Lecture notes on risk management, public policy, and the financial system

Volatility behavior and forecasting

Allan M. Malz

Columbia University

Time variation in return volatility and correlation

Volatility forecasting

Volatility forecasts

- A major departure from standard model: *risk* or *volatility changes over time*
- Volatility, unlike return, not directly observable, must be estimated
    - Challenge: method for estimating volatility that captures typical patterns of volatility
- Recent past and long-term volatility help predict future volatility
    - But: while estimators efficacious for forecasting near-term volatility, they often miss sharp changes in volatility
- **Second-moment efficiency**: option market does less-poor job forecasting return variance than forward markets of forecasting mean return

Typical patterns of volatility behavior over time

Persistence: volatility tends to stay near its current level

- Periods of high or low volatility tend to be enduring
- Once a large-magnitude return shocks volatility higher, volatility persists at its higher level
- Magnitude or square of return as well as return volatility display positive **autocorrelation**

Abrupt changes in volatility are not unusual

- Together with persistence, leads to **volatility clustering** or **volatility regimes**
- Shifts from low to high volatility are more abrupt, while shifts from high to low volatility are more gradual

Long-term mean reversion: volatility of an asset's return tends to gravitate to a long-term level

- In turn implies a **term structure of volatility:** different current estimates of volatility for different time horizons

Volatility of oil prices 1986-2018

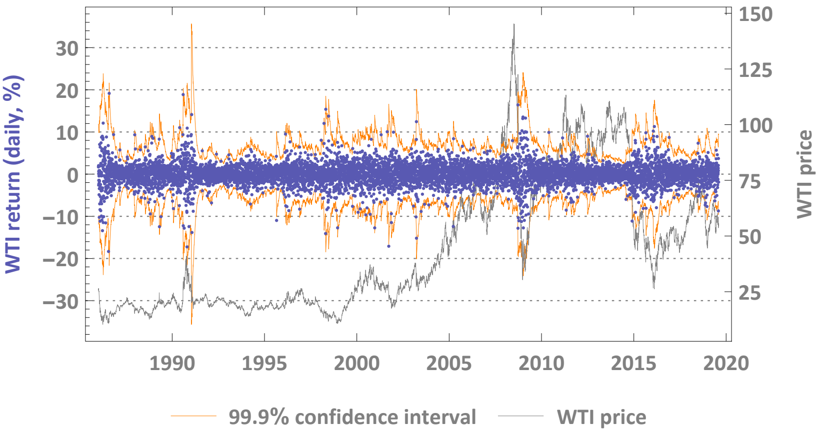

Prices and returns of Cushing, OK West Texas Intermediate (WTI) crude oil. 99.9 percent confidence interval calculated using daily EWMA volatility estimate (decay factor $\lambda = 0.94$). Daily, 02Jan1986–16Jan2018. *Data source*: U.S. Energy Information Administration (`https://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm`).

Conditional volatility

- Volatility regimes suggest use of **conditional volatility:** estimate weighted toward more recent information
- Formally, volatility forecasts based on some information ("shocks" or "innovations") up to present time t
    - $\sigma_t \equiv$ current estimate of future return volatility based on (a model and) information through time t
- What new information drives σ_t? In most models:
    - Magnitude (and possibly the sign) of recent *returns*
    - Recent estimates of *volatility*
- Term structure of volatility, e.g. weekly volatility higher or lower than daily
    - Typically, volatility expected to rise (fall) when low (high) relative to long-term average level

Impact of time-varying correlation

- Like volatility, correlations vary over time
- Correlations have strong impact on portfolio returns, hedged positions
- Abrupt changes in correlation during periods of financial stress $\rightarrow$
  - Failure of hedging strategies
  - Diminution of diversification benefits
- "Risk-on risk-off" behavior: tendency for correlations across many assets to rise when risk appetites diminish in stress periods
- **Examples**:
  - Increase in return correlations among equities
  - Higher correlation between equity returns, Treasury yields

