is 0.1.
- a_tau_st: positive, real number. The default value is 0.1.
- c_xi_st: positive, real number. The default value is 0.1. Note that the prior for c_xi is restricted to (0, 0.5).
- c_tau_st: positive, real number. The default value is 0.1. Note that the prior for c_tau is restricted to (0, 0.5).
- sv_mu_st: real number. The default value is -10.
- sv_phi_st: positive, real number between -1 and 1. The default value is 0.5.
- sv_sigma2_st : positive, real number. The default value is 1.
- C0_st: positive, real number. The default value is 1.
- sigma2_st: positive, real number if sv is FALSE, otherwise a vector of positive, real numbers of length N. The default value is 1 or a vector thereof.
- h0_st: real number. The default value is 0.

Details

For details concerning the algorithms please refer to the papers by Bitto and Frühwirth-Schnatter (2019) and Cadonna et al. (2020). For more details on the package and the usage of the functions, see Knaus et al. (2021). For the multivariate extension, see the shrinkTVPVAR function from the shrinkTVPVAR package.

Value

The value returned is a list object of class `shrinkTVP` containing

`beta`	`list` object containing an `mcmc.tvp` object for the parameter draws from the posterior distribution of the centered states, one for each covariate. In the case that there is only one covariate, this becomes just a single `mcmc.tvp` object.
`beta_mean`	`mcmc` object containing the parameter draws from the posterior distribution of beta_mean.
`theta_sr`	`mcmc` object containing the parameter draws from the posterior distribution of the square root of theta.
`tau2`	`mcmc` object containing the parameter draws from the posterior distribution of tau2.
`xi2`	`mcmc` object containing the parameter draws from the posterior distribution of xi2.
`lambda2`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of lambda2. Not returned if `mod_type` is not `"triple"`.
`kappa2`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of kappa2. Not returned if `mod_type` is not `"triple"`.
`a_xi`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of a_xi. Not returned if `learn_a_xi` is `FALSE` or `mod_type` is `"ridge"`.
`a_tau`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of a_tau. Not returned if `learn_a_tau` is `FALSE` or `mod_type` is `"ridge"`.
`c_xi`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of c_xi. Not returned if `learn_c_xi` is `FALSE` or `mod_type` is not `"triple"`.
`c_tau`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of c_tau. Not returned if `learn_c_tau` is `FALSE` or `mod_type` is not `"triple"`.
`lambda2_B`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of lambda2_B. Not returned if `learn_lambda2_B` is `FALSE` or `mod_type` is `"ridge"`.
`kappa2_B`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of kappa2_B. Not returned if `learn_kappa2_B` is `FALSE` or `mod_type` is `"ridge"`.
`sigma2`	`mcmc` object containing the parameter draws from the posterior distribution of `sigma2`. If `sv` is `TRUE`, `sigma2` is additionally an `mcmc.tvp` object.
`C0`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of C0. Not returned if `sv` is `TRUE`.
`sv_mu`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of the mu parameter for the stochastic volatility model on the errors. Not returned if `sv` is `FALSE`.
`sv_phi`	*(optional)* `mcmc` object containing the parameter draws from the posterior distribution of the phi parameter for the stochastic volatility model on the errors. Not returned if `sv` is `FALSE`.

sv_sigma2	*(optional)* mcmc object containing the parameter draws from the posterior distribution of the sigma2 parameter for the stochastic volatility model on the errors. Not returned if sv is FALSE.
MH_diag	*(optional)* named list containing statistics for assessing MH performance. Not returned if no MH steps are required or none of them are specified to be adaptive.
internals	list object containing two arrays that are required for calculating the LPDS.
priorvals	list object containing hyperparameter values of the prior distributions, as specified by the user.
model	list object containing the model matrix, model response and formula used.
summaries	list object containing a collection of summary statistics of the posterior draws.

To display the output, use plot and summary. The summary method displays the specified prior values stored in priorvals and the posterior summaries stored in summaries, while the plot method calls coda's plot.mcmc or the plot.mcmc.tvp method. Furthermore, all functions that can be applied to coda::mcmc objects (e.g. coda::acfplot) can be applied to all output elements that are coda compatible.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

References

Bitto, A., & Frühwirth-Schnatter, S. (2019). "Achieving shrinkage in a time-varying parameter model framework." *Journal of Econometrics*, 210(1), 75-97. <doi:10.1016/j.jeconom.2018.11.006>

Cadonna, A., Frühwirth-Schnatter, S., & Knaus, P. (2020). "Triple the Gamma—A Unifying Shrinkage Prior for Variance and Variable Selection in Sparse State Space and TVP Models." *Econometrics*, 8(2), 20. <doi:10.3390/econometrics8020020>

Knaus, P., Bitto-Nemling, A., Cadonna, A., & Frühwirth-Schnatter, S. (2021) "Shrinkage in the Time-Varying Parameter Model Framework Using the R Package shrinkTVP." *Journal of Statistical Software* 100(13), 1–32. <doi:10.18637/jss.v100.i13>

See Also

plot.shrinkTVP, plot.mcmc.tvp, shrinkTVPVAR

Examples

```
## Example 1, learn everything
set.seed(123)
sim <- simTVP(theta = c(0.2, 0, 0), beta_mean = c(1.5, -0.3, 0))
data <- sim$data

res <- shrinkTVP(y ~ x1 + x2, data = data)
# summarize output
summary(res)
```

```
## Example 2, hierarchical Bayesian Lasso
res <- shrinkTVP(y ~ x1 + x2, data = data,
                 learn_a_xi = FALSE, learn_a_tau = FALSE,
                 a_xi = 1, a_tau = 1)


## Example 3, non-hierarchical Bayesian Lasso
res <- shrinkTVP(y ~ x1 + x2, data = data,
                 learn_a_xi = FALSE, learn_a_tau = FALSE,
                 a_xi = 1, a_tau = 1,
                 learn_kappa2 = FALSE, learn_lambda2 = FALSE)


## Example 4, adding stochastic volatility
res <- shrinkTVP(y ~ x1 + x2, data = data,
                 sv = TRUE)


## Example 4, changing some of the default hyperparameters
res <- shrinkTVP(y ~ x1 + x2, data = data,
                 hyperprior_param = list(beta_a_xi = 5,
                                         alpha_a_xi = 10))

## Example 5, using the triple gamma prior
res <- shrinkTVP(y ~ x1 + x2, data = data,
                 mod_type = "triple")
```

`simTVP` *Generate synthetic data from a time-varying parameter model*

Description

`simTVP` generates synthetic data from a time-varying parameter model. The covariates are always generated i.i.d. from a Normal(0,1) distribution.

Usage

```
simTVP(
  N = 200,
  d = 3,
  sv = FALSE,
  DTG = FALSE,
  sigma2 = 1,
  theta,
  beta_mean,
  sv_mu = 0,
```

```
  sv_phi = 0.98,
  sv_sigma2 = 0.2,
  a_psi = 0.5,
  c_psi = 2,
  rho = 0.9
)
```

Arguments

N	integer > 2. Indicates the length of the time series to be generated. The default value is 200.
d	positive integer. Indicates the number of covariates to simulate. The default value is 3.
sv	logical value. If set to TRUE, the data will be generated with stochastic volatility for the errors of the observation equation using svsim. The default value is FALSE.
DTG	logical value. If set to TRUE, the betas will be generated as dynamic triple gamma processes. The default value is FALSE.
sigma2	positive real number. Determines the variance of the errors of the observation equation. Ignored if sv is TRUE. The default value is 1.
theta	*(optional)* vector containing positive real numbers. If supplied, these determine the variances of the innovations of the state equation. Otherwise, the elements of theta are generated from a X^2(1) distribution. Has to be of length d or an error will be thrown.
beta_mean	*(optional)* vector containing real numbers. If supplied, these determine the mean of the initial value of the state equation. Otherwise, the elements of beta_mean are generated from a Normal(0,1) distribution. Has to be of length d or an error will be thrown.
sv_mu	real number. Determines the mean of the logarithm of the volatility. Ignored if sv is FALSE. The default value is 0.
sv_phi	real number between -1 and 1. Determines the persistence of the SV process. Ignored if sv is FALSE. The default value is 0.98.
sv_sigma2	positive, real number. Determines the variance of the innovations of the logarithm of the volatility. Ignored if sv is FALSE. The default value is 0.2.
a_psi	positive, real number. Determines the pole parameter of the dynamic triple gamma process. Ignored if DTG is FALSE. The default value is 0.5.
c_psi	positive, real number. Determines the tail parameter of the dynamic triple gamma process. Ignored if DTG is FALSE. The default value is 2.
rho	real number between 0 and 1. Determines the persistence of the dynamic triple gamma process. Ignored if DTG is FALSE. The default value is 0.9

