

# CFM-Imperial Distinguished Lecture Series

## The Volatility Surface

### Lecture 3: Rough volatility and the connection between historical and implied volatility

Jim Gatheral

Department of Mathematics



The City University of New York



### Outline of lecture 3

- The time series of historical volatility
  - Scaling properties
- The RFSV model
- Forecasting realized variance
- The time series of variance swaps
- Relating historical and implied

### The time series of realized variance

- Assuming an underlying variance process  $v_s$ , integrated variance  $\frac{1}{\delta} \int_t^{t+\delta} v_s ds$  may (in principle) be estimated arbitrarily accurately given enough price data.
  - In practice, market microstructure noise makes estimation harder at very high frequency.
  - Sophisticated estimators of integrated variance have been developed to adjust for market microstructure noise. See Gatheral and Oomen <sup>[10]</sup> (for example) for details of these.
- The Oxford-Man Institute of Quantitative Finance makes historical realized variance (RV) estimates freely available at <http://realized.oxford-man.ox.ac.uk>. These estimates are updated daily.
  - Each day, for 21 different indices, all trades and quotes are used to estimate realized (or integrated) variance over the trading day from open to close.
- Using daily RV estimates as proxies for instantaneous variance, we may investigate the time series properties of  $v_t$  empirically.

```
In [1]: %load_ext rpy2.ipynthon
```

First update and save the latest Oxford-Man data:

```
In [6]: %%R
download.file(url="http://realized.oxford-man.ox.ac.uk/media/1366/oxfordmanrealizedvolatilityindices.zip", destfile="oxfordRvData.zip")
unzip(zipfile="oxfordRvData.zip")

trying URL 'http://realized.oxford-man.ox.ac.uk/media/1366/oxfordmanrealizedvolatilityindices.zip'
Content type 'application/x-zip-compressed' length 10775610 bytes (10.3 Mb)
opened URL
=====
downloaded 10.3 Mb
```

There are many different estimates of realized variance, all of them very similar. We will use the realized kernel estimates denoted by ".rk".

```
In [7]: %%R
library(xts)

rv.data <- read.csv("OxfordManRealizedVolatilityIndices.csv")
colnums <- which(sapply(rv.data, function(x) grep(".rk",x))>0)
col.names <- names(colnums)

rv1 <- rv.data[,colnums]
index.names <- rv1[2,]

datesRaw <- rv.data[-(1:2),1]
dates <- strptime(datesRaw,"%Y%m%d")

rv.list <- NULL
index.names <- as.matrix(index.names)

n <- length(index.names)
for (i in 1:n){
  tmp.krv1 <- xts(rv1[-(1:2),i],order.by=dates)
  rv.list[[i]] <- tmp.krv1[(tmp.krv1!="")&(tmp.krv1!="0")]
}

names(rv.list)<- index.names

save(rv.list, file="oxfordRV.rData")
```

Loading required package: zoo

Attaching package: 'zoo'

The following objects are masked from 'package:base':

as.Date, as.Date.numeric

Let's plot SPX realized variance.

```
In [8]: %%R
# Load Oxford-Man KRV data
load("oxfordRV.rData")
names(rv.list)

spx.rk <- rv.list[["SPX2.rk"]]
plot(spx.rk, main="SPX realized variance", plot=NULL)
lines(spx.rk, col="red")
```

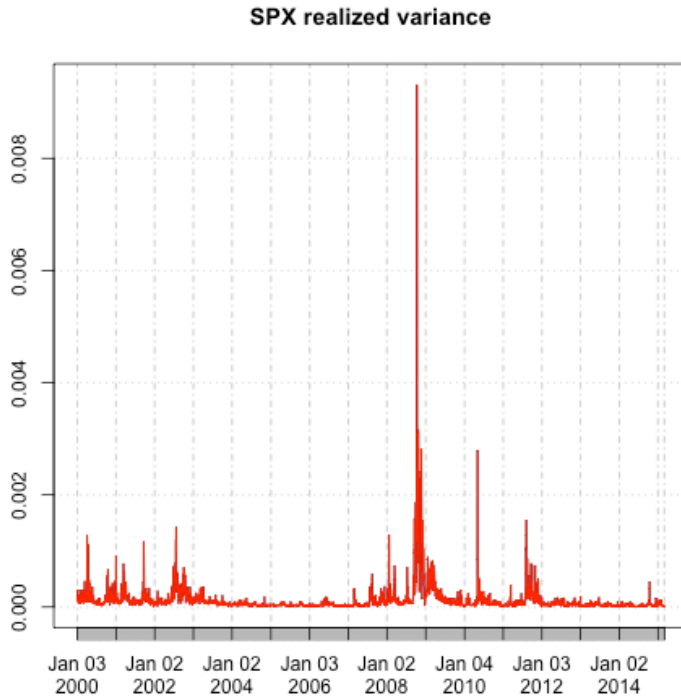


Figure 1: Oxford-Man KRV estimates of SPX realized variance from January 2000 to the current date.

```
In [9]: %%R
print(head(spx.rk))
print(tail(spx.rk))

      [,1]
2000-01-03 "0.000160726642338866"
2000-01-04 "0.000264396469319473"
2000-01-05 "0.000304650302935347"
2000-01-06 "0.000148582063339039"
2000-01-07 "0.000123266970191763"
2000-01-10 "0.000130693391920629"
      [,1]
2015-02-25 "9.06199441502613E-06"
2015-02-26 "1.70042838302636E-05"
2015-02-27 "1.31586225256439E-05"
2015-03-02 "8.36714993575352E-06"
2015-03-03 "2.24875045689952E-05"
2015-03-04 "3.12041413708352E-05"
```

We can get SPX data from Yahoo using the quantmod package:

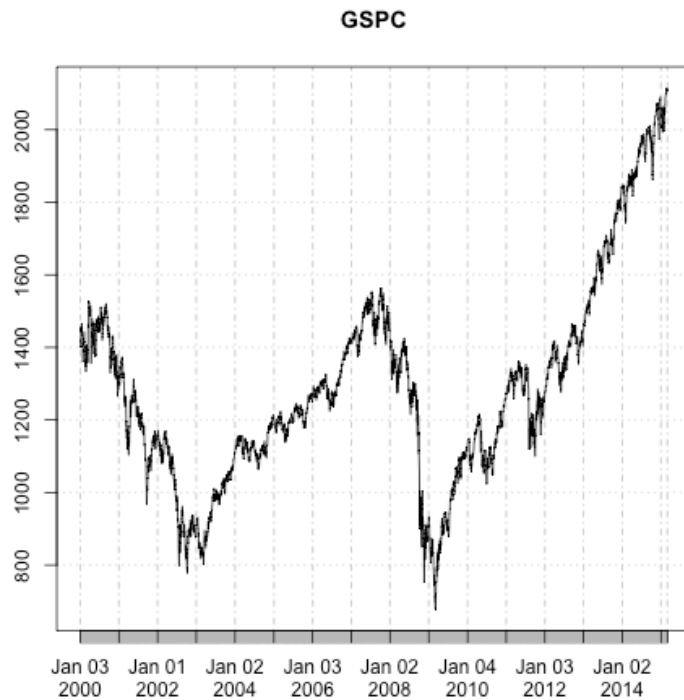
In [10]: %%R

```
library(quantmod)

options("getSymbols.warning4.0"=FALSE)
getSymbols("^GSPC",from="2000-01-01") #Creates the time series object GSPC

plot(GSPC)
```

Loading required package: TTR  
Version 0.4-0 included new data defaults. See ?getSymbols.



### The smoothness of the volatility process

For  $q \geq 0$ , we define the  $q$ th sample moment of differences of log-volatility at a given lag  $\Delta$ :

$$m(q, \Delta) = \langle |\log \sigma_{t+\Delta} - \log \sigma_t|^q \rangle$$

For example

$$m(2, \Delta) = \langle (\log \sigma_{t+\Delta} - \log \sigma_t)^2 \rangle$$

is just the sample variance of differences in log-volatility at the lag  $\Delta$ .