Time-varying correlation of stock and bond returns

- Persistent changes over time in general level of correlation between of stock and bond returns
  - 1960s-1990s: generally negative
  - 1990s-: generally positive
- Experience during the inflation and disinflation from late 1960s
  - Rising rates driven by rising inflation expectations, associated with adverse impact on economic growth
  - "Fed model": increase in discount rate for future earnings reduces present value
- Experience once low-inflation monetary policy fully credible
  - Rising rates driven by increases in anticipated real returns ($r*$), associated with poitive impact on economic growth
  - Risk-on risk-off: investors reduce allocations to risky in favor of safe assets
- N.B. positive correlation of equity returns and yield changes corresponds to negative correlation of stock and bond returns

Correlation of stock returns and rates 1962-2020

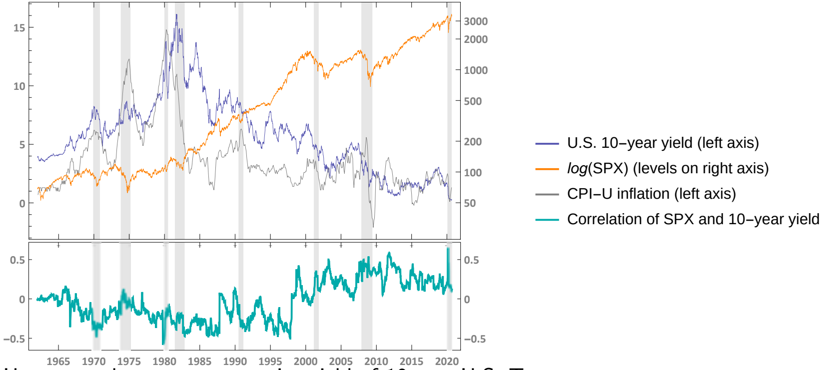

Upper panel: constant-maturity yield of 10-year U.S. Treasury note, year-over-year CPI-U inflation, log of S&P 500 index. Lower panel: correlation between log S&P price return and changes in 10-year yield (estimated via EWMA with decay factor $\lambda = 0.97$). Vertical shading represents NBER recession dates. Weekly data, 05Jan1962 to 11Sep2020. *Data source*: Bloomberg L.P.

Time variation in return volatility and correlation

Volatility forecasting

- Simple approaches conditional volatility estimation
- GARCH
- The exponentially-weighted moving average model

Using conditional volatility estimators

- General approach: revise most recent estimate of volatility based on most recent return data
- Simplified notation when working with daily data: return from yesterday's to today's close

$$r_t \equiv r_{t-1,t} \equiv \ln(S_t) - \ln(S_{t-1})$$

- At close of each day t, use r_t to update yesterday's volatility estimate σ_{t-1}
- Use the new estimate σ_t to measure risk or forecast volatility over the next business day $t+1$
- Volatility forecast horizon includes *trading days*, not calendar time
  - Price can change only when market open
  - $\not\leftrightarrow$Holding period and cash flows accrue every *calendar* day

Volatility is easier to estimate than mean

- Imagine asset return approximately follows diffusion with drift
  - Observed at regular intervals over a period of time
  - Drift and volatility may change over time, but slowly
- You only observe *one* sample path in real history
- The only information on mean/drift is return over entire period
- But finer intervals—every 5 min. instead of daily—provide more information on volatility
  - Finer intervals provide more information on tendency to wander
  - Confidence interval of volatility estimate$\rightarrow$ 0
  - But not confidence interval of mean estimate
- Tail risk very hard to estimate

Zero-mean assumption

- A typical risk-measurement modeling choice:
  - *Estimate* return volatility
  - But *assume* mean return $= 0$
- In lognormal model, assume drift $\mu = 0 \Rightarrow$:
  - Mean logarithmic return $\mu = 0$

$$r_{t,t+\tau} = \ln(S_{t+\tau}) - \ln(S_t) \sim \mathcal{N}(0, \sigma^2 \tau)$$

  - But discrete returns have non-zero mean due to Jensen's Inequality term:

$$\mathbf{E}\left[S_{t+\tau}\right] = S_t e^{\frac{1}{2}\sigma^2 \tau}$$

Why assume zero-mean returns?

