Modelling Volatility Cycles: The (MF)^2 GARCH model

Christian Conrad[1] and Robert Engle[2]

[1]Heidelberg University
[2]Stern School of Business

Macroeconomic forecasting in abnormal times
11th European Central Bank Conference on Forecasting Techniques

June 16, 2021

UNIVERSITÄT HEIDELBERG
ZUKUNFT SEIT 1386

Financial Times, June 16, 2021

Why is Wall Street's fear gauge so low?

Investors urged to have 'eyes wide open about what is coming next'

© FT montage; Bloomberg

Naomi Rovnick and **Eva Szalay** in London YESTERDAY 69

After inflation fears shocked investors in the first few months of 2021, markets have switched into a different mode: a deep slumber.

- transitory inflation narrative of the Federal Reserve
- "higher downside risk to the news flow than usual"

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Motivation

- strong empirical evidence that the conditional variance of stock returns consists of several components (Ding and Granger, 1996, Engle and Lee, 1999, Corsi, 2009)
- recent GARCH-type models assume multiplicative short- and long-term components: Spline GARCH (Engle and Rangel, 2008), GARCH-MIDAS (Engle et al., 2013)
- we propose a new approach for modelling the long-term component:
  - we document that the **daily** forecast errors of one-component GARCH models are predictable **when averaged at a lower frequency**
    - due to long-term volatility cycles that are not captured by standard GARCH models
  - long-term component: learning from past forecast errors
  - Multiplicative Factor Multi Frequency GARCH – $(MF)^2$ GARCH
    - instantaneous/short-run response to news
    - long-term forecasts
  - Volatility Lab: real-time estimates and forecasts of financial volatility

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

New Stylized Fact: Predictable Volatility Forecast Errors

GJR GARCH(1,1) for daily returns:

$$r_{i,t} = \sqrt{\tilde{h}_{i,t}}\zeta_{i,t},$$

where $\tilde{h}_{i,t}$ denotes the conditional variance and the $\zeta_{i,t}$ are $i.i.d.$ with mean zero and variance one.

- low frequency period $t = 1, \ldots, T$ (monthly or quarterly)
- high frequency period i, n days within in each period t

Misspecification testing:

- testing for an omitted multiplicative volatility component

$$\zeta_{i,t} = \sqrt{\tau_t} Z_{i,t}, \text{ where the } Z_{i,t} \text{ are } i.i.d.$$

- Lundbergh and Teräsvirta (2002): $n = 1$
  Conrad and Schienle (2020): explanatory variables, $n \geq 1$

New Stylized Fact: Predictable Volatility Forecast Errors

Tests are based on the (fitted) squared standardized residuals

$$\zeta_{i,t}^2 = \frac{r_{i,t}^2}{\tilde{h}_{i,t}}$$

What if τ_t varies smoothly over time?

We will focus on

$$\tilde{V}_t = \frac{1}{n} \sum_{i=1}^{n} \frac{r_{i,t}^2}{\tilde{h}_{i,t}}.$$

Interpretation: average of the standardized volatility forecast errors within low-frequency period t.

Example: daily S&P 500 log-returns for the period January 1971 to June 2020.

$$\tilde{h}_{i,t} = \underset{(0.003)}{0.017} + \left(\underset{(0.007)}{0.022} + \underset{(0.014)}{0.110}\ \mathbf{1}_{\{r_{i-1,t}<0\}}\right) r_{i-1,t}^2 + \underset{(0.009)}{0.905}\ \tilde{h}_{i-1,t}$$

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

New Stylized Fact: Predictable Volatility Forecast Errors

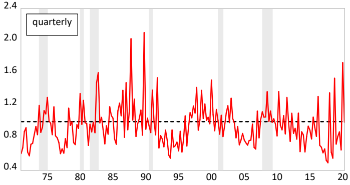

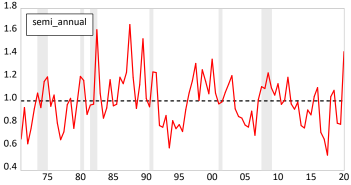

Figure 1: A GJR GARCH(1,1) is estimated for the S&P 500 using daily return data for the January 1971 to June 2020 period. The figure shows $\tilde{V}_t = n^{-1} \sum_{i=1}^{n} r_{i,t}^2 / \tilde{h}_{i,t}$ at a quarterly (left panel) and semi-annual (right panel) frequency. Grey shaded areas represent NBER recession periods.