### Scaling of $m(q, \Delta)$ with lag $\Delta$

```
In [11]: %%R

sig <- sqrt(as.numeric(spx.rk))

mq.del.Raw <- function(q,lag){mean(abs(diff(log(sig),lag=lag))^q)}
mq.del <- function(x,q){sapply(x,function(x){mq.del.Raw(q,x)}})

# Plot mq.del(1:100,q) for various q

x <- 1:100

mycol <- rainbow(5)

ylab <- expression(paste(log, " ",m(q,Delta)))
xlab <- expression(paste(log, " ", Delta))

qVec <- c(.5,1,1.5,2,3)
zeta.q <- numeric(5)

q <- qVec[1]
plot(log(x),log(mq.del(x,q)),pch=20,cex=.5,
      ylab=ylab, xlab=xlab,ylim=c(-3,-.5))
(fit.lm <- lm(log(mq.del(x,q)) ~ log(x)))
abline(fit.lm, col=mycol[1],lwd=2)
zeta.q[1] <- coef(fit.lm)[2]

for (i in 2:5){
  q <- qVec[i]
  points(log(x),log(mq.del(x,q)),pch=20,cex=.5)
  fit.lm <- lm(log(mq.del(x,q)) ~ log(x))
  abline(fit.lm, col=mycol[i],lwd=2)
  zeta.q[i] <- coef(fit.lm)[2]
}
legend("bottomright", c("q = 0.5","q = 1.0","q = 1.5","q = 2.0","q = 3.0"),inset=0.05, lty=1, col =
mycol)

print(zeta.q)
```

```
[1] 0.07261453 0.14193859 0.20803046 0.27113292 0.38949549
```




Figure 2:  $\log m(q, \Delta)$  as a function of  $\log \Delta$ , SPX.

### Monofractal scaling result

- From the above log-log plot, we see that for each  $q$ ,  $m(q, \Delta) \propto \Delta^{\zeta_q}$ .
- How does  $\zeta_q$  scale with  $q$ ?

### Scaling of $\zeta_q$ with $q$

```
In [12]: %%R
plot(qVec, zeta.q, xlab="q", ylab=expression(zeta[q]))
fit.lm <- lm(zeta.q ~ qVec+0)
abline(fit.lm, col="red", lwd=2)
(h.est <- coef(fit.lm)[1])
```

```
qVec
0.1333966
```




Figure 3: Scaling of  $\zeta_q$  with  $q$ .

We find the monofractal scaling relationship

$$\zeta_q = qH$$

with  $H \approx 0.13$ .

- Note however that  $H$  does vary over time, in a narrow range.
- Note also that our estimate of  $H$  is biased high because we proxied instantaneous variance  $v_t$  with its average over each day  $\frac{1}{T} \int_0^T v_t dt$ , where  $T$  is one trading day.

## Estimated $H$ for all indices

We now repeat this analysis for all 21 indices in the Oxford-Man dataset.

```
In [13]: %%R

n <- length(rv.list)
h <- numeric(n) # H is estimated as half of the slope
nu <- numeric(n)

for (i in 1:n){ # Run all the regressions
  v <- rv.list[[i]]
  sig <- sqrt(as.numeric(v))

  x <- 1:100
  dlsig2 <- function(lag){mean((diff(log(sig),lag=lag))^2)}
  dlsig2Vec <- function(x){sapply(x,dlsig2)}

  fit.lm <- lm(log(dlsig2Vec(x)) ~ log(x))

  nu[i] <- sqrt(exp(coef(fit.lm)[1]))
  h[i] <- coef(fit.lm)[2]/2
}

(OxfordH <- data.frame(names(rv.list),h.est=h,nu.est=nu))

names.rv.list.    h.est    nu.est
1      SPX2.rk 0.1355665 0.3159217
2      FTSE2.rk 0.1453551 0.2692957
3      N2252.rk 0.1135003 0.3097937
4      GDAXI2.rk 0.1562330 0.2698523
5      RUT2.rk 0.1220715 0.3310271
6      AORD2.rk 0.0855579 0.3581023
7      DJI2.rk 0.1341714 0.3089454
8      IXIC2.rk 0.1277426 0.2908918
9      FCHI2.rk 0.1351559 0.2857176
10     HSI2.rk 0.1029477 0.2758013
11     KS11.rk 0.1277870 0.2731603
12     AEX.rk 0.1516508 0.2832848
13     SSMI.rk 0.1827945 0.2197018
14     IBEX2.rk 0.1312482 0.2781439
15     NSEI.rk 0.1137149 0.3228384
16     MXX.rk 0.0935143 0.3229200
17     BVSP.rk 0.1091998 0.3145880
18     GSPTSE.rk 0.1189021 0.3058733
19     STOXX50E.rk 0.1232648 0.3321809
20     FTSTI.rk 0.1271110 0.2300200
21     FTSEMIB.rk 0.1334866 0.2976255
```

## Distributions of $(\log \sigma_{t+\Delta} - \log \sigma_t)$ for various lags $\Delta$

Having established these beautiful scaling results for the moments, how do the histograms look?



In [14]: %%R

```

plotScaling <- function(j,scaleFactor){
  v <- as.numeric(rv.list[[j]])
  x <- 1:100

  xDel <- function(x,lag){diff(x,lag=lag)}
  sd1 <- sd(xDel(log(v),1))
  sd1 <- function(lag){sd(xDel(log(v),lag))}

  h <- OxfordH$h.est[j]

  plotLag <- function(lag){
    y <- xDel(log(v),lag)
    hist(y,breaks=100,freq=F,main=paste("Lag =",lag,"Days"),xlab=NA)# Very long tailed!
    curve(dnorm(x,mean=mean(y),sd=sd(y)),add=T,col="red",lwd=2)
    curve(dnorm(x,mean=0,sd=sd1*lag^h),add=T,lty=2,lwd=2,col="blue")
    lines(density(y)$x,density(y)$y,col="dark green",lty=3,lwd=3)
  }

  (lags <- scaleFactor^(0:3))
  print(names(rv.list)[j])
  par(mfrow=c(2,2))
  par(mar=c(3,2,1,3))
  for (i in 1:4){plotLag(lags[i])}
  par(mfrow=c(1,1))
}

plotScaling(1,5)

```

[1] "SPX2.rk"




Figure 4: Histograms of  $(\log \sigma_{t+\Delta} - \log \sigma_t)$  for various lags  $\Delta$ ; normal fit in red;  $\Delta = 1$  normal fit scaled by  $\Delta^{0.14}$  in blue.

## Universality?

- [Gatheral, Jaisson and Rosenbaum]<sup>[9]</sup> compute daily realized variance estimates over one hour windows for DAX and Bund futures contracts, finding similar scaling relationships.
- We have also checked that Gold and Crude Oil futures scale similarly.
  - Although the increments  $(\log \sigma_{t+\Delta} - \log \sigma_t)$  seem to be fatter tailed than Gaussian.

## A natural model of realized volatility

- Distributions of differences in the log of realized volatility are close to Gaussian.
  - This motivates us to model  $\sigma_t$  as a lognormal random variable.
- Moreover, the scaling property of variance of RV differences suggests the model:

(1)

$$\log \sigma_{t+\Delta} - \log \sigma_t = \nu (W_{t+\Delta}^H - W_t^H)$$

where  $W^H$  is fractional Brownian motion.

- In [Gatheral, Jaisson and Rosenbaum]<sup>[9]</sup>, we refer to a stationary version of (1) as the RFSV (for Rough Fractional Stochastic Volatility) model.

## Fractional Brownian motion (fBm)

- *Fractional Brownian motion* (fBm)  $\{W_t^H; t \in \mathbb{R}\}$  is the unique Gaussian process with mean zero and autocovariance function

$$\mathbb{E} [W_t^H W_s^H] = \frac{1}{2} \{ |t|^{2H} + |s|^{2H} - |t - s|^{2H} \}$$

where  $H \in (0, 1)$  is called the *Hurst index* or parameter.