- Because we can:
  - Small impact of mean on volatility over short intervals
  - But: mean return increases linearly with time, return volatility increases as square root of time
  - $\Rightarrow$Over longer periods, mean has larger impact than volatility
- Because we must:
  - Expected return very hard to measure
  - Estimation of mean introduces additional source of statistical error into variance estimate
  - Bad enough to assume return normality, let's not also invent mean

Square-root-of-time rule

- In standard (→)geometric Brownian motion model, *variance* (vol squared) of price change proportional to time elapsed
  - Position after t time units

$$S_t \sim \mathcal{N}(0, t)$$

  - Together with martingale property

$$S_{t+\tau} - S_t \sim \mathcal{N}(0, \tau)$$

- Carries over to standard lognormal/geometric Brownian motion model: variance increases in proportion to time elapsed
  - ⇒Vol increases in proportion to square root of time elapsed
- Useful rule-of-thumb even if returns only approximately lognormal
  - But assumes constant return volatility, i.e. flat term structure of volatility
  - At odds with changes in volatility over time and with long-term mean reversion

Applying the square-root-of-time rule

- Volatility forecast horizon includes *trading days*, not calendar time
  - Typical year includes about 250-255 trading days
  - Assume 256 trading days, $\sqrt{256} = 16$
  - Annualized volatility $\approx 16 \times$daily volatility
- **Examples**:
  - Long-term average annual volatility of U.S. stock indexes $\approx 16 - 20$ percent$\Rightarrow$daily vol $\approx 1 - 1.25$ percent
  - Swaption **normal volatility** 80 bps$\Rightarrow$daily vol 5 bps

Simple conditional volatility estimators

Use moving window incorporating past m trading days' returns

Root mean square: square root of the sum of squared returns (deviations from zero) divided by the number of observations

$$\sigma_t = \sqrt{\frac{1}{m}\sum_{\tau=1}^{m} r_{t-m+\tau}^2}$$

- Incorporates assumption of zero mean return

Standard deviation: the square root of the sum of squared deviations from the mean return $\bar{r}_t = \frac{1}{m}\sum_{\tau=1}^{m} r_{t-m+\tau}$ divided by the number of observations minus 1

$$\sigma_t = \sqrt{\frac{1}{m-1}\sum_{\tau=1}^{m}(r_{t-m+\tau} - \bar{r}_t)^2},$$

- Bias-corrected for 1 degree of freedom lost due to use of $\bar{r}_t$

GARCH model of volatility

- **Generalized autoregressive conditionally heteroscedastic** model
- Volatility driven by
  - Recent volatility
  - Recent returns
  - Long-term "point of rest" of volatility or "forever vol" $\bar{\sigma}$
- Estimate σ_t made at today's close updates yesterday's estimate σ_{t-1} with latest return r_t
- Look back one period→**GARCH(1,1):**

$$\sigma_t^2 = \alpha r_t^2 + \beta \sigma_{t-1}^2 + \gamma \bar{\sigma}^2$$

- Feedback to returns via "shock" or "innovation" ϵ_t

$$r_t = \epsilon_t \sigma_{t-1},$$

- Today's return r_t the only pertinent new information on date t
  - ϵ_t assumed i.i.d. with mean 0 and variance 1
  - ϵ_t together with current volatility σ_{t-1} determines new return r_t
  - r_t random but not "free," set by current vol and random shock ϵ_t
- The weights satisfy $\alpha, \beta, \gamma > 0$ and $\alpha + \beta + \gamma = 1$