- low-frequency forecast errors predictable (not *i.i.d.*)
- low-frequency forecast errors are counter-cyclical

New Stylized Fact: Predictable Volatility Forecast Errors

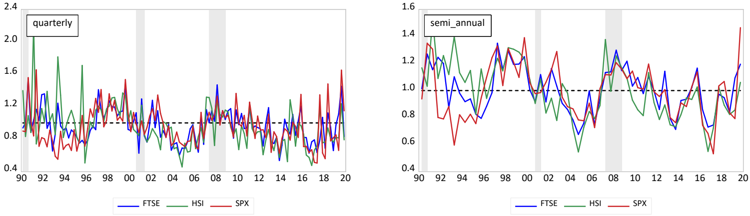

Figure 2: Volatility forecast errors for international stock markets.

- Engle and Campos-Martins (2020): Geopolitical events

$(MF)^2$ GARCH

The $(MF)^2$ GARCH

We focus on the model with daily τ_t (i.e. we assume that $n = 1$).

$$r_t = \sigma_t Z_t = \sqrt{h_t \tau_t} Z_t.$$

Assumption 1
Let Z_t be i.i.d. with $\mathbf{E}[Z_t] = 0$, $\mathbf{E}[Z_t^2] = 1$ and $1 < \kappa < \infty$, where $\kappa = \mathbf{E}[Z_t^4]$.

Short-term volatility component

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

where $\phi = \alpha + \gamma/2 + \beta$.

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

(MF)^2 GARCH

Assumption 2
The parameters of the short-term GJR-GARCH component satisfy the conditions $\alpha > 0$, $\alpha + \gamma > 0$, $\beta > 0$ *and* $\phi = \alpha + \gamma/2 + \beta < 1$.

Hence,

$$\frac{r_t}{\sqrt{\tau_t}} = \sqrt{h_t} Z_t,$$

follows a covariance stationary GJR GARCH with $\mathbf{E}[h_t Z_t^2] = \mathbf{E}[h_t] = 1$.

How to specify τ_t?

Consider the squared deGARCHed returns:

$$V_t = \frac{r_t^2}{h_t} = \tau_t Z_t^2.$$

If the GARCH component fully captures the conditional heteroskedasticity, then the V_t are *i.i.d*.

$(MF)^2$ GARCH

Long-term volatility component

We parameterize the long-term component as

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

where

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

interpretation:

- $V_{t-1}^{(m)}$ as rolling window measure of recent forecast performance of the GARCH component
- MEM$(1, m)$ with the restriction that the m "ARCH" coefficients are given by λ_1/m.
- τ_t as scaling factor: scales volatility up/down if the short-term component underestimated/overestimated volatility in the recent past

(MF)^2 GARCH

Assumption 3
The parameters in the long-term component satisfy the conditions $\lambda_0 > 0$, $\lambda_1 > 0$, $\lambda_2 > 0$ *and* $\lambda_1 + \lambda_2 < 1$.

Under Assumptions 1 and 3 it holds that

$$\begin{aligned} \mathbf{E}[V_t|\mathcal{F}_{t-1}] &= \tau_t \\ \mathbf{E}[V_t] &= \lambda_0/(1-\lambda_1-\lambda_2) \end{aligned}$$

- $V_t = \tau_t Z_t^2$ is a covariance stationary MEM

We refer to this parametrization as (MF)^2 GARCH-rw-m.

- can be estimated by QMLE
- select m by BIC

(MF)^2 GARCH

Table 1: (MF)^2 GARCH models: S&P 500.