- In particular, when  $H = 1/2$ , fBm is just Brownian motion.
- If  $H > 1/2$ , increments are positively correlated. % so the process is trending.
- If  $H < 1/2$ , increments are negatively correlated. % so the process is reverting.

## Representations of fBm

There are infinitely many possible representations of fBm in terms of Brownian motion. For example, with  $\gamma = \frac{1}{2} - H$ ,

### Mandelbrot-Van Ness

$$W_t^H = C_H \left\{ \int_{-\infty}^t \frac{dW_s}{(t-s)^\gamma} - \int_{-\infty}^0 \frac{dW_s}{(-s)^\gamma} \right\}.$$

where the choice

$$C_H = \sqrt{\frac{2H\Gamma(3/2-H)}{\Gamma(H+1/2)\Gamma(2-2H)}}$$

ensures that

$$\mathbb{E} [W_t^H W_s^H] = \frac{1}{2} \{ t^{2H} + s^{2H} - |t - s|^{2H} \}.$$

## Comte and Renault: FSV model

[Comte and Renault]<sup>[6]</sup> were perhaps the first to model volatility using fractional Brownian motion.

In their fractional stochastic volatility (FSV) model,

$$\frac{dS_t}{S_t} = \sigma_t dZ_t$$

$$d \log \sigma_t = -\alpha (\log \sigma_t - \theta) dt + \gamma d\hat{W}_t^H$$

with

$$\hat{W}_t^H = \int_0^t \frac{(t-s)^{H-1/2}}{\Gamma(H+1/2)} dW_s, \quad 1/2 \leq H < 1$$

and  $\mathbb{E}[dW_t dZ_t] = \rho dt$ .

- The FSV model is a generalization of the Hull-White stochastic volatility model.

## RFSV and FSV

- The model (1):

$$\log \sigma_{t+\Delta} - \log \sigma_t = \nu (W_{t+\Delta}^H - W_t^H)$$

is not stationary.

- Stationarity is desirable both for mathematical tractability and also to ensure reasonableness of the model at very large times.
- The RFSV model (the stationary version of (1)) is formally identical to the FSV model. Except that
- $H < 1/2$  in RFSV vs  $H > 1/2$  in FSV.
- $\alpha T \gg 1$  in RFSV vs  $\alpha T \sim 1$  in FSV, where  $T$  is a typical timescale of interest.

## FSV and long memory

- Why did [Comte and Renault]<sup>[6]</sup> choose  $H > 1/2$ ?
- Because it has been a widely-accepted stylized fact that the volatility time series exhibits long memory.
- In this technical sense, *long memory* means that the autocorrelation function of volatility decays as a power-law.
- One of the influential papers that established this was [Andersen, Bollerslev, Diebold and Ebens]<sup>[11]</sup> which estimated the degree  $d$  of fractional integration from daily realized variance data for the 30 DJIA stocks.
- Using the GPH estimator, they found  $d$  around 0.35 which implies that the ACF  $\rho(\tau) \sim \tau^{2d-1} = \tau^{-0.3}$  as  $\tau \rightarrow \infty$ .
- But every statistical estimator assumes the validity of some underlying model!
- In the RFSV model,

$$\rho(\Delta) \sim \exp \left\{ -\frac{\eta^2}{2} \Delta^{2H} \right\}.$$

- Using the same GPH estimator on the Oxford-Man RV data we find  $d = 0.48$ . But our model (1) is different from that of [Andersen, Bollerslev, Diebold and Ebens]<sup>[11]</sup>. In our case,  $\rho(\tau) \sim \tau^{2H-1} = \tau^{-0.72}$  as  $\tau \rightarrow \infty$ .

## Correlogram and test of scaling

In [15]: %%R

```
v <- rv.list[[1]] # Pick spx.rk
sig <- sqrt(as.numeric(v))

aclog <- acf(log(sig), lag=100, plot=F)
plot(aclog$lag[-1], aclog$acf[-1], pch=20, ylab=expression(rho(Delta)), xlab=expression(paste(Delta, " (
days)")), log="xy", col="blue")
esig2 <- mean(sig)^2
covdel <- acf(sig, lag.max=100, type="covariance", plot=F)$acf[-1]
x <- (1:100)^.28
plot(x, log(covdel+esig2), pch=20, col="dark green", ylab=expression(phi(Delta)), xlab=expression(Delta^
0.28 ))
abline(lm(log(covdel+esig2)~x), col="red", lwd=2)
```





Figure 5: The upper plot is a conventional correlogram of  $\sigma_t = \sqrt{RV_t}$ ; the lower plot is of  $\phi(\Delta) := \langle \log(\text{cov}(\sigma_{t+\Delta}, \sigma_t) + \langle \sigma_t \rangle^2) \rangle$  vs  $\Delta^{2H}$  with  $H = 0.14$ . The lower plot again supports the scaling relationship  $m(2, \Delta) \propto \Delta^{2H}$ .

## Heuristic derivation of autocorrelation function

We assume that  $\sigma_t = \bar{\sigma}_t e^{\eta W_t^H}$ . Then

$$\begin{aligned} & \text{cov} [\sigma_t, \sigma_{t+\Delta}] \\ &= \bar{\sigma}_t \bar{\sigma}_{t+\Delta} \left[ \exp \left\{ \frac{\eta^2}{2} (t^{2H} + (t+\Delta)^{2H} - \Delta^{2H}) \right\} - 1 \right] \\ &\sim \bar{\sigma}_t \bar{\sigma}_{t+\Delta} \exp \left\{ \frac{\eta^2}{2} (t^{2H} + (t+\Delta)^{2H} - \Delta^{2H}) \right\} \text{ as } t \rightarrow \infty. \end{aligned}$$

Similarly,

$$\text{var} [\sigma_t] \sim \bar{\sigma}_t^2 \exp \{ \eta^2 t^{2H} \}.$$

Thus

$$\rho(\Delta) = \frac{\text{cov} [\sigma_t, \sigma_{t+\Delta}]}{\sqrt{\text{var} [\sigma_t] \text{var} [\sigma_{t+\Delta}]}} \sim \exp \left\{ -\frac{\eta^2}{2} \Delta^{2H} \right\}.$$

## Model vs empirical autocorrelation functions

```

In [16]: %%R

v <- rv.list[[1]] # Pick spx.rk
sig <- sqrt(as.numeric(v))

aclog <- acf(log(sig), lag=100, plot=F)
plot(aclog$lag[-1], aclog$acf[-1], pch=20, ylab=expression(rho(Delta)),
      xlab=expression(paste(Delta, " (days)")), log="xy", col="blue")
y <- log(aclog$acf)
(h.spx <- OxfordH$h.est[1])
x <- aclog$lag^(2*h.spx)
fit.lm <- lm(y[-1]~x[-1])
a <- fit.lm$coef[1]
b <- fit.lm$coef[2]
curve(exp(a+b*x^(2*h.est)), from=0.001, to=100, col="red", add=T, lwd=2, log="xy")

```




Figure 6: Here we superimpose the RFSV functional form  $\rho(\Delta) \sim \exp\left\{-\frac{\eta^2}{2} \Delta^{2H}\right\}$  (in red) on the empirical curve (in blue).

## Volatility is not long memory

- It's clear from Figures 5 and 6 that volatility is not long memory.
- Moreover, the RFSV model reproduces the observed autocorrelation function very closely.
- [Gatheral, Jaisson and Rosenbaum]<sup>[6]</sup> further simulate volatility in the RFSV model and apply standard estimators to the simulated data.
  - Real data and simulated data generate very similar plots and similar estimates of the long memory parameter to those found in the prior literature.
- The RFSV model does not have the long memory property.
- Classical estimation procedures seem to identify spurious long memory of volatility.

