MF2-GARCH Toolbox for Matlab

Christian Conrad[*] and Julius Theodor Schoelkopf[†]

September 8, 2025

A Matlab package for estimating and forecasting volatility using the multiplicative factor multi-frequency GARCH (MF2-GARCH) proposed in *"Modelling Volatility Cycles: The MF2-GARCH Model"* by Conrad and Engle (2025). The MF2-GARCH model is used in *"Long-term Volatility shapes the Stock Market's Sensitivity to News" by Conrad et al. (2025).*

- A comprehensive toolbox for estimating and forecasting volatility using the MF2-GARCH-rw-m.
- Code for four applications: estimation, news impact curve, out-of-sample forecasting, and illustration of forecasting behavior.

We do not assume any responsibilities for results produced with the available code. Please let us know if you have suggestions for further versions or find any bugs.

Download the Repository: github.com/juliustheodor/mf2garch

Suggested Citation: Please cite as:

> Conrad, Christian and Engle, Robert F., *Modelling Volatility Cycles: The MF2-GARCH Model* (2005). *Journal of Applied Econometrics*, 40(4): 438–454. https://doi.org/10.1002/jae.3118.

and

> Conrad, Christian and Schoelkopf, Julius Theodor and Tushteva, Nikoleta, *Long-Term Volatility Shapes the Stock Market's Sensitivity to News* (2025). Available at SSRN: http://dx.doi.org/10.2139/ssrn.4632733.

and

> Conrad, Christian and Schoelkopf, Julius Theodor (2025). *MF2-GARCH Toolbox for Matlab*. Matlab package version 0.1.0.

[*]Heidelberg University, Department of Economics, Bergheimer Straße 58, 69115 Heidelberg, Germany; KOF Swiss Economic Institute; HEiKA - Heidelberg Karlsruhe Strategic Partnership, Heidelberg University, Karlsruhe Institute of Technology. Email: christian.conrad@awi.uni-heidelberg.de. Website: uni-heidelberg.de/conrad.

[†]Heidelberg University, Department of Economics, Bergheimer Straße 58, 69115 Heidelberg, Germany. Email: julius.schoelkopf@awi.uni-heidelberg.de. Website: julius-schoelkopf.com.

1 Applications

1.1 Estimation of the MF2-GARCH-rw-m model in Matlab for S&P 500 stock returns

Define daily log-returns as $y_t = \mu + \sigma_t Z_t = \mu + \sqrt{h_t \tau_t}\, Z_t$, where the Z_t are i.i.d. with mean zero, variance one, and symmetric density. The fourth moment of the Z_t is denoted by κ. In the following, we write the demeaned returns as $r_t = y_t - \mu$. The conditional variance is denoted by σ_t^2 and the short- and long-term volatility components are given by h_t and τ_t. Let y be a $(T \times 1)$ vector of daily log-returns.

The short-term volatility component is defined as a unit variance GJR-GARCH(1,1)

$$h_t = (1-\phi) + \left(\alpha + \gamma \mathbf{1}_{\{r_{t-1}<0\}}\right) \frac{r_{t-1}^2}{\tau_{t-1}} + \beta h_{t-1} \tag{1}$$

and the long-term component is specified as a MEM equation for the conditional expectation of $V_t = r_t^2/h_t$ (squared deGARCHed returns):

$$\tau_t = \lambda_0 + \lambda_1 V_{t-1}^{(m)} + \lambda_2 \tau_{t-1}, \tag{2}$$

where

$$V_{t-1}^{(m)} = \frac{1}{m}\sum_{j=1}^{m} V_{t-j} = \frac{1}{m}\sum_{j=1}^{m} \frac{r_{t-j}^2}{h_{t-j}}\,. \tag{3}$$

The MF2-GARCH can be estimated using the following function from the toolbox in Matlab:

```
[coeff, qmle_se, p_value_qmle, Z, h, tau, sigma_annual, tau_annual, ...
 annual_unconditional_vola, foptions] = mf2_garch_estimation(y, foptions);
```

The function `mf2_garch_estimation(y, foptions)` prints estimation output for the seven parameters

$$(\mu, \alpha, \gamma, \beta, \lambda_0, \lambda_1, \lambda_2)$$

of the short- and long-term component by maximizing the log-likelihood. The outputs include the coefficient estimates (`coeff`), Bollerslev-Wooldridge robust standard errors (`qmle_se`), and p-values (`p_value_qmle`).