Role of parameters in the GARCH model

- Impact of α: high $\alpha \Rightarrow$
    - Large r_t causes large, immediate change in estimated return volatility σ_t
    - Wider range of variation of σ_t over time
- Impact of β: high $\beta \Rightarrow$
    - σ_t and deviations from $\bar{\sigma}^2$ very persistent
    - Less variation of σ_t over time
- Long-term variance $\bar{\sigma}^2 > 0$
    - Presence of $\bar{\sigma}^2$ generates a term structure of volatility
    - **Example:** $\bar{\sigma}^2$ approximately 1.0–1.15 percent for U.S. equity market (at daily rate)
- Low $\gamma \Rightarrow$ little mean reversion
- Estimates of β generally not very far from 1, $\alpha + \beta$ quite close to 1
- Estimated parameter values lead to (hopefully realistic) behavior of volatility

Estimating GARCH(1,1) model parameters

- **Maximum likelihood method** a standard approach
  - Assume **conditional normality**, shocks ϵ_t normally distributed, a stronger assumption than i.i.d.:

$$\epsilon_t \sim \mathcal{N}(0,1) \quad \forall t$$

  - Joint normal density of m return observations⇒**log likelihood function**

$$\sum_{\tau=1}^{m} \left[-\ln(\sigma^2_{t-\tau}) - \frac{r_t^2}{\sigma^2_{t-\tau}} \right],$$

    with initial volatility value σ_0
- Use numerical search procedure to find parameters that maximize log likelihood function
  - Numerical search procedure can be sensitive to initial trial guess
  - $\omega \equiv \gamma \bar{\sigma}^2$ treated as a single parameter
  - γ can then be recovered as $1 - \alpha - \beta$ and

$$\bar{\sigma} = \sqrt{\frac{\omega}{1-\alpha-\beta}}$$

Influence of past returns in GARCH model

- GARCH(1,1) formula can be recast in terms of most recent and past squared returns (setting $m = t$):

$$
\begin{aligned}
\sigma_1^2 &= \alpha r_1^2 + \beta \sigma_0^2 + \omega \\
\sigma_2^2 &= \alpha r_2^2 + \beta \sigma_1^2 + \omega = \alpha r_2^2 + \beta(\alpha r_1^2 + \beta \sigma_0^2 + \omega) + \omega \\
&= \alpha r_2^2 + \alpha\beta r_1^2 + \beta^2 \sigma_0^2 + (1+\beta)\omega \\
&\vdots \\
\sigma_t^2 &= \alpha \Sigma_{\tau=1}^{t} \beta^{t-\tau} r_\tau^2 + \Sigma_{\tau=1}^{t} \beta^{t-\tau} \omega + \beta^t \sigma_0^2 \\
&\approx \alpha \Sigma_{\tau=1}^{t} \beta^{t-\tau} r_\tau^2 + \frac{1}{1-\beta}\omega
\end{aligned}
$$

- $\alpha < 1, \beta < 1 \Rightarrow$ small influence of more remote past returns, starting value σ_0
- Tradeoff bet influence of long-term variance and that of most recent volatility estimate

Example of GARCH(1,1) model estimation

- Applied to S&P 500 index, using $m+1=3651$ closing-price observations 30Jun2005 to 31Dec2019
- r_1^2 used as starting value σ_0
  - Can also use sample variance of entire time series
- For each pass of the search procedure, successively apply GARCH(1,1) formula to calculate trial values $\sigma_1, \sigma_2, \ldots, \sigma_t$:

Parameter estimates			
α	0.12195	β	0.85609
ω	2.40805×10^{-6}	$\bar{\sigma}$ (%)	1.04715
γ	0.02196		

- Practical application: (re-)estimate parameters infrequently, but use estimated model regularly to forecast volatility

Exponentially-weighted moving average model

- **Exponentially-weighted moving average** (EWMA)
  - A.k.a. **RiskMetrics model**
  - Variance a weighted average of past returns
  - Weights smaller for more-remote past returns
- Single parameter: **decay factor** λ
  - Low λ: rapid adaptation to recent returns
  - High λ: slow adaptation to recent returns
- EWMA implies a flat term structure of volatility
  - Volatility follows square-root-of-time rule
  - Volatility behaves as a random walk, subject to shocks
- Decay factor estimation: λ that minimizes forecast errors, e.g. RMS criterion
- Decay factor may also be chosen judgmentally

Estimating volatility with the EWMA model

- Typically, assume a value for parameter λ rather than estimate it, and apply a formula
- Current volatility estimate σ_t uses m most recent observed returns $r_{t-m+1}, \ldots, r_t$
  - Treat λ as known parameter
  - Weight on each squared return $\frac{1-\lambda}{1-\lambda^m}\lambda^{m-\tau}, \tau = 1, \ldots, m$
    - Apply $\lambda^{m-m} = 1$ for $\tau = m$, most recent (time t) return
    - Apply $\lambda^{m-1} \approx 0$ for $\tau = 1$, most remote (time $t - m + 1$) return

$$\boxed{\sigma_t^2 = \frac{1-\lambda}{1-\lambda^m} \sum_{\tau=1}^{m} \lambda^{m-\tau} r_{t-m+\tau}^2}$$

  - $1 - \lambda^m \approx 1 \Rightarrow$

$$\sigma_t^2 \approx (1-\lambda) \sum_{\tau=1}^{m} \lambda^{m-\tau} r_{t-m+\tau}^2$$

- m doesn't have to be large
  - $m \approx 100$ more than adequate unless λ quite close to 1

The EWMA model weighting scheme

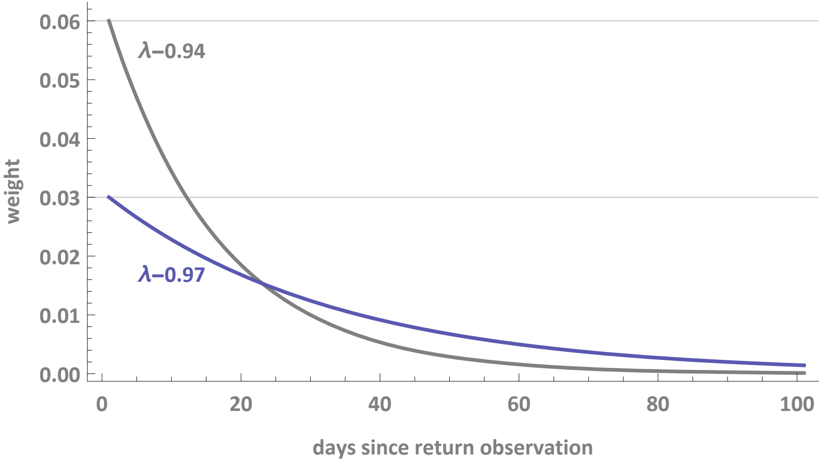

The graph displays the values of the last 100 of $m = 250$ EWMA weights $\frac{1-\lambda}{1-\lambda^m}\lambda^{m-\tau}$ for $\lambda = 0.94$ and $\lambda = 0.97$.

Choosing the decay factor

- Low λ $\Rightarrow$recent observations have greater weight:
  - Volatility changes rapidly
  - Recent observations have most information useful for short-term conditional volatility forecasting
- Low λ effectively shortens historical sample size compared to high λ
- Estimates using low λ much more variable than those using high λ
- Estimates using low λ respond more rapidly to new information
- Estimates using low λ may move in the opposite direction from those using high λ
  - Estimates using low λ decline after a sequence of high-magnitude returns, while those using high λ still rising in response
- No agreed method for estimating λ
- Widely adopted standard settings for decay factor:
  - $\lambda = 0.94$ for short-term (e.g. one-day) forecasts
  - $\lambda = 0.97$ for medium-term (e.g. one-month) forecasts
  - Minimizes RMS of forecast errors for range of assets in original 1994 RiskMetrics study