## Incompatibility of FSV with realized variance (RV) data}

- In Figure 7, we demonstrate graphically that long memory volatility models such as FSV with  $H > 1/2$  are not compatible with the RV data.
- In the FSV model, the autocorrelation function  $\rho(\Delta) \propto \Delta^{2H-2}$ . Then, for long memory, we must have  $1/2 < H < 1$ .
  - For  $\Delta \gg 1/\alpha$ , stationarity kicks in and  $m(2, \Delta)$  tends to a constant as  $\Delta \rightarrow \infty$ .
  - For  $\Delta \ll 1/\alpha$ , mean reversion is not significant and  $m(2, \Delta) \propto \Delta^{2H}$ .

## RFSV vs FSV

- We can compute  $m(2, \Delta)$  explicitly in both the FSV and RFSV models.
- The smallest possible value of  $H$  in FSV is  $H = 1/2$ . One empirical estimate in the literature says that  $H \approx 0.53$  some time in 2008.
- Let's see how the theoretical estimates of  $m(2, \Delta)$  compare with data.




Figure 7: Black points are empirical estimates of  $m(2, \Delta)$ ; the blue line is the FSV model with  $\alpha = 0.5$  and  $H = 0.53$ ; the orange line is the RFSV model with  $\alpha = 0$  and  $H = 0.14$ .

## Does simulated RSV data look real?






Figure 8: Volatility of SPX (above) and of the RFSV model (below).

### Remarks on the comparison

- The simulated and actual graphs look very alike.
- Persistent periods of high volatility alternate with low volatility periods.
- $H \sim 0.1$  generates very rough looking sample paths (compared with  $H = 1/2$  for Brownian motion).
- Hence *rough volatility*.
- On closer inspection, we observe fractal-type behavior.
- The graph of volatility over a small time period looks like the same graph over a much longer time period.
- This feature of volatility has been investigated both empirically and theoretically in, for example, [Bacry and Muzy]<sup>[3]</sup>.
- In particular, their Multifractal Random Walk (MRW) is related to a limiting case of the RSFV model as  $H \rightarrow 0$ .

## Pricing under rough volatility

The foregoing behavior suggest the following model for volatility under the real (or historical or physical) measure  $\mathbb{P}$ :

$$\log \sigma_t = \nu W_t^H.$$

Let  $\gamma = \frac{1}{2} - H$ . We choose the Mandelbrot-Van Ness representation of fractional Brownian motion  $W^H$  as follows:

$$W_t^H = C_H \left\{ \int_{-\infty}^t \frac{dW_s^{\mathbb{P}}}{(t-s)^\gamma} - \int_{-\infty}^0 \frac{dW_s^{\mathbb{P}}}{(-s)^\gamma} \right\}$$

where the choice

$$C_H = \sqrt{\frac{2H\Gamma(3/2-H)}{\Gamma(H+1/2)\Gamma(2-2H)}}$$

ensures that

$$\mathbb{E} [W_t^H W_s^H] = \frac{1}{2} \{t^{2H} + s^{2H} - |t-s|^{2H}\}.$$

Then

$$\begin{aligned} & \log v_u - \log v_t \\ &= \nu C_H \left\{ \int_t^u \frac{1}{(u-s)^\gamma} dW_s^{\mathbb{P}} + \int_{-\infty}^t \left[ \frac{1}{(u-s)^\gamma} - \frac{1}{(t-s)^\gamma} \right] dW_s^{\mathbb{P}} \right\} \\ &=: 2\nu C_H [M_t(u) + Z_t(u)]. \end{aligned}$$

- Note that  $\mathbb{E}^{\mathbb{P}} [M_t(u) | \mathcal{F}_t] = 0$  and  $Z_t(u)$  is  $\mathcal{F}_t$ -measurable.
- To price options, it would seem that we would need to know  $\mathcal{F}_t$ , the entire history of the Brownian motion  $W_s$  for  $s < t$ !

## Pricing under $\mathbb{P}$

Let

$$\tilde{W}_t^{\mathbb{P}}(u) := \sqrt{2H} \int_t^u \frac{dW_s^{\mathbb{P}}}{(u-s)^\gamma}$$

With  $\eta := 2\nu C_H / \sqrt{2H}$  we have  $2\nu C_H M_t(u) = \eta \tilde{W}_t^{\mathbb{P}}(u)$  so denoting the stochastic exponential by  $\mathcal{E}(\cdot)$ , we may write

$$\begin{aligned} v_u &= v_t \exp \left\{ \eta \tilde{W}_t^{\mathbb{P}}(u) + 2\nu C_H Z_t(u) \right\} \\ &= \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left( \eta \tilde{W}_t^{\mathbb{P}}(u) \right). \end{aligned}$$

- The conditional distribution of  $v_u$  depends on  $\mathcal{F}_t$  only through the variance forecasts  $\mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t]$ ,
- To price options, one does not need to know  $\mathcal{F}_t$ , the entire history of the Brownian motion  $W_s^{\mathbb{P}}$  for  $s < t$ .

## Pricing under $\mathbb{Q}$

Our model under  $\mathbb{P}$  reads:

$$(2) \quad v_u = \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left( \eta \tilde{W}_t^{\mathbb{P}}(u) \right).$$

Consider some general change of measure

$$dW_s^{\mathbb{P}} = dW_s^{\mathbb{Q}} + \lambda_s ds,$$

where  $\{\lambda_s : s > t\}$  has a natural interpretation as the price of volatility risk. We may then rewrite (2) as

$$v_u = \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left( \eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \exp \left\{ \eta \sqrt{2H} \int_t^u \frac{\lambda_s}{(u-s)^{\gamma}} ds \right\}.$$

- Although the conditional distribution of  $v_u$  under  $\mathbb{P}$  is lognormal, it will not be lognormal in general under  $\mathbb{Q}$ .
- The upward sloping smile in VIX options means  $\lambda_s$  cannot be deterministic in this picture.

## The rough Bergomi (rBergomi) model

Let's nevertheless consider the simplest change of measure

$$dW_s^{\mathbb{P}} = dW_s^{\mathbb{Q}} + \lambda(s) ds,$$

where  $\lambda(s)$  is a deterministic function of  $s$ . Then from (2), we would have

$$\begin{aligned} v_u &= \mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t] \mathcal{E} \left( \eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \exp \left\{ \eta \sqrt{2H} \int_t^u \frac{1}{(u-s)^{\gamma}} \lambda(s) ds \right\} \\ &= \xi_t(u) \mathcal{E} \left( \eta \tilde{W}_t^{\mathbb{Q}}(u) \right) \end{aligned}$$

where the forward variances  $\xi_t(u) = \mathbb{E}^{\mathbb{Q}} [v_u | \mathcal{F}_t]$  are (at least in principle) tradable and observed in the market.

- $\xi_t(u)$  is the product of two terms:
- which depends on the historical path  $\{W_s, s < t\}$  of the Brownian motion
- a term which depends on the price of risk  $\lambda(s)$ .

## Features of the rough Bergomi model

- The rBergomi model is a non-Markovian generalization of the Bergomi model:

$$\mathbb{E} [v_u | \mathcal{F}_t] \neq \mathbb{E} [v_u | v_t].$$

- The rBergomi model is Markovian in the (infinite-dimensional) state vector  $\mathbb{E}^{\mathbb{Q}} [v_u | \mathcal{F}_t] = \xi_t(u)$ .
- We have achieved our aim from Lecture 1 of replacing the exponential kernels in the Bergomi model with a power-law kernel.
- We may therefore expect that the rBergomi model will generate a realistic term structure of ATM volatility skew.

## Re-interpretation of the conventional Bergomi model

- A conventional  $n$ -factor Bergomi model is not self-consistent for an arbitrary choice of the initial forward variance curve  $\xi_t(u)$ .
- $\xi_t(u) = \mathbb{E}[v_u | \mathcal{F}_t]$  should be consistent with the assumed dynamics.
- Viewed from the perspective of the fractional Bergomi model however:
  - The initial curve  $\xi_t(u)$  reflects the history  $\{W_s; s < t\}$  of the driving Brownian motion up to time  $t$ .
  - The exponential kernels in the exponent of the conventional Bergomi model approximate more realistic power-law kernels.
  - The conventional two-factor Bergomi model is then justified in practice as a tractable Markovian engineering approximation to a more realistic fractional Bergomi model.