Moreover, `mf2_garch_estimation(y, foptions)` returns standardized residuals Z, fitted values for the short (`h`) and long-term (`tau`, or annualized `tau_annual`) components, the annualized conditional volatility time series (`sigma_annual`), and the estimate of annualized unconditional volatility (`annual_unconditional_vola`).

For the long-term component, you need to specify m, i.e., the number of days over which $V_t^{(m)}$ is computed. Choose whether you want to use a fixed value of m or let the optimal m be selected as the one that minimizes the BIC. The `foptions` structure contains the researcher's choice for m. Either specify `foptions.choice = 'BIC'` to search over m, or use `foptions.choice = 'fix'` together with `foptions.m = 63`. When `foptions.choice = 'BIC'`, the code estimates models for values of m between 20 and 150 and determines the optimal m as the one minimizing the BIC (Schwarz, 1978). The code also generates a figure of BIC versus m. As shown in Figure 1 (see Figure 3 in Conrad and Engle (2025)), the lowest BIC materializes for $m = 63$ (red line).

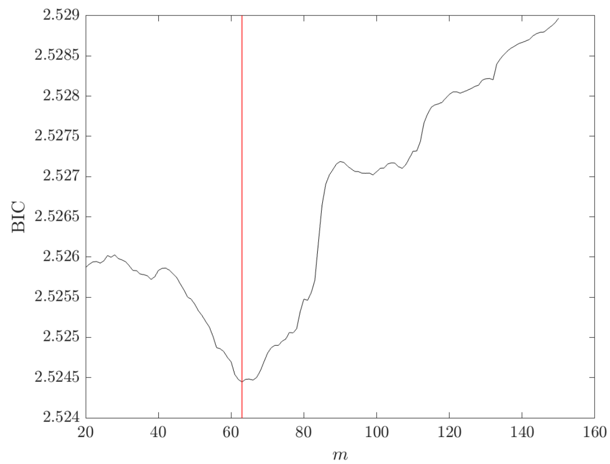

Figure 1: BIC as a function of m. See Figure 3 in Conrad and Engle (2025).

When computing the likelihood, the first two years of y (i.e., 2×252 trading days) are discarded to account for lags of the squared deGARCHed returns in the long-term component. This allows comparing the BIC of models with different values of m. You could decrease this, but you need to discard at least $2m$ values. The Matlab function uses parameter constraints following Assumption 2 (short-term component) and Assumption 3 (long-term component) of Conrad and Engle (2025). For details on the estimation, see Section A.1.1 in Conrad and Engle (2025).

The following application replicates the second panel in Table 2 in Conrad and Engle (2025) for the MF2-GARCH-rw-m. In Conrad and Engle (2025), all models were estimated using OxMetrics. We use daily S&P 500 log-return data from January 1971 to June 2023. For the sub-period 1971–1983, the return data were initially obtained from the Federal Reserve Bank of St. Louis database. Data from 1983 onwards are from TickData.

```
%% Import the return data to Matlab (S&P500 returns from 1971-2023)
% Read the data into a table
Returns = readtable('data/SP500_1971_2023_06_30_ret.xlsx');

% Extract the column 'RET_SPX' from the table and store it
y = Returns.RET_SPX;

%% Select the m for the estimation
foptions.choice = 'fix'; % choices: 'BIC' or 'fix' (specify m)

% If foptions.choice = 'fix', please specify the m you choose here:
foptions.m = 63;

% Example A (Estimation) for regression output:
mf2_garch_estimation(y, foptions);
```