τ_t	α	γ	β	λ_0	λ_1	λ_2	LLF	BIC
				GJR GARCH				
const.	0.022*** (0.005)	0.114*** (0.020)	0.902*** (0.014)	1.008*** (0.140)			-15639.02	2.512
				(MF)^2 GARCH-rw-*m*				
$m = 126$	0.011* (0.006)	0.142*** (0.023)	0.868*** (0.018)	0.017 (0.014)	0.097 (0.088)	0.885*** (0.102)	-15591.25	2.506
$m = 63$	0.004 (0.007)	0.158*** (0.021)	0.843*** (0.019)	0.015* (0.008)	0.100** (0.050)	0.885*** (0.057)	-15572.62	**2.503**
$m = 21$	0.002 (0.007)	0.170*** (0.022)	0.817*** (0.026)	0.004** (0.002)	0.034*** (0.013)	0.962*** (0.014)	-15580.74	2.505

Notes: The table reports estimation results for GJR GARCH and (MF)^2 GARCH-rw-*m* models. All models are estimated using daily return data for the period January 1971 to June 2020. The numbers in parentheses are Bollerslev–Wooldridge robust standard errors. ***, ** and * indicate significance at the 1%, 5% and 10% level. LLF is the value of the maximized log-likelihood function. BIC is the Bayesian information criterion.

- Persistence: GJR GARCH - 0.981, (MF)^2 GARCH-rw-63 - 0.926 (short-term component)

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

V-Lab

- (MF)^2 GARCH implemented in V-Lab (https://vlab.stern.nyu.edu/)
- estimated for more than 18,000 assets from different asset classes on a weekly basis
- volatility forecasts for one day up to one year

Table 2: Summary Statistics: Median parameter estimates, persistence and optimal *m*.

	α	γ	β	$\alpha + \gamma/2 + \beta$	$\lambda_1 + \lambda_2$	*m*
	70 international equity indices					
GJR GARCH	0.030	0.110	0.900	0.982	-	-
(MF)^2 GARCH	0.010	0.145	0.842	0.928	0.991	41
	1240 US equities					
GJR GARCH	0.068	0.055	0.864	0.971	-	-
(MF)^2 GARCH	0.095	0.071	0.693	0.841	0.988	66
	3852 international equities					
GJR GARCH	0.083	0.015	0.884	0.974	-	-
(MF)^2 GARCH	0.112	0.029	0.704	0.841	0.984	46

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

V-Lab

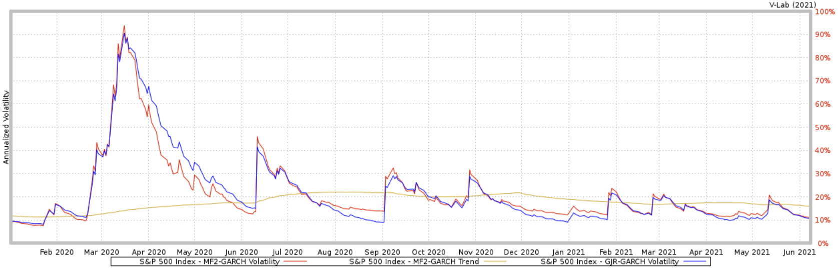

Figure 3: Conditional volatility of S&P 500: $(MF)^2$ GARCH (red: conditional volatility, yellow: long-term component) and GJR GARCH (blue).

Volatility Prediction for Wednesday, June 16th, 2021: 10.17% (-0.04

V-Lab

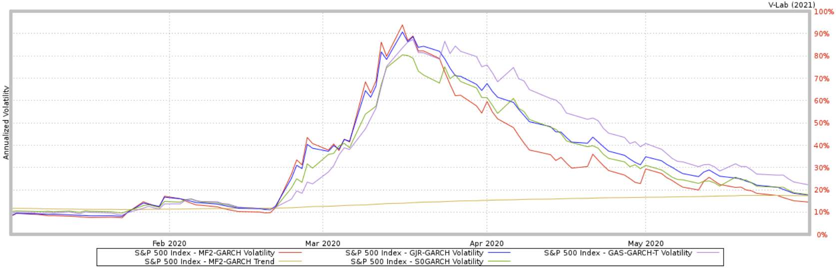

Figure 4: Conditional volatility of S&P 500: (MF)^2 GARCH (red), GJR GARCH (blue), Spline GARCH (green), GAS GARCH (purple).

(MF)^2 GARCH

Unconditional variance

Theorem 1: Let Assumptions 1-3 be satisfied. If the (MF)^2 GARCH-rw-m process, $(r_t)_{t\in\mathbb{Z}}$, is covariance stationary, then

$$\Gamma_m = \left(\lambda_1 \frac{1}{m}\phi_\kappa + \lambda_2\phi\right) + \lambda_1\phi_\kappa \frac{1}{m}\sum_{j=2}^{m}\phi^{j-1} < 1 \quad (2)$$

with $\phi_\kappa = (\alpha + \gamma/2)\kappa + \beta$. The unconditional variance of the daily returns is given by

$$\mathbf{Var}[r_t] = \frac{\lambda_0 + \frac{\lambda_0}{1-\lambda_1-\lambda_2}(1-\phi)(\lambda_1+\lambda_2) + \Delta_m}{1-\Gamma_m}, \quad (3)$$

where

$$\Delta_m = (1-\phi)\lambda_1\phi\frac{\lambda_0}{1-\lambda_1-\lambda_2}\left(\frac{m-1}{m} + \frac{1}{m}\sum_{j=2}^{m}\sum_{k=1}^{j-2}\phi^k\right).$$

- feedback: short- and long-term component are correlated
- $\mathbf{Var}[r_t]$ depends on κ ($\nearrow$) and m ($\searrow$)

(MF)^2 GARCH

News impact function

NIC for GJR GARCH

$$NIC_{t+1}^{GA} = \tilde{h}_{t+1}(r_t|\tilde{h}_t) = A_t^{GA} + (\alpha + \gamma \mathbf{1}_{\{r_t<0\}})r_t^2,$$

Effect of r_t does not depend on current level of volatility!

At odds with empirical observation for options implied volatility:

> *"..., if VIX indicates the market is going to be volatile it may not react too much when the market is volatile. And if VIX indicates the market is expected to be less volatile, if we do get high volatility, VIX adjusts accordingly." (Rhoades, 2020, p.60)*

NIC for (MF)^2 GARCH: the effect of r_t depends on σ_t^2, h_t and τ_t

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

(MF)^2 GARCH

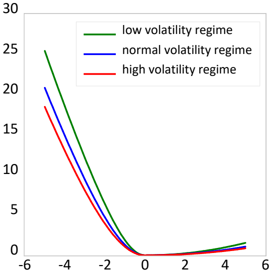

Figure 5: NIF for (MF)^2 GARCH-rw-63 with parameters as before. Short-term component fixed at unconditional expectation ($h_t = 1$). We plot the NIC as a function of the return, r_t, and for $\tau_t = 0.5$ (green), $\tau_t = 1$ (blue) and $\tau_t = 1.5$ (red). NIC's are standardized such that news impact is zero for $r_t = 0$ and presented as annualized volatilities, i.e. we plot $\sqrt{252(\sigma^2_{t+1}(r_t|\tau_t, h_t = 1) - \sigma^2_{t+1}(r_t = 0|\tau_t, h_t = 1))}$.

(MF)^2 GARCH

Forecasting Volatility

- short-term component of (MF)^2 GARCH less persistent than one-component GJR GARCH

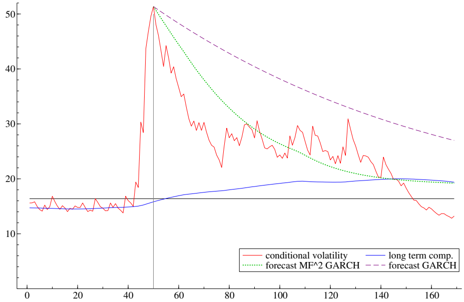

Figure 6: Annualized volatility forecasts from (MF)^2 GARCH-rw-63 (green dotted line) and GJR GARCH (purple dashed line) for S&P 500. $t = 50$ corresponds to August 12, 2011.