## The stock price process

- The observed anticorrelation between price moves and volatility moves may be modeled naturally by anticorrelating the Brownian motion  $W$  that drives the volatility process with the Brownian motion driving the price process.
- Thus

$$\frac{dS_t}{S_t} = \sqrt{v_t} dZ_t$$

with

$$dZ_t = \rho dW_t + \sqrt{1 - \rho^2} dW_t^\perp$$

where  $\rho$  is the correlation between volatility moves and price moves.

## Simulation of the rBergomi model

We simulate the rBergomi model as follows:

- Construct the joint covariance matrix for the Volterra process  $\tilde{W}$  and the Brownian motion  $Z$  and compute its Cholesky decomposition.
- For each time, generate iid normal random vectors {and multiply them by the lower-triangular matrix obtained by the Cholesky decomposition} to get a  $m \times 2n$  matrix of paths of  $\tilde{W}$  and  $Z$  with the correct joint marginals.
- With these paths held in memory, we may evaluate the expectation under  $\mathbb{Q}$  of any payoff of interest.
- This procedure is very slow!
  - Speeding up the simulation is work in progress.

## Guessing rBergomi model parameters

- The rBergomi model has only three parameters:  $H$ ,  $\eta$  and  $\rho$ .
- If we had a fast simulation, we could just iterate on these parameters to find the best fit to observed option prices. But we don't.
- However, the model parameters  $H$ ,  $\eta$  and  $\rho$  have very direct interpretations:
  - $H$  controls the decay of ATM skew  $\psi(\tau)$  for very short expirations.
  - The product  $\rho \eta$  sets the level of the ATM skew for longer expirations.
  - Keeping  $\rho \eta$  constant but decreasing  $\rho$  (so as to make it more negative) pushes the minimum of each smile towards higher strikes.
- So we can guess parameters in practice.
- As we will see, even without proper calibration (*i.e.* just guessing parameters), rBergomi model fits to the volatility surface are amazingly good.

### Parameter estimation from historical data

- Both the roughness parameter (or Hurst parameter)  $H$  and the volatility of volatility  $\eta$  should be the same under  $\mathbb{P}$  and  $\mathbb{Q}$ .
- Earlier, using the Oxford-Man realized variance dataset, we estimated the Hurst parameter  $H_{eff} \approx 0.14$  and volatility of volatility  $\nu_{eff} \approx 0.3$ .
- However, we not observe the instantaneous volatility  $\sigma_t$ , only  $\frac{1}{\delta} \int_0^\delta \sigma_t^2 dt$  where  $\delta$  is roughly 3/4 of a whole day from close to close.
- Using Appendix C of [Gatheral, Jaisson and Rosenbaum]<sup>[9]</sup>, we rescale finding  $H \approx 0.05$  and  $\nu \approx 1.7$ .
- Also, recall that

$$\eta = 2\nu \frac{C_H}{\sqrt{2H}} = 2\nu \sqrt{\frac{\Gamma(3/2 - H)}{\Gamma(H + 1/2)\Gamma(2 - 2H)}}$$

which yields the estimate 2.5.

### SPX smiles in the rBergomi model

- In Figures 9 and 10, we show how well a rBergomi model simulation with guessed parameters fits the SPX option market as of February 4, 2010, a day when the ATM volatility term structure happened to be pretty flat.
- rBergomi parameters were:  $H = 0.07, \eta = 1.9, \rho = -0.9$ .
- Only three parameters to get a very good fit to the whole SPX volatility surface!

### rBergomi fits to SPX smiles as of 04-Feb-2010

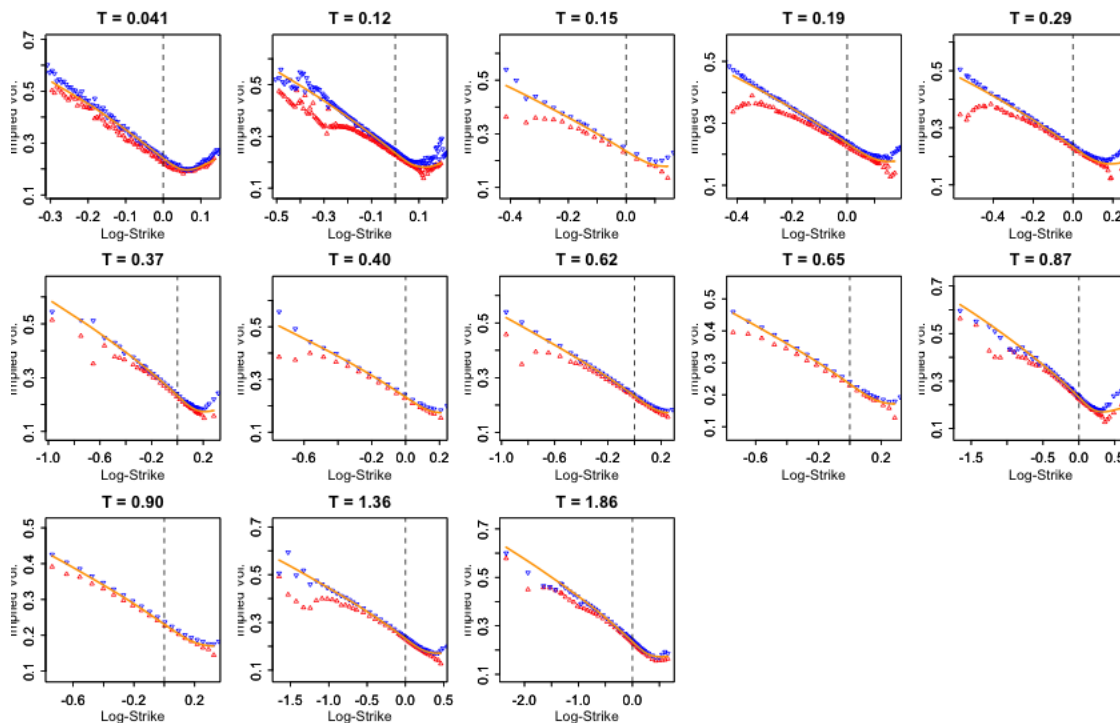


Figure 9: Red and blue points represent bid and offer SPX implied volatilities; orange smiles are from the rBergomi simulation.

### Shortest dated smile as of February 4, 2010




Figure 10: Red and blue points represent bid and offer SPX implied volatilities; orange smile is from the rBergomi simulation.

### ATM volatilities and skews

In Figures 11 and 12, we see just how well the rBergomi model can match empirical skews and vols. Recall also that the parameters we used are just guesses!

### Term structure of ATM skew as of February 4, 2010




Figure 11: Blue points are empirical skews; the red line is from the rBergomi simulation.

## Term structure of ATM vol as of February 4, 2010




Figure 12: Blue points are empirical ATM volatilities; the red line is from the rBergomi simulation.

## Another date

- Now we take a look at another date: August 14, 2013, two days before the last expiration date in our dataset.
- Options set at the open of August 16, 2013 so only one trading day left.
- Note in particular that the extreme short-dated smile is well reproduced by the rBergomi model.
- There is no need to add jumps!