This yields output in the Matlab command window (example excerpt):

```
=============== Estimation results MF2-GARCH-rw-m ===============
The optimal m was specified by the user: m = 63
Log-Likelihood Function = -16678.611, BIC = 2.524
Estimated fourth moment of the innovations: kappa = 5.441
     Parameter      Coefficient    Standard Error  p-value     Significance

    ____________     _________     ___________     __________    ________

    {'mu'       }     0.030395      0.0069646       1.2758e-05      ***
    {'alpha'    }     0.0032236     0.0026327          0.22078
    {'gamma'    }     0.16169       0.020378        2.2204e-15      ***
    {'beta'     }     0.83956       0.017385                 0      ***
    {'lambda_0'}      0.017512      0.0072155         0.015226       **
    {'lambda_1'}      0.11183       0.046429          0.016013       **
    {'lambda_2'}      0.87014       0.051675                 0      ***
Output reports Bollerslev-Wooldridge robust standard errors (see Conrad and
Engle (2025), equation (27)).
Covariance stationarity condition satisfied (see Conrad and Engle (2025),
equation (7)): Gamma_m = 0.778
Annualized unconditional volatility = 16.043
==========================================================
```

If you additionally want to store the fitted values, specify the output of the function as follows:

```
[coeff, qmle_se, p_value_qmle, Z, h, tau, sigma_annual, tau_annual, ...
 annual_unconditional_vola, foptions]  = mf2_garch_estimation(y, foptions);
```

You can use this output to create figures for estimated conditional and long-term volatilities over the full sample. Grey shaded areas represent NBER recession periods for the US. The function exports the Figure 2 in the figures folder.

```
% Extract the date column (not required for estimation, only for figure)
dates = datetime(Returns.OBS, 'InputFormat', 'MM/dd/yyyy');

% Figure of time series
mf2_garch_time_series(dates, sigma_annual, tau_annual);
```

Additionally, you can plot the news impact curve for the estimated model. Following Engle and Ng (1993), we use the NIC to illustrate how the conditional volatility is updated in response to new information. The NIC is presented in terms of annualized volatilities (see equation (10) in Conrad and Engle (2025)). The following function provides a figure for the NIC (see Figure 3):

```
 [r, NIC] = mf2_garch_nic(Z, h, tau, foptions, coeff);
```

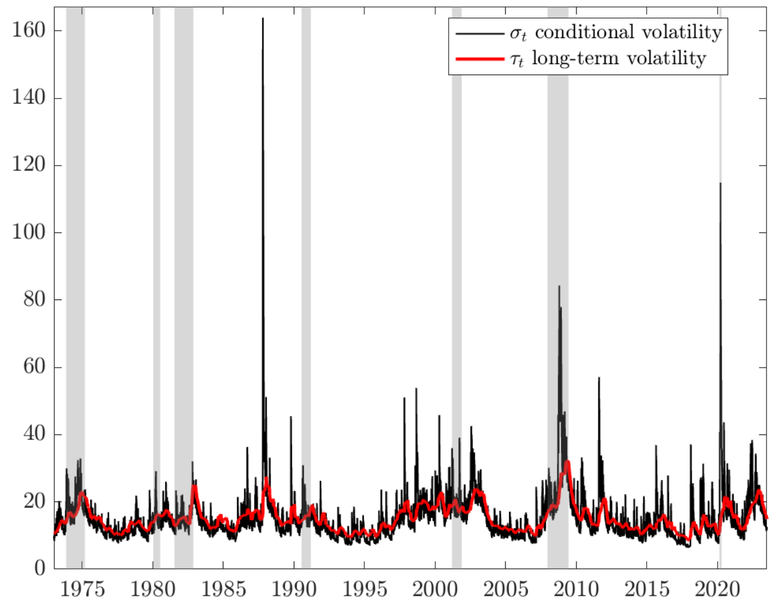

Figure 2: The figure shows the estimated conditional volatility, $\sqrt{h_t \tau_t}$, (black line) and long-term volatility, $\sqrt{\tau_t}$, (red line) from the MF2-GARCH-rw-63 model for the daily S&P 500 returns. All quantities are annualized. Gray shaded areas represent NBER recession periods.

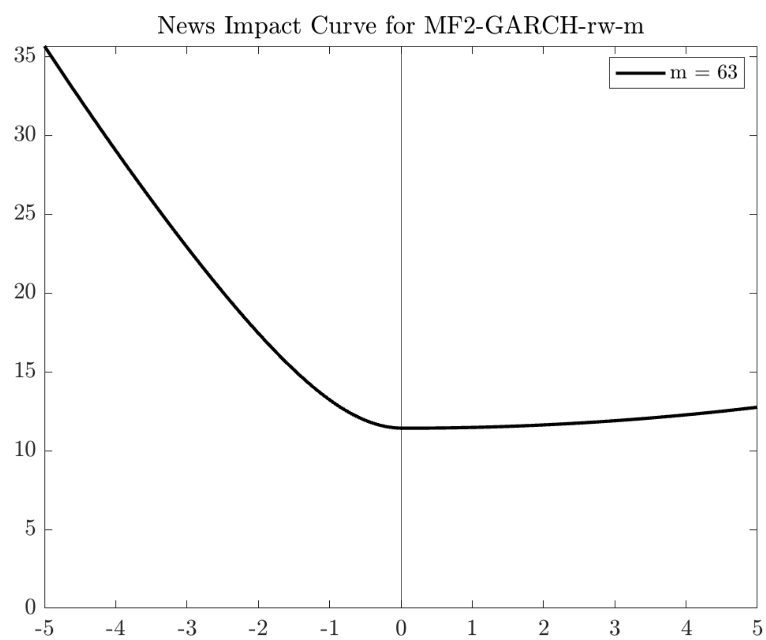

Figure 3: The figure shows the NIC for an MF2-GARCH-rw-m model with $m = 63$.

1.2 Forecasting at the end of the sample

First, you need to specify the maximum forecasting horizon using `foptions.S`, e.g. `foptions.S = 250` if you want to forecast the next 250 days. Next, you can use the forecasting function that provides forecasts for the (annualized) conditional volatility, the short- and (annualized) long-term component:

```
[horizon, forecast, an_vola_forecast, h_forecast, tau_forecast,
   tau_forecast_annual]  = mf2_garch_forecasting(y, Z, h, tau, coeff,
   foptions);
```

You must use the same sample as in the testimation function for the forecasting function. Moreover, the function displays in the command window the forecasts (from the end of the sample) for the annualized volatility on the next day, next week (5 days), next month (21 days), next 6 months (126 days), and 12 months (252 days) based on the estimated parameters.

```
annualized volatility forecast 1 day: 11.4935
annualized volatility forecast 1 week (5 days): 12.4475
annualized volatility forecast 1 month (21 days): 14.5593
annualized volatility forecast 6 months (126 days): 15.4574
annualized volatility forecast 1 year (252 days): 15.6806
```

We now want to illustrate out-of-sample forecasting using a figure. The following code yields a figure of the forecasts of long-term volatility and conditional volatility in the last 50 days of the sample and the forecasts for the next S days:

```
mf2_garch_out_of_sample_figure(sigma_annual, an_vola_forecast,
   tau_forecast_annual, annual_unconditional_vola, foptions)
```

1.3 Illustration of Forecasting behavior

Last, we want to illustrate the MF2-GARCH's out-of-sample forecast performance, as in Figure 5 in Conrad and Engle (2025). We want to forecast volatility from August 10, 2011 (10249 in dates vector) 120 days into the future and use the forecasting function:

```
% Specify the maximum forecasting horizon:
foptions.S = 120;

% Estimation of the MF2-GARCH We want to forecast volatility from August
   10, 2011 (10249 in dates vector) 150 days into the future. Specify the
   cutoff from where you want to forecast:

foptions.cutoff_date = datetime(2011,8,10);
foptions.cutoff = 10249;
```

```
% Therefore, we need to reestimate the model using data until August 10,
    2011.
[coeff, ~, ~, Z, h, tau, ~, tau_annual, annual_unconditional_vola, foptions
    ]  = mf2_garch_estimation(y(1:foptions.cutoff),foptions);

% Forecasting exercise: This function provides forecasts for the annualized
     volatility, h and tau for the next S days from the end of the specified
     sample.
[horizon, forecast, an_vola_forecast, h_forecast, tau_forecast,
    tau_forecast_annual]  = mf2_garch_forecasting(y(1:cutoff), Z, h, tau,
    coeff, foptions);

% Illustration of forecasting behaviour as in Figure 5 from Conrad & Engle
    (2025):
mf2_garch_illustration_forecasting_figure(sigma_annual, an_vola_forecast,
    tau_forecast_annual, annual_unconditional_vola, foptions, dates)
```