(MF)^2 GARCH

Forecasting Volatility

- in the medium term, the (MF)^2 GARCH forecast approaches the forecast of the long-term component

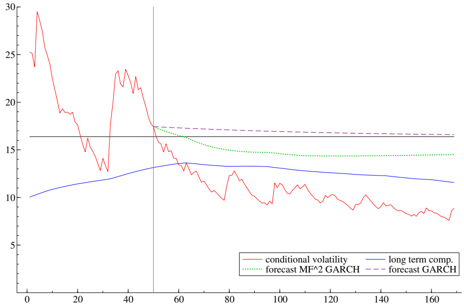

Figure 7: Annualized volatility forecasts from (MF)^2 GARCH-rw-63 (green dotted line) and GJR GARCH (purple dashed line) for S&P 500. $t = 50$ corresponds to April 16, 2018.

(MF)^2 GARCH

Extension: (MF)^2 GARCH-rw-m with Beta-weights

Imposing equal weights in $V_t^{(m)}$ might be too restrictive. Alternatively, we can consider a weighted average of the form

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

We parsimoniously model the weights $w_j(\omega)$ according to a restricted Beta weighting scheme:

$$w_j(\omega) = \frac{(1 - j/(m+1))^{\omega-1}}{\sum_{k=1}^{m} (1 - j/(m+1))^{\omega-1}}.$$

Refer to this parametrization as (MF)^2 GARCH-bw-m.

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

(MF)^2 GARCH

Relation to other multiplicative component models

- GARCH-type for volatility
  - Spline-GARCH
    - returns are non-stationary, long-term forecasting?
  - GARCH-MIDAS
    - Conrad and Loch (2015), Conrad and Kleen (2020)
      - long-term: term spread, housing starts, SPF expectations for output growth, corporate profits, unemployment rate
      - short-term: NFCI, VIX
    - how to select relevant variables?
    - not dynamically complete
  - no feedback from short- to long-term component
- intraday volume: Composite-MEM (Brownlees et al., 2011)
- realized volatilities: Component-MEM (Amendola et al., 2020)

News Effects

Instantaneous and intermediate effects of major news events

- Baker et al. (2019) consider days with returns of at least +/- 2.5%
- cause of large return is determined based on newspaper coverage (Wall Street Journal, New York Times, Chicago Tribune, Washington Post, LA Times) the following day
- human readers assign (primary/secondary) categorical cause
- indictors for 16 categories (between zero and one) based on average reader assignments
- we focus on the news categories: 'Macroeconomic News & Outlook', 'Corporate Earnings & Outlook' and 'Monetary Policy & Central Banking'

$$\begin{aligned}
\sqrt{\frac{1}{S}\sum_{s=1}^{S}\sigma_{t+s}^2} \quad = \quad & b_1 + b_2 \text{Macro}_t \times \mathbf{1}_{(r_t>0)} + b_3 \text{Macro}_t \times \mathbf{1}_{(r_t<0)} \\
& + b_4 \text{Corp. Earn.}_t \times \mathbf{1}_{(r_t>0)} + b_5 \text{Corp. Earn.}_t \times \mathbf{1}_{(r_t<0)} \\
& + b_6 \text{Mon. Pol.}_t \times \mathbf{1}_{(r_t>0)} + b_7 \text{Mon. Pol.}_t \times \mathbf{1}_{(r_t>0)} \\
& + b_8 \sqrt{\sigma_t^2} + e_t
\end{aligned}$$

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

News Effects

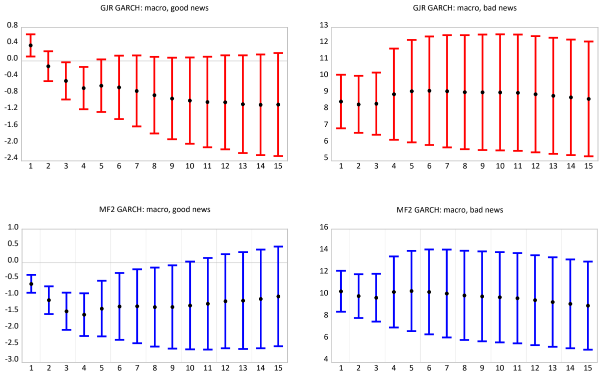

Figure 8: Effect of good (left panels)/bad (right panels) macroeconomic news on the conditional volatility of the GJR GARCH (upper panels) and $(MF)^2$ GJR-rw-63 (lower panels). The effect is plotted for 1 up to 15 days after the news event.