### SPX smiles as of August 14, 2013

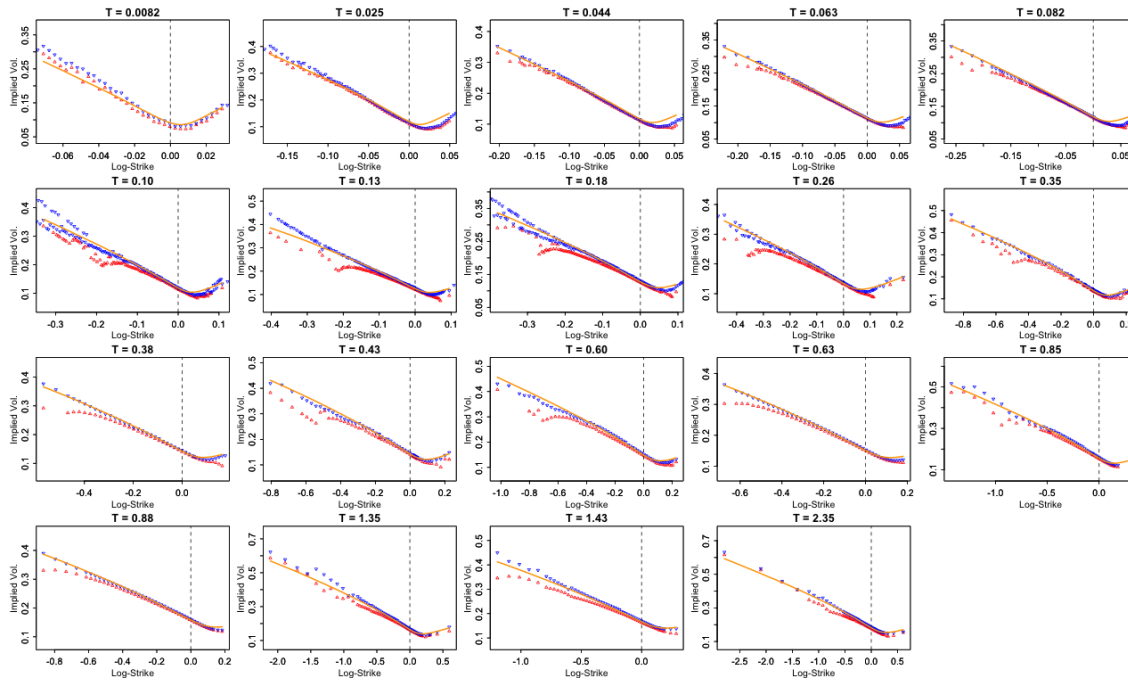


Figure 13: Red and blue points represent bid and offer SPX implied volatilities; orange smiles are from the rBergomi simulation.

### The forecast formula

- In the RFSV model (1),  $\log \sigma_t \approx \nu W_t^H + C$  for some constant  $C$ .
- [Nuzman and Poor]<sup>[11]</sup> show that  $W_{t+\Delta}^H$  is conditionally Gaussian with conditional expectation

$$\mathbb{E}[W_{t+\Delta}^H | \mathcal{F}_t] = \frac{\cos(H\pi)}{\pi} \Delta^{H+1/2} \int_{-\infty}^t \frac{W_s^H}{(t-s+\Delta)(t-s)^{H+1/2}} ds$$

and conditional variance

$$\text{Var}[W_{t+\Delta}^H | \mathcal{F}_t] = c \Delta^{2H}.$$

where

$$c = \frac{\Gamma(3/2 - H)}{\Gamma(H + 1/2)\Gamma(2 - 2H)}.$$



## The forecast formula

Thus, we obtain

### Variance forecast formula

$$(3) \quad \mathbb{E}^{\mathbb{P}} [v_{t+\Delta} | \mathcal{F}_t] = \exp \left\{ \mathbb{E}^{\mathbb{P}} [\log(v_{t+\Delta}) | \mathcal{F}_t] + 2 c \nu^2 \Delta^{2H} \right\}$$

where

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}} [\log v_{t+\Delta} | \mathcal{F}_t] \\ &= \frac{\cos(H\pi)}{\pi} \Delta^{H+1/2} \int_{-\infty}^t \frac{\log v_s}{(t-s+\Delta)(t-s)^{H+1/2}} ds. \end{aligned}$$

## Implement variance forecast in R

```
In [17]: %%R
# Find all of the dates
dateIndex <- substr(as.character(index(spx.rk)),1,10) # Create index of dates

cTilde <- function(h){gamma(3/2-h)/(gamma(h+1/2)*gamma(2-2*h))} # Factor because we are computing c
conditional on \cF_t

# XTS compatible version of forecast
rv.forecast.XTS <- function(rvdata,h,date,nLags,delta,nu){
  i <- (1:nLags)-1
  cf <- 1/((i+1/2)^(h+1/2)*(i+1/2+delta)) # Lowest number should apply to latest date
  datepos <- which(dateIndex==date)
  ldata <- log(as.numeric(rvdata[datepos-i])) # Note that this object is ordered from earlier to la
ter
  pick <- which(!is.na(ldata))
  norm <- sum(cf[pick])
  fcst <- cf[pick]**%ldata[rev(pick)]/norm # Most recent dates get the highest weight
  return(exp(fcst+2*nu^2*cTilde(h)*delta^(2*h)))
}
```

## SPX actual vs forecast variance

```
In [18]: %%R
rvdata <- spx.rk
nu <- OxfordH$nu.est[1] # Vol of vol estimate for SPX
n <- length(rvdata)
delta <- 1
nLags <- 500
range <- nLags:(n-delta)
rv.predict <- sapply(dateIndex[range],function(d){rv.forecast.XTS(rvdata,h=h.est,d,nLags=nLags,delt
a=delta,nu=nu)})
rv.actual <- rvdata[range+delta]
```

## Scatter plot of delta days ahead predictions

```
In [19]: %%R
x <- as.numeric(rv.predict)
y <- as.numeric(rv.actual)
plot(x,y,col="blue",pch=20, ylab="Actual vol.", xlab="Predicted vol.")
```




Figure 14: Actual vols vs predicted vols.

## Superimpose actual and predicted vols

```
In [20]: %%R
vol.actual <- sqrt(as.numeric(rv.actual)*252)
vol.predict <- sqrt(rv.predict*252)
plot(vol.actual, col="blue",type="l")
lines(vol.predict, col="red",type="l")
```




Figure 15: Actual volatilities in blue; predicted vols in red.

### Forecasting the variance swap curve

Finally, we forecast the whole variance swap curve using the variance forecasting formula (3).

```

In [21]: %%R

library(stinepack)

xil <- function(date,nu,h,dt, tscale){ # dt=(u-t) is in units of years
  xi <- rv.forecast.XTS(spx.rk,h=h,date=date,nLags=500,delta=dt*tscale,nu)
  return(xi)
}

# Forward variance curve (again the array tt should be in units of years)
xi <- function(date,tt,nu,h, tscale){sapply(tt,function(x){xil(date,nu=nu,h=h,x,tscale)}})

nu <- OxfordH$nu.est[1]
h <- OxfordH$h.est[1]

varSwapCurve <- function(date,bigT,nSteps,nu,h,tscale,onFactor){
  # Make vector of fwd variances
  tt <- seq(0,bigT,length.out=(nSteps+1))
  dt <- tt[2]
  xicurve <- xi(date,tt,nu,h,tscale)
  xicurve.mid <- (xicurve[1:nSteps]+xicurve[2:(nSteps+1)])/2
  int.xicurve <- cumsum(xicurve.mid)*dt
  varcurve <- int.xicurve/tt[-1]
  varcurve <- c(xicurve[1], varcurve)*onFactor*tscale #onFactor is to compensate for overnight move
  s
  res <- data.frame(tt,sqrt(varcurve))
  names(res) <- c("texp","vsQuote")
  return(res)
}

varSwapForecast <- function(date,tau,nu,h,tscale,onFactor){
  vsc <- varSwapCurve(date,bigT=2.5,nSteps=100,nu=nu,h=h,tscale,onFactor) # Creates the whole curve
  x <- vsc$texp
  y <- vsc$vsQuote
  res <- stinterp(x,y,tau)$y
  return(res)
}

# Test the function
tau <- c(.25,.5,1,2)
date <- "2008-09-08"
varSwapForecast(date,tau,nu=nu,h=h,tscale=252,onFactor=1)

[1] 0.2195294 0.2159190 0.2156538 0.2186971

```

## Constructing a time series of variance swap curves

For each of 2,658 days from Jan 27, 2003 to August 31, 2013:

- We compute proxy variance swaps from closing prices of SPX options sourced from OptionMetrics (www.optionmetrics.com) via WRDS.
- We form the forecasts  $\mathbb{E}^{\mathbb{P}} [v_u | \mathcal{F}_t]$  using (3) with 500 lags of SPX RV data sourced from The Oxford-Man Institute of Quantitative Finance (<http://realized.oxford-man.ox.ac.uk>).
- We note that the actual variance swap curve is a factor (of roughly 1.4) higher than the forecast, which we may attribute to a combination of overnight movements of the index and the price of volatility risk.
- Forecasts must therefore be rescaled to obtain close-to-close realized variance forecasts.