Effect of the decay factor on the volatility forecast

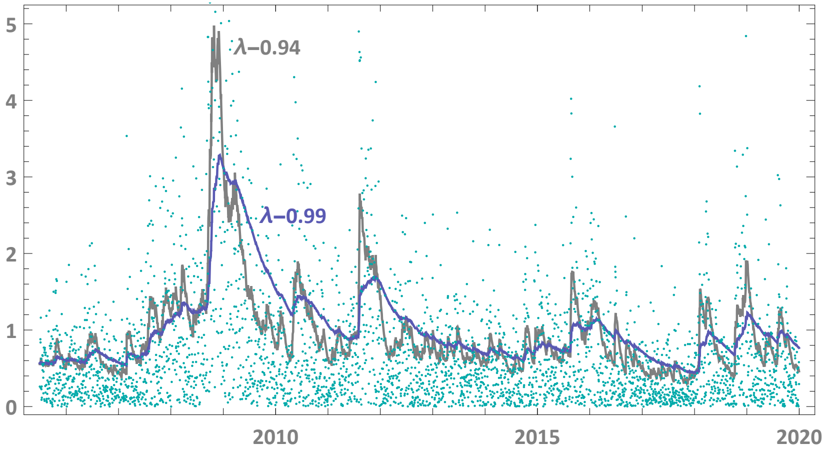

EWMA estimates of the volatility of daily S&P 500 index returns 01Jul2005 to 31Dec2019, at a daily rate in percent, using decay factors of $\lambda = 0.94$ and $\lambda = 0.99$. Points represent the absolute value of daily return observations.

Estimating volatility with the EWMA model

τ	Date	$S_{t+\tau-m}$	$r_{t+\tau-m}$	$\frac{1-\lambda}{1-\lambda^m}\lambda^{m-\tau}$	$\frac{1-\lambda}{1-\lambda^m}\lambda^{m-\tau} r^2_{t+\tau-m}$
0	21Jul2014	1973.63	NA	NA	NA
1	22Jul2014	1983.53	0.00500	0.00000	0.00000×10^{-6}
2	23Jul2014	1987.01	0.00175	0.00000	0.00000×10^{-6}
3	24Jul2014	1987.98	0.00049	0.00000	0.00000×10^{-6}
4	25Jul2014	1978.34	-0.00486	0.00000	0.00000×10^{-6}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
173	27Mar2015	2061.02	0.00237	0.00051	0.00286×10^{-6}
174	30Mar2015	2086.24	0.01216	0.00054	0.08052×10^{-6}
175	31Mar2015	2067.89	-0.00883	0.00058	0.04520×10^{-6}
176	01Apr2015	2059.69	-0.00397	0.00062	0.00973×10^{-6}
177	02Apr2015	2066.96	0.00352	0.00066	0.00814×10^{-6}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
246	13Jul2015	2099.60	0.01101	0.04684	5.67368×10^{-6}
247	14Jul2015	2108.95	0.00444	0.04984	0.98391×10^{-6}
248	15Jul2015	2107.40	-0.00074	0.05302	0.02866×10^{-6}
249	16Jul2015	2124.29	0.00798	0.05640	3.59398×10^{-6}
250	17Jul2015	2126.64	0.00111	0.06000	0.07335×10^{-6}

Return vol of the S&P 500 index, estimated after the close on 17Jul2015 (date t), with $m = 250$, $\lambda = 0.94$. Return (4th column) expressed as a decimal. Add the 250 values in the last column to get the estimated variance σ_t^2.

Recursive formula for EWMA volatility estimates

- Recursive formula updates most recent volatility estimate with new data on return magnitude

$$\boxed{\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)r_t^2}$$

  - $\rightarrow$Easy computation technique, very close to result of full EWMA weighting scheme
- Shows similarity of EWMA to "one-parameter" GARCH
  - But with long-term volatility term $\gamma = 0, \alpha + \beta = 1$
  - λ analogous to β, $1 - \lambda$ analogous to α
  - Shocks to volatility permanent, no long-term "forever" vol
  - Also known as **integrated GARCH** or IGARCH(1,1)
- EWMA estimate usually close to unrestricted GARCH(1,1) estimate
- "Starter value" (orange in example on next slide):
  - Root mean square, using first 21 days of data
  - Starter value method not crucial, converges quickly (esp. for low λ)