Value

The value returned is a list object containing:

data	data frame that holds the simulated data.

`true_vals` list object containing:

- `theta`: the values of theta used in the data generating process.
- `beta_mean`: the values of beta_mean used in the data generating process.
- `beta`: the true paths of beta used for the data generating process.
- `sigma2`: the value(s) of sigma2 used in the data generating process.
- `lambda_p`: the true paths of lambda_p used for the data generating process. Not returned if DTG is `FALSE`.
- `lambda_p_0`: the values of lambda_p_0 used for the data generating process. Not returned if DTG is `FALSE`.
- `kappa_p`: the true paths of kappa_p used for the data generating process. Not returned if DTG is `FALSE`.
- `psi`: the true paths of psi used for the data generating process. Not returned if DTG is `FALSE`.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

Examples

```
# Generate a time series of length 300
res <- simTVP(N = 300)

# Extract the generated data
data <- res$data

# Now with stochastic volatility
res_sv <- simTVP(N = 300, sv = TRUE)

# Now with dynamic triple gamma process
res_DTG <- simTVP(N = 300, DTG = TRUE, c_psi = 1)
```

---

updateTVP *One step update version of* shrinkTVP *with minimal overhead*

Description

`updateTVP` draws a single sample from the joint posterior distribution of the parameters of a time-varying parameter model with shrinkage potentially including stochastic volatility (SV). It performs **no input checks** and must therefore be used with caution. It is designed to be used in a modular fashion within other samplers, where speed is important. As such, no draws are saved and must be stored manually if the user wants to analyze them further.

Usage

```
updateTVP(
  y,
  x,
  curr_draws,
  mod_type = "double",
  learn_a_xi = TRUE,
  learn_a_tau = TRUE,
  a_xi = 0.1,
  a_tau = 0.1,
  learn_c_xi = TRUE,
  learn_c_tau = TRUE,
  c_xi = 0.1,
  c_tau = 0.1,
  a_eq_c_xi = FALSE,
  a_eq_c_tau = FALSE,
  learn_kappa2_B = TRUE,
  learn_lambda2_B = TRUE,
  kappa2_B = 20,
  lambda2_B = 20,
  hyperprior_param,
  sv = FALSE,
  sv_param,
  MH_tuning
)
```

Arguments

y	vector of length N containing the response variable.
x	matrix of dimension Nxd containing the covariates.
curr_draws	named list containing all the current draws from the joint posterior of the parameters. Not all values are required for all model setups. The following elements can be supplied: • beta_mean_st: vector of length d containing single numbers. • theta_sr_st: vector of length d containing single, positive numbers. • tau2_st: *optional* vector of length d containing single, positive numbers. Not required if mod_type is "ridge". • xi2_st: *optional* vector of length d containing single, positive numbers. Not required if mod_type is "ridge". • kappa2_st: *optional* vector of length d containing single, positive numbers. Only required if mod_type is "triple". • lambda2_st: *optional* vector of length d containing single, positive numbers. Only required if mod_type is "triple". • kappa2_B_st: *optional* positive, real number. Not required if mod_type is "ridge" or learn_kappa2_B is FALSE. • lambda2_B_st: *optional* positive, real number. Not required if mod_type is "ridge" or learn_lambda2_B is FALSE.

- a_xi_st: *optional* positive, real number. Not required if mod_type is "ridge" or learn_a_xi is FALSE.
- a_tau_st: *optional* positive, real number. Not required if mod_type is "ridge" or learn_a_tau is FALSE.
- c_xi_st: *optional* positive, real number. Note that the prior for c_xi is restricted to (0, 0.5). Not required if mod_type is not "triple" or learn_c_xi is FALSE.
- c_tau_st: *optional* positive, real number. Note that the prior for c_tau is restricted to (0, 0.5). Not required if mod_type is not "triple" or learn_c_tau is FALSE.
- sv_mu_st: *optional* real number. Not required if sv is FALSE.
- sv_phi_st: *optional* positive, real number between -1 and 1. Not required if sv is FALSE.
- sv_sigma2_st : *optional* positive, real number. Not required if sv is FALSE.
- C0_st: *optional* positive, real number. Not required if sv is TRUE.
- sigma2_st: positive, real number if sv is FALSE, otherwise a vector of positive, real numbers of length N. The default value is 1 or a vector thereof.
- h0_st: *optional* real number. The default value is 0. Not required if sv is FALSE.

mod_type — character string that reads either "triple", "double" or "ridge". Determines whether the triple gamma, double gamma or ridge prior are used for theta_sr and beta_mean. The default is "double".

learn_a_xi — logical value indicating whether to learn a_xi, the spike parameter of the state variances. Ignored if mod_type is set to "ridge". The default value is TRUE.

learn_a_tau — logical value indicating whether to learn a_tau, the spike parameter of the mean of the initial values of the states. Ignored if mod_type is set to "ridge". The default value is TRUE.

a_xi — positive, real number, indicating the (fixed) value for a_xi. Ignored if learn_a_xi is TRUE or mod_type is set to "ridge". The default value is 0.1.

a_tau — positive, real number, indicating the (fixed) value for a_tau. Ignored if learn_a_tau is TRUE or mod_type is set to "ridge". The default value is 0.1.

learn_c_xi — logical value indicating whether to learn c_xi, the tail parameter of the state variances. Ignored if mod_type is not set to "triple". The default value is TRUE.

learn_c_tau — logical value indicating whether to learn c_tau, the tail parameter of the mean of the initial values of the states. Ignored if mod_type is not set to "triple". The default value is TRUE.

c_xi — positive, real number, indicating the (fixed) value for c_xi. Ignored if learn_c_xi is TRUE or mod_type is not set to "triple". The default value is 0.1.

c_tau — positive, real number, indicating the (fixed) value for c_tau. Ignored if learn_c_xi is TRUE or mod_type is not set to "triple". The default value is 0.1.

a_eq_c_xi — logical value indicating whether to force a_xi and c_xi to be equal. Ignored if mod_type is not set to "triple". The default value is FALSE.

a_eq_c_tau logical value indicating whether to force a_tau and c_tau to be equal. Ignored if mod_type is not set to "triple". The default value is FALSE.

learn_kappa2_B logical value indicating whether to learn kappa2_B, the global level of shrinkage for the state variances. The default value is TRUE.

learn_lambda2_B logical value indicating whether to learn the lambda2_B parameter, the global level of shrinkage for the mean of the initial values of the states. The default value is TRUE.

kappa2_B positive, real number, indicating the (fixed) value for kappa2_B. Ignored if learn_kappa2_B is TRUE. The default value is 20.

lambda2_B positive, real number, indicating the (fixed) value for lambda2_B. Ignored if learn_lambda2_B is TRUE. The default value is 20.

hyperprior_param *optional* named list containing hyperparameter values. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. All hyperparameter values have to be positive, real numbers. The following hyperparameters can be supplied:

- c0: The default value is 2.5.
- g0: The default value is 5.
- G0: The default value is 5 / (2.5 - 1).
- e1: The default value is 0.001.
- e2: The default value is 0.001.
- d1: The default value is 0.001.
- d2: The default value is 0.001.
- beta_a_xi: The default value is 10.
- beta_a_tau: The default value is 10.
- alpha_a_xi: The default value is 5.
- alpha_a_tau: The default value is 5.
- beta_c_xi: The default value is 2.
- alpha_c_xi: The default value is 5.
- beta_c_tau: The default value is 2.
- alpha_c_tau: The default value is 5.

sv logical value indicating whether to use stochastic volatility for the error of the observation equation. For details please see stochvol, in particular svsample. The default value is FALSE.

sv_param *optional* named list containing hyperparameter values for the stochastic volatility parameters. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. Ignored if sv is FALSE. The following elements can be supplied:

- Bsigma_sv: positive, real number. The default value is 1.
- a0_sv: positive, real number. The default value is 5.
- b0_sv: positive, real number. The default value is 1.5.

- bmu: real number. The default value is 0.
- Bmu: real number. larger than 0. The default value is 1.

MH_tuning *optional* named list containing values used to tune the MH steps for a_xi, a_tau, c_xi and c_tau. Not all have to be supplied, with those missing being replaced by the default values. Any list elements that are misnamed will be ignored and a warning will be thrown. The arguments for a_xi(a_tau) are only used if learn_a_xi(learn_a_tau) is set to TRUE and mod_type is not equal to "ridge". The arguments for c_xi(c_tau) are only used if learn_c_xi(learn_c_tau) is set to TRUE and mod_type is equal to "triple". Arguments ending in "adaptive" are logical values indicating whether or not to make the MH step for the respective parameter adaptive. Arguments ending in "tuning_par" serve two different purposes. If the respective MH step is not set to be adaptive, it acts as the standard deviation of the proposal distribution. If the respective MH step is set to be adaptive, it acts as the initial standard deviation. Arguments ending in "target_rate" define the acceptance rate the algorithm aims to achieve. Arguments ending in "max_adapt" set the maximum value by which the logarithm of the standard deviation of the proposal distribution is adjusted. Finally, arguments ending in "batch_size" set the batch size after which the standard deviation of the proposal distribution is adjusted. The following elements can be supplied:

- a_xi_adaptive: logical value. The default is TRUE.
- a_xi_tuning_par: positive, real number. The default value is 1.
- a_xi_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- a_xi_max_adapt: positive, real number. The default value is 0.01.
- a_xi_batch_size: positive integer. The default value is 50.
- a_tau_adaptive: logical value. The default is TRUE.
- a_tau_tuning_par: positive, real number. The default value is 1.
- a_tau_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- a_tau_max_adapt: positive, real number. The default value is 0.01.
- a_tau_batch_size: positive integer. The default value is 50.
- c_xi_adaptive: logical value. The default is TRUE.
- c_xi_tuning_par: positive, real number. The default value is 1.
- c_xi_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- c_xi_max_adapt: positive, real number. The default value is 0.01.
- c_xi_batch_size: positive integer. The default value is 50.
- c_tau_adaptive: logical value. The default is TRUE.
- c_tau_tuning_par: positive, real number. The default value is 1.
- c_tau_target_rate: positive, real number, between 0 and 1. The default value is 0.44.
- c_tau_max_adapt: positive, real number. The default value is 0.01.
- c_tau_batch_size: positive integer. The default value is 50.

Value

The value returned is a named list object which can be immediately used as the `curr_draws` argument for another draw from the posterior with `updateTVP`. Note that, depending on the model setup, some elements may be matrices of dimension zero. It contains the following elements:

`beta_st`	dx(N + 1) matrix containing the current draw from the posterior distribution of beta.
`beta_mean_st`	dx1 matrix containing the current draws from the posterior distribution of beta_mean.
`theta_sr_st`	dx1 matrix containing the current draws from the posterior distribution of the square root of theta.
`tau2_st`	dx1 matrix containing the current draws from the posterior distribution of tau2.
`xi2_st`	dx1 matrix containing the current draws from the posterior distribution of xi2.
`lambda2_st`	dx1 matrix containing the current draws from the posterior distribution of lambda2.
`kappa2_st`	dx1 matrix containing the current draws from the posterior distribution of kappa2.
`a_xi_st`	number representing the current draw from the posterior distribution of a_xi.
`a_tau_st`	number representing the current draw from the posterior distribution of a_tau.
`c_xi_st`	number representing the current draw from the posterior distribution of c_xi.
`c_tau_st`	number representing the current draw from the posterior distribution of c_tau.
`lambda2_B_st`	number representing the current draw from the posterior distribution of lambda2_B.
`kappa2_B_st`	`mcmc` object containing the parameter draws from the posterior distribution of kappa2_B.
`sigma2_st`	number if `sv` is `FALSE`, otherwise a vector of length `N` containing the current draws from the posterior distribution of sigma2.
`C0_st`	number representing the current draw from the posterior distribution of C0.
`sv_mu_st`	number representing the current draw from the posterior distribution of the mu parameter for the stochastic volatility model on the errors.
`sv_phi_st`	number representing the current draw from the posterior distribution of the phi parameter for the stochastic volatility model on the errors.
`sv_sigma2_st`	number representing the current draw from the posterior distribution of the sigma2 parameter for the stochastic volatility model on the errors.
`h0_st`	number representing the current draw from the posterior distribution of the h0 parameter for the stochastic volatility model on the errors.
`internals`	`list` object containing two arrays that are required for calculating the LPDS and bookkeeping objects required for the adaptive MH algorithm to work.

Author(s)

Peter Knaus <peter.knaus@wu.ac.at>

References

Bitto, A., & Frühwirth-Schnatter, S. (2019). "Achieving shrinkage in a time-varying parameter model framework." *Journal of Econometrics*, 210(1), 75-97. <doi:10.1016/j.jeconom.2018.11.006>

Cadonna, A., Frühwirth-Schnatter, S., & Knaus, P. (2020). "Triple the Gamma—A Unifying Shrinkage Prior for Variance and Variable Selection in Sparse State Space and TVP Models." *Econometrics*, 8(2), 20. <doi:10.3390/econometrics8020020>

Knaus, P., Bitto-Nemling, A., Cadonna, A., & Frühwirth-Schnatter, S. (2021) "Shrinkage in the Time-Varying Parameter Model Framework Using the R Package shrinkTVP." *Journal of Statistical Software* 100(13), 1–32.<doi:10.18637/jss.v100.i13>

Examples

```
# Simulate data
sim <- simTVP()
y <- sim$data$y
x <- as.matrix(sim$data[,2:4])

# Create starting values
d <- ncol(x)
curr_draws <- list(beta_mean_st = rep(0, d),
                   theta_sr_st = rep(1, d),
                   tau2_st = rep(1, d),
                   xi2_st = rep(1, d),
                   lambda2_st = rep(1, d),
                   kappa2_B_st = 20,
                   lambda2_B_st = 20,
                   a_xi_st = 0.1,
                   a_tau_st = 0.1,
                   c_tau_st = 0.1,
                   sv_mu_st = -10,
                   sv_phi_st = 0.5,
                   sv_sigma2_st = 1,
                   C0_st = 1,
                   sigma2_st = 1,
                   h0_st = 0)

# Run the algorithm for 1000 iterations
# Note that curr_draws is always re-written and immediately re-used
for (i in 1:1000){
  curr_draws <- updateTVP(y, x, curr_draws)
}
```

Index

∗ **plotting functions**
plot.mcmc.tvp, 6
plot.shrinkTVP, 9
plot.shrinkTVP_forc, 10
∗ **prediction functions**
eval_pred_dens, 2
fitted.shrinkTVP, 3
forecast_shrinkTVP, 4
LPDS, 5
predict.shrinkTVP, 11
residuals.shrinkTVP, 13

eval_pred_dens, 2, *4–6*, *12, 13*

fitted.shrinkTVP, 3, *3*, *5, 6*, *11–13*
forecast_shrinkTVP, *3*, 4, *4*, *6*, *12, 13*
formula, *15*, *23*

LPDS, *3, 4*, 5, *5*, *12, 13*

plot.mcmc.tvp, 6, *10, 11*, *21*, *28*
plot.shrinkTVP, *8*, 9, *11*, *21*, *28*
plot.shrinkTVP_forc, *8*, 10, *10*
predict.shrinkTVP, *3–6*, 11, *13*
print.shrinkTVP, 12

residuals.shrinkTVP, *3–6*, *12*, 13

shrinkDTVP, 14
shrinkTVP, 22, *31*
shrinkTVPVAR, *26*, *28*
simTVP, 29
stochvol, *17*, *24*, *34*
svsample, *17*, *24*, *34*
svsim, *30*

updateTVP, 31