### 3-month forecast vs actual variance swaps




Figure 16: Actual (proxy) 3-month variance swap quotes in blue vs forecast in red (with no scaling factor).

### Ratio of actual to forecast




Figure 17: The ratio between 3-month actual variance swap quotes and 3-month forecasts.

## The Lehman weekend

- Empirically, it seems that the variance curve is a simple scaling factor times the forecast, but that this scaling factor is time-varying.
  - We can think of this factor as having two multiplicative components: the overnight factor, and the price of volatility risk.
- Recall that as of the close on Friday September 12, 2008, it was widely believed that Lehman Brothers would be rescued over the weekend. By Monday morning, we knew that Lehman had failed.
- In Figure 18, we see that variance swap curves just before and just after the collapse of Lehman are just rescaled versions of the RFSV forecast curves.

## We need variance swap estimates for 12-Sep-2008 and 15-Sep-2008

We proxy these by taking SVI fits for the two dates and computing the log-strips.

```
In [22]: %%R

varSwaps12 <- c(
  0.2872021, 0.2754535, 0.2601864, 0.2544684, 0.2513854, 0.2515314,
  0.2508418, 0.2520099, 0.2502763, 0.2503309, 0.2580933, 0.2588361,
  0.2565093)

texp12 <- c(
  0.01916496, 0.04654346, 0.09582478, 0.19164956, 0.26830938, 0.29842574,
  0.51745380, 0.54483231, 0.76659822, 0.79397673, 1.26488706, 1.76317591,
  2.26146475)

varSwaps15 <- c(
  0.4410505, 0.3485560, 0.3083603, 0.2944378, 0.2756881, 0.2747838,
  0.2682212, 0.2679770, 0.2668113, 0.2706713, 0.2729533, 0.2689598,
  0.2733176)

texp15 <- c(
  0.01095140, 0.03832991, 0.08761123, 0.18343600, 0.26009582, 0.29021218,
  0.50924025, 0.53661875, 0.75838467, 0.78576318, 1.25667351, 1.75496235,
  2.25325120)
```

## Actual vs predicted over the Lehman weekend

```
In [23]: %%R

nu <- OxfordH$nu.est[1]
h <- OxfordH$h.est[1]

# Variance curve fV model forecasts
tau1000 <- seq(0,2.5,length.out=1001)[-1]
vs1 <- varSwapForecast("2008-09-12",tau1000,nu=nu,h=h,tscale=252,onFactor=1.29)
vs2 <- varSwapForecast("2008-09-15",tau1000,nu=nu,h=h,tscale=252,onFactor=1.29)
```

```
In [21]: %%R
# Plot results
plot(texp12,varSwaps12,type="b",col="red",ylim=c(0.2,0.45),xlab="Maturity",ylab="Variance swap quote",lwd=2)
lines(texp15,varSwaps15,type="b",col="blue",lwd=2)
lines(tau1000,vs1,col="red",type="l",lty=2,lwd=2)
lines(tau1000,vs2,col="blue",type="l",lty=2,lwd=2)
```




Figure 18: SPX variance swap curves as of September 12, 2008 (red) and September 15, 2008 (blue). The dashed curves are RFSV model forecasts rescaled by the 3-month ratio (1.29) as of the Friday close.

## Remarks

We note that

- The actual variance swaps curves are very close to the forecast curves, up to a scaling factor.
- We are able to explain the change in the variance swap curve with only one extra observation: daily variance over the trading day on Monday 15-Sep-2008.
- The SPX options market appears to be backward-looking in a very sophisticated way.

## The Flash Crash

- The so-called Flash Crash of Thursday May 6, 2010 caused intraday realized variance to be much higher than normal.
- In Figure 19, we plot the actual variance swap curves as of the Wednesday and Friday market closes together with forecast curves rescaled by the 3-month ratio as of the close on Wednesday May 5 (which was 2.52).
- We see that the actual variance curve as of the close on Friday is consistent with a forecast from the time series of realized variance that {} the anomalous price action of Thursday May 6.
- In Figure 20, we see that the actual variance swap curve on Monday, May 10 is consistent with a forecast that excludes the Flash Crash.
- Volatility traders realized that the Flash Crash should not influence future realized variance projections.

## Variance swap estimates

We again proxy variance swaps for 05-May-2010, 07-May-2010 and 10-May-2010 by taking SVI fits for the three dates and computing the log-strips.

In [24]: %%R

```
varSwaps5 <- c(
  0.4250369, 0.2552473, 0.2492892, 0.2564899, 0.2612677, 0.2659618, 0.2705928, 0.2761203,
  0.2828139, 0.2841165, 0.2884955, 0.2895839, 0.2927817, 0.2992602, 0.3116500)

texp5 <- c(
  0.002737851, 0.043805613, 0.120465435, 0.150581793, 0.197125257, 0.292950034,
  0.369609856, 0.402464066, 0.618754278, 0.654346338, 0.867898700, 0.900752909,
  1.117043121, 1.615331964, 2.631074606)

varSwaps7 <- c(
  0.5469727, 0.4641713, 0.3963352, 0.3888213, 0.3762354, 0.3666858, 0.3615814, 0.3627013,
  0.3563324, 0.3573946, 0.3495730, 0.3533829, 0.3521515, 0.3506186, 0.3594066)

texp7 <- c(
  0.01642710, 0.03832991, 0.11498973, 0.14510609, 0.19164956, 0.28747433, 0.36413415,
  0.39698836, 0.61327858, 0.64887064, 0.86242300, 0.89527721, 1.11156742, 1.60985626,
  2.62559890)

varSwaps10 <- c(
  0.3718439, 0.3023223, 0.2844810, 0.2869835, 0.2886912, 0.2905637, 0.2957070, 0.2960737,
  0.3005086, 0.3031188, 0.3058492, 0.3065815, 0.3072041, 0.3122905, 0.3299425)

texp10 <- c(
  0.008213552, 0.030116359, 0.106776181, 0.136892539, 0.183436003, 0.279260780,
  0.355920602, 0.388774812, 0.605065024, 0.640657084, 0.854209446, 0.887063655,
  1.103353867, 1.601642710, 2.617385352)
```



In [23]: %%R

```

# Variance curve fV model forecasts
vsf5 <- varSwapCurve("2010-05-05",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)
vsf7 <- varSwapCurve("2010-05-07",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)

plot(texp5,varSwaps5,type="b",col="red",xlab=expression(paste("Time to maturity ",tau)),ylab="Variance swap quote",lwd=2,ylim=c(0.2,.55))
lines(texp7,varSwaps7,type="b",col="green4",lwd=2)
legend("topright",inset=.02,c("May 5","May 7"),lty=1,col=c("red","green4"))

lines(vsf5,col="red",type="l",lty=2,lwd=2)
lines(vsf7,col="green4",type="l",lty=2,lwd=2)

```




Figure 19: SPX variance swap curves as of May 5, 2010 (red) and May 7, 2010 (green). The dashed curves are RFSV model forecasts rescaled by the 3-month ratio (2.52) as of the close on Wednesday May 5. The curve as of the close on May 7 is consistent with the forecast **including** the crazy moves on May 6.

## The weekend after the Flash Crash

Now we plot forecast and actual variance swap curves as of the close on Friday May 7 and Monday May 10.