Recursive formula for EWMA volatility estimates

t	Date	S_t	r_t (%)	$\lambda\sigma^2_{t-1}$	$(1-\lambda)r^2_t$	σ_t (%)
1	30Jun2005	1191.33	NA	NA	NA	0.55583
2	01Jul2005	1194.44	0.2607	0.29041×10^{-4}	0.40783×10^{-6}	0.54267
3	05Jul2005	1204.99	0.8794	0.27682×10^{-4}	4.63987×10^{-6}	0.56853
4	06Jul2005	1194.94	-0.8375	0.30383×10^{-4}	4.20873×10^{-6}	0.58815
5	07Jul2005	1197.87	0.2449	0.32516×10^{-4}	0.35986×10^{-6}	0.57338
6	08Jul2005	1211.86	1.1611	0.30903×10^{-4}	8.08946×10^{-6}	0.62444
7	11Jul2005	1219.44	0.6235	0.36653×10^{-4}	2.33279×10^{-6}	0.62439
⋮	⋮	⋮	⋮	⋮	⋮	⋮
3644	19Dec2019	3205.37	0.4449	0.24206×10^{-4}	1.18778×10^{-6}	0.50392
3645	20Dec2019	3221.22	0.4933	0.23870×10^{-4}	1.45986×10^{-6}	0.50329
3646	23Dec2019	3224.01	0.0866	0.23810×10^{-4}	0.04497×10^{-6}	0.48842
3647	24Dec2019	3223.38	-0.0195	0.22424×10^{-4}	0.00229×10^{-6}	0.47356
3648	26Dec2019	3239.91	0.5115	0.21080×10^{-4}	1.56983×10^{-6}	0.47592
3649	27Dec2019	3240.02	0.0034	0.21291×10^{-4}	0.00007×10^{-6}	0.46142
3650	30Dec2019	3221.29	-0.5798	0.20014×10^{-4}	2.01673×10^{-6}	0.46937
3651	31Dec2019	3230.78	0.2942	0.20709×10^{-4}	0.51921×10^{-6}	0.46074

Return vol of the S&P 500 index, estimated daily using the recursive formula, with $\lambda = 0.94$. Initial vol estimate: RMS of the 20 daily returns 01Jul2005–29Jul2005.

GARCH(1,1) and EWMA volatility estimates

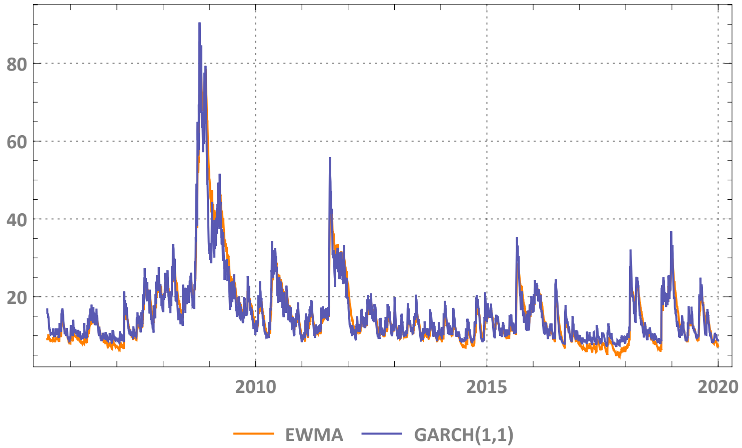

Daily estimates of S&P 500 index's annualized return volatility, 30Jun2005 to 31Dec2019. EWMA estimates with $\lambda = 0.94$ GARCH(1,1) estimates use parameters $\alpha = 0.12195, \beta = 0.85609, \gamma\bar{\sigma}^2 = 2.40805 \times 10^{-6}$. The annualized realized return volatility was 15.69 percent over the period.