The following figure is saved as `'ForecastIllustration.png'` in the figures folder:

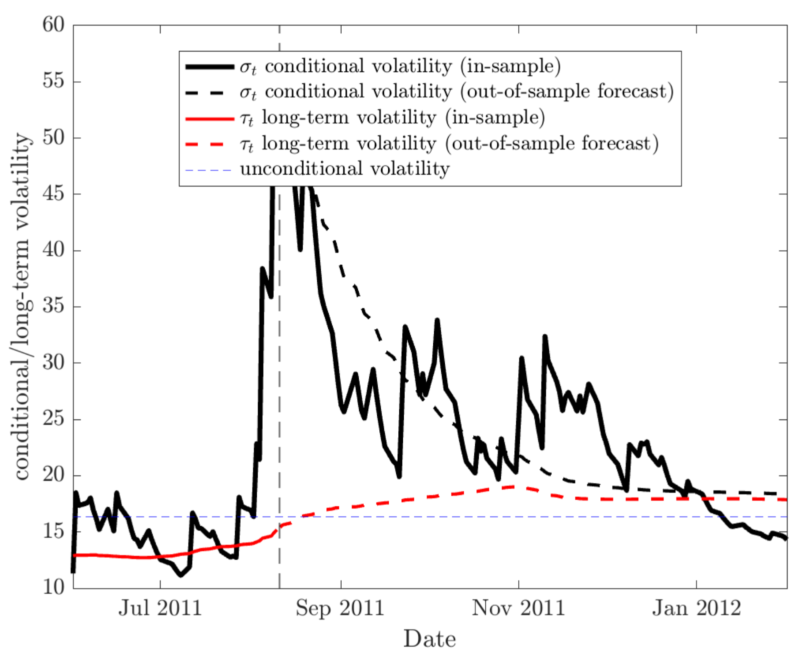

Figure 4: The figure shows the conditional volatility (solid black line) from an MF2-GARCH-rw-63 estimated for S&P 500 returns. From day $t = 0$ (August 10, 2011, indicated by the black vertical line) onwards, we compute volatility forecasts (dashed black line). The figure also shows the long-term components (red line) and the forecast of long-term volatility (dashed red line). All quantities are annualized.

The figure shows the conditional volatility (solid black line) from an MF2-GARCH-rw-m model with $m = 63$ estimated for S&P 500 returns. From August 10, 2011 (indicated by the black vertical line) onwards, we compute volatility forecasts (dashed black line) for 120 days in the future. The plot also shows long-term volatility (red line) and the forecast of long-term volatility (dashed red line). All quantities are annualized. The conditional volatility as well as the long-term volatility are below the unconditional volatility until there is a jump in volatility up to a level above 50%, driven by the European sovereign debt crisis and a downgrade of the U.S.'s credit rating by Standard & Poor's. In the medium run, the forecast for the conditional volatility converges towards the forecast of the long-term volatility (dashed red line). In the very long run, the MF2-GARCH forecast will converge towards the unconditional volatility (blue dashed line).

References

Conrad, Christian, and Robert F. Engle. 2025. "Modelling Volatility Cycles: The MF2-GARCH Model." *Journal of Applied Econometrics*, 40(4): 438–454.

Conrad, Christian, Julius Theodor Schoelkopf, and Nikoleta Tushteva. 2025. "Long-Term Volatility Shapes the Stock Market's Sensitivity to News." *Working Paper*. Available at SSRN: https://ssrn.com/abstract=4632733.

Schwarz, Gideon. 1978. "Estimating the Dimension of a Model." *Annals of Statistics*, 6(2): 461–464.