In [25]: %%R

```

# Variance curve fV model forecasts
vsf7 <- varSwapCurve("2010-05-07",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)
vsf10 <- varSwapCurve("2010-05-10",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)

plot(texp7,varSwaps7,type="b",col="green4",xlab=expression(paste("Time to maturity ",tau)),ylab="Va
riance swap quote",lwd=2,ylim=c(0.2,.55),plot=NA)
lines(texp10,varSwaps10,type="b",col="orange",lwd=2)
legend("topright",inset=.02,c("May 7","May 10"),lty=1,col=c("green4","orange"))

lines(vsf7,col="green4",type="l",lty=2,lwd=2)
lines(vsf10,col="orange",type="l",lty=2,lwd=2)

```




Figure 20: The May 10 actual curve is inconsistent with a forecast that includes the Flash Crash.

Now let's see what happens if we exclude the Flash Crash from the time series used to generate the variance curve forecast.

```
In [26]: %%R
(fl原因.day <- which(index(spx.rk)=="2010-05-06"))
spx.rk.p <- spx.rk[-fl原因.day]
plot(spx.rk["2010-05-04::2010-05-10"],type="b")
lines(spx.rk.p["2010-05-04::2010-05-10"],col="red",type="b")
```

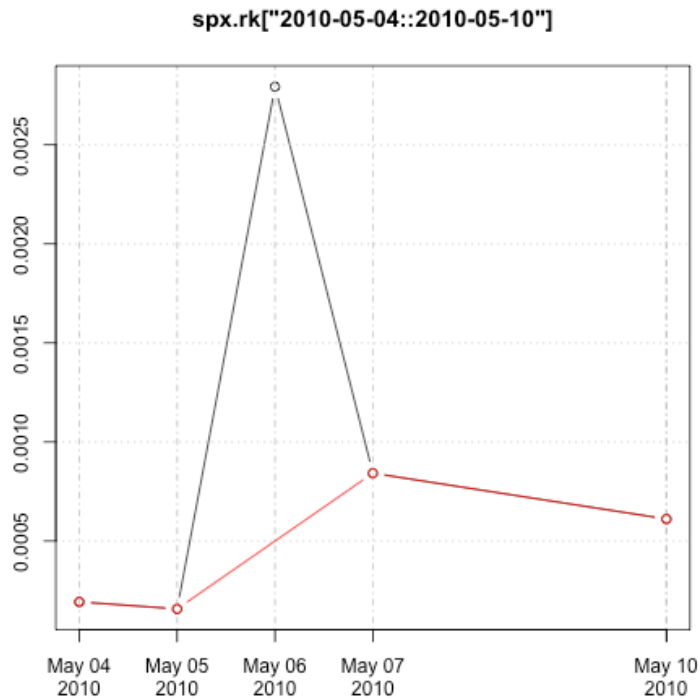


Figure 21: `spx.rk.p` has the May 6 realized variance datapoint eliminated. Notice the crazy realized variance estimate for May 6!

We need a new variance curve forecast function that uses the new time series.

```

In [27]: %%R

xilp <- function(date,nu,h,dt, tscale){ # dt=(u-t) is in units of years
  xi <- rv.forecast.XTS(spx.rk.p,h=h,date=date,nLags=500,delta=dt*tscale,nu)
  return(xi)
}

# Forward variance curve (again the array tt should be in units of years)
xip <- function(date,tt,nu,h, tscale){sapply(tt,function(x){xilp(date,nu=nu,h=h,x,tscale)}})

varSwapCurve.p <- function(date,bigT,nSteps,nu,h,tscale,onFactor){
  # Make vector of fwd variances
  tt <- seq(0,bigT,length.out=(nSteps+1))
  dt <- tt[2]
  xicurve <- xip(date,tt,nu,h,tscale)
  xicurve.mid <- (xicurve[1:nSteps]+xicurve[2:(nSteps+1)])/2
  int.xicurve <- cumsum(xicurve.mid)*dt
  varcurve <- int.xicurve/tt[-1]
  varcurve <- c(xicurve[1], varcurve)*onFactor*tscale #onFactor is to compensate for overnight move
  res <- data.frame(tt,sqrt(varcurve))
  names(res) <- c("texp","vsQuote")
  return(res)
}

varSwapForecast.p <- function(date,tau,nu,h,tscale,onFactor){
  vsc <- varSwapCurve.p(date,bigT=2.5,nSteps=100,nu=nu,h=h,tscale,onFactor) # Creates the whole curve
  x <- vsc$texp
  y <- vsc$vsQuote
  res <- stinterp(x,y,tau)$y
  return(res)
}

# Test the function
tau <- c(.25,.5,1,2)
date <- "2010-05-10"
varSwapForecast.p(date,tau,nu=nu,h=h,tscale=252,onFactor=1/(1-.35))

[1] 0.2411204 0.2418901 0.2488975 0.2623605

```

Finally, we compare our new forecast curves with the actuals.

In [28]: %%R

```

# Variance curve fV model forecasts
vsf7 <- varSwapCurve("2010-05-07",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)
vsf10p <- varSwapCurve.p("2010-05-10",bigT=2.5,nSteps=100,nu=nu,h=h,tscale=252,onFactor=2.52)

plot(texp7,varSwaps7,type="b",col="green4",xlab=expression(paste("Time to maturity ",tau)),ylab="Va
riance swap quote",lwd=2,ylim=c(0.2,.55),plot=NA)
lines(texp10,varSwaps10,type="b",col="orange",lwd=2)
legend("topright",inset=.02,c("May 7","May 10"),lty=1,col=c("green4","orange"))

lines(vsf7,col="green4",type="l",lty=2,lwd=2)
lines(vsf10p,col="orange",type="l",lty=2,lwd=2)

```




Figure 22: The May 10 actual curve is consistent with a forecast that excludes the Flash Crash.

## Summary

- We uncovered a remarkable monofractal scaling relationship in historical volatility.
  - A corollary is that volatility is not a long memory process, as widely believed.
- This leads to a natural non-Markovian stochastic volatility model under  $\mathbb{P}$ .
- The simplest specification of  $\frac{dQ}{dP}$  gives a non-Markovian generalization of the Bergomi model.
- The history of the Brownian motion  $W_s, s < t$  required for pricing is encoded in the forward variance curve, which is observed in the market.
- This model fits the observed volatility surface surprisingly well with very few parameters.
- For perhaps the first time, we have a simple consistent model of historical and implied volatility.

## References

1. [^](#) Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Heiko Ebens, The distribution of realized stock return volatility, *Journal of Financial Economics* **61**(1) 43-76 (2001).
2. [^](#) Christian Bayer, Peter Friz and Jim Gatheral, Pricing under rough volatility, Available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2554754](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2554754), (2015).
3. [^](#) Emmanuel Bacry and Jean-François Muzy, Log-infinitely divisible multifractal processes, *Communications in Mathematical Physics* **236**(3) 449-475 (2003).
4. [^](#) Lorenzo Bergomi, Smile dynamics II, *Risk Magazine* 67-73 (October 2005).
5. [^](#) Lorenzo Bergomi and Julien Guyon, Stochastic volatility's orderly smiles. *Risk Magazine* 60-66, (May 2012).
6. [^](#) Fabienne Comte and Eric Renault, Long memory in continuous-time stochastic volatility models, *Mathematical Finance* **8** 29-323(1998).
7. [^](#) Jim Gatheral, *The Volatility Surface: A Practitioner's Guide*, John Wiley and Sons, Hoboken, NJ (2006).
8. [^](#) Jim Gatheral and Antoine Jacquier, Arbitrage-free SVI volatility surfaces, *Quantitative Finance*, **14**(1) 59-71 (2014).
9. [^](#) Jim Gatheral, Thibault Jaisson and Mathieu Rosenbaum, Volatility is rough, available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2509457](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2509457), (2014).
10. [^](#) Jim Gatheral and Roel Oomen, Zero-intelligence realized variance estimation, *Finance and Stochastics* **14**(2) 249-283 (2010).
11. [^](#) Carl J. Nuzman and H. Vincent Poor, Linear estimation of self-similar processes via Lamperti's transformation, *Journal of Applied Probability* **37**(2) 429-452 (2000).

In []: