

Package ‘bidask’

October 13, 2025

Type Package

Title Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices

Version 2.1.5

Description

Implements the efficient estimator of bid-ask spreads from open, high, low, and close prices described in Ardia, Guidotti, & Kroencke (JFE, 2024) <[doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)>. It also provides an implementation of the estimators described in Roll (JF, 1984) <[doi:10.1111/j.1540-6261.1984.tb03897.x](https://doi.org/10.1111/j.1540-6261.1984.tb03897.x)>, Corwin & Schultz (JF, 2012) <[doi:10.1111/j.1540-6261.2012.01729.x](https://doi.org/10.1111/j.1540-6261.2012.01729.x)>, and Abdi & Rinaldo (RFS, 2017) <[doi:10.1093/rfs/hhx084](https://doi.org/10.1093/rfs/hhx084)>.

License MIT + file LICENSE

URL <https://github.com/eguidotti/bidask>

BugReports <https://github.com/eguidotti/bidask/issues>

Encoding UTF-8

Imports data.table

RoxygenNote 7.2.3

Suggests xts, zoo, dplyr, crypto2, quantmod, ggplot2, knitr, rmarkdown, testthat (>= 3.0.0)

Config/testthat/edition 3

VignetteBuilder knitr

NeedsCompilation no

Author Emanuele Guidotti [aut, cre] (ORCID: <<https://orcid.org/0000-0002-8961-6623>>),
David Ardia [ctb] (ORCID: <<https://orcid.org/0000-0003-2823-782X>>),
Tim Kroencke [ctb] (ORCID: <<https://orcid.org/0000-0001-8700-356X>>)

Maintainer Emanuele Guidotti <emanuele.guidotti@usi.ch>

Repository CRAN

Date/Publication 2025-10-13 13:30:15 UTC

Contents

edge	2
edge_expanding	3
edge_rolling	4
sim	5
spread	7
Index	9

edge	<i>Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices</i>
------	---

Description

Implements the efficient estimator of bid-ask spreads from open, high, low, and close prices described in Ardia, Guidotti, & Kroencke (JFE, 2024): [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)

Usage

```
edge(open, high, low, close, sign = FALSE)
```

Arguments

open	numeric vector of open prices.
high	numeric vector of high prices.
low	numeric vector of low prices.
close	numeric vector of close prices.
sign	whether to return signed estimates.

Details

Prices must be sorted in ascending order of the timestamp.

Value

The spread estimate. A value of 0.01 corresponds to a spread of 1%.

References

Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. Journal of Financial Economics, 161, 103916. [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)

Examples

```
# reduce number of threads to pass CRAN checks (you can ignore this)
data.table::setDTthreads(1)

# simulate open, high, low, and close prices with spread 1%
x <- sim(n = 1000, spread = 0.01)

# estimate the spread
edge(x$Open, x$High, x$Low, x$Close)
```

edge_expanding	<i>Expanding Estimates of Bid-Ask Spreads from Open, High, Low, and Close Prices</i>
----------------	--

Description

Implements an expanding window calculation of the efficient estimator of bid-ask spreads from open, high, low, and close prices described in Ardia, Guidotti, & Kroencke (JFE, 2024): [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916).

Usage

```
edge_expanding(open, high, low, close, sign = FALSE, na.rm = TRUE)
```

Arguments

open	numeric vector of open prices.
high	numeric vector of high prices.
low	numeric vector of low prices.
close	numeric vector of close prices.
sign	whether to return signed estimates.
na.rm	whether to ignore missing values.

Details

Prices must be sorted in ascending order of the timestamp.

Value

Vector of spread estimates. A value of 0.01 corresponds to a spread of 1%. This function always returns a result of the same length as the input prices.

References

Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. Journal of Financial Economics, 161, 103916. [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)

Examples

```
# reduce number of threads to pass CRAN checks (you can ignore this)
data.table::setDTthreads(1)

# simulate open, high, low, and close prices with spread 1%
x <- sim(n = 1000, spread = 0.01)

# estimate the spread using an expanding window
s <- edge_expanding(x$Open, x$High, x$Low, x$Close)
tail(s)
# equivalent to
s <- edge_rolling(x$Open, x$High, x$Low, x$Close, width = 1:nrow(x), na.rm = TRUE)
tail(s)
```

edge_rolling	<i>Rolling Estimates of Bid-Ask Spreads from Open, High, Low, and Close Prices</i>
--------------	--

Description

Implements a rolling window calculation of the efficient estimator of bid-ask spreads from open, high, low, and close prices described in Ardia, Guidotti, & Kroencke (JFE, 2024): [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916).

Usage

```
edge_rolling(open, high, low, close, width, sign = FALSE, na.rm = FALSE)
```

Arguments

open	numeric vector of open prices.
high	numeric vector of high prices.
low	numeric vector of low prices.
close	numeric vector of close prices.
width	if an integer, the width of the rolling window. If a vector with the same length of the input prices, the width of the window corresponding to each observation. Otherwise, a vector of endpoints. See examples.
sign	whether to return signed estimates.
na.rm	whether to ignore missing values.

Details

Prices must be sorted in ascending order of the timestamp.

Value

Vector of spread estimates. A value of 0.01 corresponds to a spread of 1%. This function always returns a result of the same length as the input prices.

References

Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. *Journal of Financial Economics*, 161, 103916. [doi:10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)

Examples

```
# reduce number of threads to pass CRAN checks (you can ignore this)
data.table::setDTthreads(1)

# simulate open, high, low, and close prices with spread 1%
x <- sim(n = 1000, spread = 0.01)

# estimate the spread using a rolling window
s <- edge_rolling(x$Open, x$High, x$Low, x$Close, width = 21)
tail(s)

# estimate the spread using custom endpoints
ep <- c(3, 35, 100)
s <- edge_rolling(x$Open, x$High, x$Low, x$Close, width = ep)
s[c(35, 100)]
# equivalent to
edge(x$Open[3:35], x$High[3:35], x$Low[3:35], x$Close[3:35])
edge(x$Open[35:100], x$High[35:100], x$Low[35:100], x$Close[35:100])

# estimate the spread using an expanding window
s <- edge_rolling(x$Open, x$High, x$Low, x$Close, width = 1:nrow(x))
tail(s)
# equivalent to
s <- edge_expanding(x$Open, x$High, x$Low, x$Close, na.rm = FALSE)
tail(s)
```

sim

*Simulation of Open, High, Low, and Close Prices***Description**

This function performs simulations consisting of n periods and where each period consists of a given number of trades. For each trade, the actual price P_t is simulated as $P_t = P_{t-1}e^{\sigma x}$, where σ is the standard deviation per trade and x is a random draw from a unit normal distribution. The standard deviation per trade equals the volatility divided by the square root of the number of trades. Trades are assumed to be observed with a given probability. The bid (ask) for each trade is defined as P_t multiplied by one minus (plus) half the spread and we assume a 50% chance

that a bid (ask) is observed. High and low prices equal the highest and lowest prices observed during the period. Open and Close prices equal the first and the last price observed in the period. If no trade is observed for a period, then the previous Close is used as the Open, High, Low, and Close prices for that period.

Usage

```
sim(
  n = 10000,
  trades = 390,
  prob = 1,
  spread = 0.01,
  volatility = 0.03,
  overnight = 0,
  drift = 0,
  units = 1,
  sign = FALSE
)
```

Arguments

<code>n</code>	the number of periods to simulate.
<code>trades</code>	the number of trades per period.
<code>prob</code>	the probability to observe a trade.
<code>spread</code>	the bid-ask spread.
<code>volatility</code>	the open-to-close volatility.
<code>overnight</code>	the close-to-open volatility.
<code>drift</code>	the expected return per period.
<code>units</code>	the units of the time period. One of: 1, sec, min, hour, day, week, month, year.
<code>sign</code>	whether to return positive prices for buys and negative prices for sells.

Value

A data.frame of open, high, low, and close prices if `units=1` (default). Otherwise, an xts object is returned (requires the xts package to be installed).

References

Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. *Journal of Financial Economics*, 161, 103916. doi:10.1016/j.jfineco.2024.103916

Examples

```
# reduce number of threads to pass CRAN checks (you can ignore this)
data.table::setDTthreads(1)
```

```
# simulate 10 open, high, low, and close prices with spread 1%
sim(n = 10, spread = 0.01)
```

spread	<i>Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices</i>
--------	---

Description

This function implements several methods to estimate bid-ask spreads from open, high, low, and close prices and it is optimized for fast calculations over rolling and expanding windows.

Usage

```
spread(x, width = nrow(x), method = "EDGE", sign = FALSE, na.rm = FALSE)
```

Arguments

x	tabular data with columns named open, high, low, close (case-insensitive).
width	if an integer, the width of the rolling window. If a vector with the same length of the input prices, the width of the window corresponding to each observation. Otherwise, a vector of endpoints. By default, the full sample is used to compute a single spread estimate. See examples.
method	the estimators to use. See details.
sign	whether to return signed estimates.
na.rm	whether to ignore missing values.

Details

The method EDGE implements the Efficient Discrete Generalized Estimator described in Ardia, Guidotti, & Kroencke (JFE, 2024).

The methods OHL, OHLC, CHL, CHLO implement the generalized estimators described in Ardia, Guidotti, & Kroencke (JFE, 2024). They can be combined by concatenating their identifiers, e.g., OHLC.CHLO uses an average of the OHLC and CHLO estimators.

The method AR implements the estimator described in Abdi & Ranaldo (RFS, 2017). AR2 implements their 2-period version.

The method CS implements the estimator described in Corwin & Schultz (JF, 2012). CS2 implements their 2-period version. Both versions are adjusted for overnight (close-to-open) returns as described in the paper.

The method ROLL implements the estimator described in Roll (JF, 1984).

Value

A data.frame of spread estimates, or an xts object if x is of class xts. A value of 0.01 corresponds to a spread of 1%.

References

- Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. *Journal of Financial Economics*, 161, 103916. doi:[10.1016/j.jfineco.2024.103916](https://doi.org/10.1016/j.jfineco.2024.103916)
- Abdi, F., & Ranaldo, A. (2017). A simple estimation of bid-ask spreads from daily close, high, and low prices. *Review of Financial Studies*, 30 (12), 4437-4480. doi:[10.1093/rfs/hhx084](https://doi.org/10.1093/rfs/hhx084)
- Corwin, S. A., & Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67 (2), 719-760. doi:[10.1111/j.15406261.2012.01729.x](https://doi.org/10.1111/j.15406261.2012.01729.x)
- Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39 (4), 1127-1139. doi:[10.1111/j.15406261.1984.tb03897.x](https://doi.org/10.1111/j.15406261.1984.tb03897.x)

Examples

```
# reduce number of threads to pass CRAN checks (you can ignore this)
data.table::setDTthreads(1)

# simulate open, high, low, and close prices with spread 1%
x <- sim(n = 1000, spread = 0.01)

# estimate the spread
spread(x)
# equivalent to
edge(x$Open, x$High, x$Low, x$Close)

# estimate the spread using a rolling window of 21 periods
s <- spread(x, width = 21)
tail(s)
# equivalent to
s <- edge_rolling(x$Open, x$High, x$Low, x$Close, width = 21)
tail(s)

# estimate the spread using an expanding window
s <- spread(x, width = 1:nrow(x))
tail(s)
# equivalent to
s <- edge_expanding(x$Open, x$High, x$Low, x$Close, na.rm = FALSE)
tail(s)

# estimate the spread using custom endpoints
ep <- c(3, 35, 100)
spread(x, width = ep)
# equivalent to
edge(x$Open[3:35], x$High[3:35], x$Low[3:35], x$Close[3:35])
edge(x$Open[35:100], x$High[35:100], x$Low[35:100], x$Close[35:100])

# use multiple estimators
spread(x, method = c("EDGE", "AR", "CS", "ROLL", "OHLC", "OHL.CHL"))
```


Index

edge, [2](#)
edge_expanding, [3](#)
edge_rolling, [4](#)

sim, [5](#)
spread, [7](#)