

Package ‘vasicekfit’

July 10, 2026

Title Extended Vasicek Credit Loss Model with Macroeconomic Factors

Version 0.2.0

Description Fits the extended Vasicek single-factor credit loss model where the probability of default depends on macroeconomic covariates. Maximum likelihood estimates of all parameters, including asset value correlation, are obtained via closed-form probit-transformed OLS regression; see Mayorov (2026) <[doi:10.2139/ssrn.6506378](https://doi.org/10.2139/ssrn.6506378)> for derivation.

License MIT + file LICENSE

Encoding UTF-8

RoxygenNote 7.3.2

Depends R (>= 3.5.0)

Imports stats

Suggests testthat (>= 3.0.0), vasicek, sandwich

Config/testthat/edition 3

URL <https://github.com/externalmemory/vasicekfit-cran-package>

BugReports <https://github.com/externalmemory/vasicekfit-cran-package/issues>

NeedsCompilation no

Author Dmitry Mayorov [aut, cre]

Maintainer Dmitry Mayorov <cran@dimview.org>

Repository CRAN

Date/Publication 2026-07-10 20:20:02 UTC

Contents

vasicekfit-package	2
predict.vasicekfit	2
Vasicek	3
vasicekfit	5
vasicekfit-methods	7

Index	9
--------------	----------

vasicekfit-package

Extended Vasicek Credit Loss Model with Macroeconomic Factors

Description

Fits the extended Vasicek single-factor credit loss model where the probability of default depends on macroeconomic covariates. Maximum likelihood estimates of all parameters, including asset value correlation, are obtained via closed-form probit-transformed OLS regression. Also provides density, distribution, quantile, and random generation functions for the Vasicek loss distribution.

Author(s)

Dmitriy Mayorov

References

- Vasicek, O. A. (2002). The distribution of loan portfolio value. *Risk*, 15(12), 160–162.
- Yang, B. H. (2014). Estimating Long-Run PD, Asset Correlation, and Portfolio Level PD by Vasicek Models. MPRA Paper No. 57244.

See Also

[vasicekfit](#), [predict.vasicekfit](#), [dvasicek](#)

predict.vasicekfit

Predict Method for vasicekfit Objects

Description

Obtain predictions from a fitted Vasicek model, optionally at specified confidence levels.

Usage

```
## S3 method for class 'vasicekfit'
predict(object, newdata = NULL,
        type = c("link", "response"), alpha = NULL, ...)
```

Arguments

object	a vasicekfit object.
newdata	an optional data frame of new predictor values. If omitted, the training data are used.
type	character; "link" (default) returns predictions in probit space, "response" returns the conditional mean loss rate (the effective PD) on the (0, 1) scale.
alpha	optional numeric vector of confidence levels in (0, 1). When supplied, predictions are conditional quantiles of the loss distribution at each confidence level.
...	additional arguments (currently unused).

Details

With `type = "response"` and `alpha = NULL`, the prediction is the conditional mean of the loss distribution given the covariates (the effective PD), $\Phi(\Phi^{-1}(p) + \sum_j \kappa_j u_j)$, matching the convention of `predict.glm`. Note that this is not `pnorm` of the "link" prediction: the model has an error term in probit space and the probit transform is nonlinear, so the back-transformed linear predictor $\Phi(\hat{\eta})$ is the conditional *median* of the loss distribution (available as `alpha = 0.5`), which for $p < 0.5$ is smaller than the mean.

Value

If `alpha` is `NULL`, a numeric vector of predictions. If `alpha` is supplied, a matrix with rows corresponding to observations and columns to confidence levels. When `alpha` is a scalar, the result is a vector.

See Also

[vasicekfit](#)

Examples

```
set.seed(42)
n <- 100
u <- rnorm(n)
x <- pnorm((qnorm(0.03) + 0.1 * u + sqrt(0.02) * rnorm(n)) / sqrt(1 - 0.02))
d <- data.frame(default_rate = x, macro = u)
fit <- vasicekfit(default_rate ~ macro, data = d)

# Conditional mean PD (expected loss rate given the covariates)
predict(fit, type = "response")

# Conditional median loss rate
predict(fit, alpha = 0.5, type = "response")

# 99th percentile loss rate under stress
predict(fit, newdata = data.frame(macro = 2), alpha = 0.99,
        type = "response")

# Multiple confidence levels
predict(fit, newdata = data.frame(macro = c(0, 1, 2)),
        alpha = c(0.95, 0.99, 0.999), type = "response")
```

Description

Density, distribution function, quantile function, and random generation for the Vasicek credit loss distribution.

Usage

```
dvasicek(x, p, rho, kappa = NULL, u = NULL, log = FALSE)
pvasicek(q, p, rho, kappa = NULL, u = NULL, lower.tail = TRUE, log.p = FALSE)
qvasicek(prob, p, rho, kappa = NULL, u = NULL, lower.tail = TRUE, log.p = FALSE)
rvasicek(n, p, rho, kappa = NULL, u = NULL)
```

Arguments

x, q	vector of quantiles (loss rates in (0, 1)).
prob	vector of probabilities.
n	number of observations to generate.
p	probability of default, in (0, 1).
rho	asset value correlation, in (0, 1).
kappa	optional numeric vector of macro-factor sensitivities.
u	optional numeric vector of macro-factor values, same length as kappa.
log, log.p	logical; if TRUE, probabilities/densities are given as log(p).
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise $P[X > x]$.

Details

The Vasicek loss distribution arises from the single-factor Gaussian copula model. Its probability density function is

$$f(x) = \sqrt{\frac{1-\rho}{\rho}} \exp \left\{ \frac{1}{2} \left(\Phi^{-1}(x)^2 - \left[\frac{\sqrt{1-\rho} \Phi^{-1}(x) - \sum \kappa_j u_j - \Phi^{-1}(p)}{\sqrt{\rho}} \right]^2 \right) \right\}$$

When kappa and u are NULL, this reduces to the standard two-parameter Vasicek distribution.

Value

dvasicek gives the density, pvasicek gives the distribution function, qvasicek gives the quantile function, and rvasicek generates random deviates.

References

Vasicek, O. A. (2002). The distribution of loan portfolio value. *Risk*, 15(12), 160–162.

See Also

[vasicekfit](#)

Examples

```
# Standard Vasicek density
curve(dvasicek(x, p = 0.03, rho = 0.10), from = 0.001, to = 0.20,
      ylab = "Density", main = "Vasicek loss distribution")

# 99th percentile (Basel-style)
qvasicek(0.99, p = 0.03, rho = 0.10)

# Check CDF and quantile are inverses
qvasicek(pvasicek(0.05, p = 0.03, rho = 0.10), p = 0.03, rho = 0.10)

# Random sample
set.seed(1)
hist(rvasicek(10000, p = 0.03, rho = 0.10), breaks = 100,
     main = "Vasicek loss distribution samples")
```

vasicekfit

Fit an Extended Vasicek Credit Loss Model

Description

Fits the extended Vasicek single-factor credit loss model where the probability of default depends on macroeconomic covariates. All parameters, including asset value correlation, are estimated via closed-form probit-transformed OLS.

Usage

```
vasicekfit(formula, data, bias_correct = FALSE, portfolio_size = NULL)
```

Arguments

formula	a formula of the form <code>response ~ predictors</code> where the response is an observed default/loss rate in (0, 1) (or [0, 1] when <code>portfolio_size</code> is supplied). The model must include an intercept.
data	a data frame containing the variables in formula.
bias_correct	logical; if TRUE, apply the small-sample bias correction to the variance estimate by multiplying by $N/(N - m - 1)$, where N is number of observations and m is number of predictors.
portfolio_size	optional positive integer. If supplied, a finite-portfolio variance correction is applied to the response before fitting (see Yang, 2014, section 4.3). The correction shrinks the observed rates towards their mean, mapping rates of exactly 0 or 1 into (0, 1), so such observations are permitted when <code>portfolio_size</code> is given.

Value

An object of class "vasicekfit", which is a list containing:

p estimated probability of default (baseline PD).

rho estimated asset value correlation.

kappa named numeric vector of macro-factor sensitivities.

sigma2 MLE variance estimate in probit space.

lm_fit the underlying `lm` object.

fitted.values fitted values in probit (Y) space.

residuals residuals in probit (Y) space.

formula the model formula.

call the matched call.

terms the `terms` object.

model the model frame.

bias_correct logical flag used.

portfolio_size portfolio size used, or NULL.

References

Vasicek, O. A. (2002). The distribution of loan portfolio value. *Risk*, 15(12), 160–162.

Yang, B. H. (2014). Estimating Long-Run PD, Asset Correlation, and Portfolio Level PD by Vasicek Models. MPRA Paper No. 57244.

See Also

`predict.vasicekfit`, `dvasicek`

Examples

```
set.seed(42)
n <- 100
u1 <- rnorm(n)
u2 <- rnorm(n)
p_true <- 0.03; rho_true <- 0.02
kappa_true <- c(0.13, -0.07)
z <- rnorm(n)
x <- pnorm((qnorm(p_true) + kappa_true[1] * u1 + kappa_true[2] * u2 +
  sqrt(rho_true) * z) / sqrt(1 - rho_true))
d <- data.frame(default_rate = x, unemp = u1, hpi = u2)
fit <- vasicekfit(default_rate ~ unemp + hpi, data = d)
fit
```

Description

Print, summary, coefficient extraction, covariance, confidence interval, and accessor methods for "vasicekfit" objects.

Usage

```
## S3 method for class 'vasicekfit'
print(x, ...)
## S3 method for class 'vasicekfit'
summary(object, type = c("iid", "HAC"), ...)
## S3 method for class 'vasicekfit'
coef(object, ...)
## S3 method for class 'vasicekfit'
fitted(object, ...)
## S3 method for class 'vasicekfit'
residuals(object, ...)
## S3 method for class 'vasicekfit'
vcov(object, type = c("iid", "HAC"), ...)
## S3 method for class 'vasicekfit'
confint(object, parm, level = 0.95,
        type = c("iid", "HAC"), ...)
```

Arguments

x, object	a vasicekfit object.
parm	a specification of which parameters to compute confidence intervals for, either a vector of names or indices. If missing, all parameters are considered.
level	the confidence level required.
type	covariance type. "iid" (default) uses the closed-form block-diagonal Gaussian covariance. "HAC" uses a kernel long-run covariance of the joint influence functions for $(\hat{\beta}, \hat{\sigma}^2)$ and requires the sandwich package.
...	when type = "HAC", extra arguments are forwarded to lrvar , e.g. prewhite, kernel, bw, or adjust. (lrvar's own type argument cannot be set this way because the name collides; if you need Newey-West-style bandwidth selection, pass kernel = "Bartlett" and an explicit bw.) Ignored when type = "iid".

Details

coef returns the recovered Vasicek parameters (p, rho, and kappas) in the original parameter space. fitted and residuals return values in probit (Y) space, consistent with the underlying OLS regression.

`vcov` returns the delta-method covariance matrix for the original-space parameters $(p, \rho, \kappa_1, \dots, \kappa_m)$, obtained by propagating the joint covariance of $(\hat{\beta}, \hat{\sigma}^2)$ through the closed-form recovery transform.

With `type = "iid"` the $(\hat{\beta}, \hat{\sigma}^2)$ block is treated as block-diagonal by the independence of sample mean and variance under normality. The variance of $\hat{\sigma}^2$ uses $2\sigma^4(N - m - 1)/N^2$ for the MLE and $2\sigma^4/(N - m - 1)$ when `bias_correct = TRUE`.

With `type = "HAC"` the joint covariance of $(\hat{\beta}, \hat{\sigma}^2)$ is estimated as the long-run variance of the stacked influence functions $\psi_{\beta,i} = (X'X/N)^{-1}x_ie_i$ and $\psi_{\sigma^2,i} = e_i^2 - \hat{\sigma}^2$, computed via `lrvar`. This is appropriate when the data exhibit serial correlation (e.g. quarterly macro covariates), at the cost of some finite-sample bias. Defaults follow `lrvar` (Andrews quadratic-spectral kernel, prewhitened); pass `prewhite`, `kernel`, `bw`, or `adjust` via `...` to override.

The HAC standard errors do not change with `bias_correct`: the long-run variance is estimated empirically from the influence functions rather than from the parametric Gaussian formula, so the $N/(N - m - 1)$ scaling does not apply. The `bias_correct` flag still affects the point estimates of p , ρ , and κ (via $\hat{\sigma}^2$) regardless of `type`.

`confint` returns Wald intervals based on `vcov`.

Value

`print` and `summary` return object invisibly.

`coef` returns a named numeric vector.

`fitted` and `residuals` return numeric vectors.

`vcov` returns a $(m + 2) \times (m + 2)$ covariance matrix.

`confint` returns a two-column matrix of lower and upper bounds.

See Also

`vasicekfit`, `predict.vasicekfit`

Index

`coef.vasicekfit (vasicekfit-methods)`, 7
`confint.vasicekfit`
 `(vasicekfit-methods)`, 7

`dvasicek`, 2, 6
`dvasicek (Vasicek)`, 3

`fitted.vasicekfit (vasicekfit-methods)`,
 7

`lm`, 6
`lrvar`, 7, 8

`predict.glm`, 3
`predict.vasicekfit`, 2, 2, 6, 8
`print.vasicekfit (vasicekfit-methods)`, 7
`pvasicek (Vasicek)`, 3

`qvasicek (Vasicek)`, 3

`residuals.vasicekfit`
 `(vasicekfit-methods)`, 7
`rvasicek (Vasicek)`, 3

`summary.vasicekfit`
 `(vasicekfit-methods)`, 7

`terms`, 6

`Vasicek`, 3
`vasicekfit`, 2–4, 5, 7, 8
`vasicekfit-methods`, 7
`vasicekfit-package`, 2
`vcov.vasicekfit (vasicekfit-methods)`, 7