News Effects

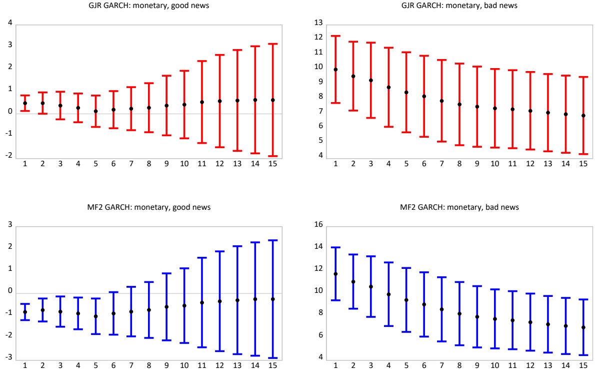

Figure 9: Effect of good (left panels)/bad (right panels) news about monetary policy on the conditional volatility of the GJR GARCH (upper panels) and (MF)^2 GJR-rw-63 (lower panels). The effect is plotted for 1 up to 15 days after the news event.

Forecasting

Data

- S&P 500: January 01, 1970 – June 30, 2020
- FTSE 100 and HSI: March 01, 1990 – June 30, 2020
- in-sample / out-of-sample:
  - monthly rolling window estimation
  - out-of-sample period: 01/01/2010 to 06/30/2020
- forecast evaluation based on MSE loss
- *RV* is realized variance (based on 5-minute intra-day returns plus overnight return)

Table 3: Out-of-sample forecasting: S&P 500 – relative RMSE.

	forecast horizon				
	daily	monthly	quarterly	semi-annual	annual
(MF)^2 GARCH-rw-126	0.751	0.693	0.672	0.657	0.688
(MF)^2 GARCH-rw-63	0.741	0.684	0.692	0.687	0.705
(MF)^2 GARCH-rw-21	0.734	0.684	0.703	0.701	0.715
(MF)^2 GARCH-bw-252	0.737	0.684	0.703	0.700	0.714
(MF)^2 GARCH-bw-126	0.737	0.684	0.702	0.699	0.713
GJR	0.782	0.788	0.806	0.810	0.771
HAR	0.771	0.711	0.759	0.800	0.875
log-HAR	0.760	0.687	0.740	0.781	0.846
log-HAR leverage	0.756	0.713	0.729	0.784	1.123
Random Walk Typ 1	0.907	1.278	1.520	1.730	1.761
Random Walk Typ 2	0.907	0.771	0.857	0.858	0.793

Notes: Numbers reported are the RMSE losses for each model and forecast horizon relative to the respective RMSE of historical volatility. Red-colored numbers indicate the model with the lowest RMSE. Blue-shaded numbers indicate that the respective model is included in the 85% model confidence set.

What is coming next?

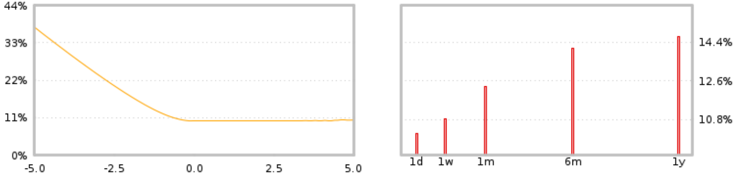

Figure 10: VLab: NIC and volatility forecast for S&P 500, June 16th, 2021

Conclusions

- (MF)^2 GARCH
  - exploits predictability in standardized volatility forecast errors
  - feedback from short- to long-term component
  - properties: stationary returns, news impact depends on current volatility, dynamically complete, simple to estimate by QMLE
- empirical evidence: (MF)^2 GARCH ...
  - is preferred to nested GJR GARCH
  - response to extreme events is more realistic than in one-component GARCH
  - outperforms standard models at intermediate and long forecast horizons
- applications
  - long-run value at risk (SRISK), portfolio choice