Financial Time Series

Topic 7: ARCH Related Models

Hung Chen
Department of Mathematics
National Taiwan University
5/15/2002

OUTLINE

1. Stock Volatility
2. Empirical Properties of Returns
3. Martingales and Random Walks
4. Testing the Random Walk Hypothesis
5. Stochastic Volatility Model
6. ARCH Processes
7. Detour to Pricing Theory
8. ARCH and Asset Pricing
9. Estimation

Stock Volatility

- Volatility: the conditional variance of the underlying asset returns
- Black-Scholes option pricing formula states that the price of a European call option is

$$c_t = P_t\Phi(x) - Kr^{-\ell}\Phi(x - \sigma_t\sqrt{\ell}),$$

where

$$x = \frac{\ln(P_t/Kr^{-\ell})}{\sigma_t\sqrt{\ell}} + \frac{1}{2}\sigma_t\sqrt{\ell}.$$

  - K: striking price
  - ℓ: the time to expiration
  - P_t: current price of the underlying stock
  - r; risk-free interest rate
  - σ_t: the conditional standard deviation of the log return of the specified stock

- The conditional variance of a stock return plays an important role in the pricing formula.
- How do we model the evolution of stock volatility?

- Conditional heteroscedastic models
  - Shocks of asset returns are NOT serially correlated, but dependent.
  - See ACF of squared and absolute returns of some stocks.
- Univariate volatility models:
  - ARCH: Engle (1992, *Econometrica*)
  - G(eneralized)ARCH: Bollerslev (1986, J. of *Econometrics*)
  - E(xponential)GARCH: Nelson (1991, *Econometric Theory*)
    Modeling asymetry in volatility
  - Stochastic volatility model: Melino and Turnball (1990, J. of *Econometrics*)
- Stock volatility is not directly observable.
  - The daily volatility is not directly observable from the daily returns.
  - If intra-daily prices are available, one can discuss the daily volatility.

Empirical Properties of Returns

- Empirical research on returns distributions has been ongoing since the early 1960s.
  - Daily returns of the market indexes and individual stocks tend to have high excess kurtoses.
  - Monthly returns have higher standard deviations than daily returns.
  - The skewness is not a serious problem for both daily and monthly returns.
- Volatility process: Study the evolution of conditional variances of the return over time.
  - Figures 2 and 3: The variabilities of returns vary over time and appear in clusters.
  - Extremes of a return series: large positive or negative returns
  - Volatility clusters: high for certain time periods and low for other periods
  - Volatility evolves over time in a continuous manner.

– Volatility varies within some fixed range. Volatility is stationary.

– Volatility seems to react differently to a big positive return and a big negative return.
EGARCH: capture the asymmetry in volatility induced by big positive and negative asset returns.

Study on Volatility

- x_t: the log return of a stock at time index t
- x_t is serially uncorrelated or with minor lower lag serial correlations, but it is dependent.
- Volatility models attempt to capture such dependence in the return series.
- Let $\mathcal{F}_{t-1}$ be the information available at time $t-1$. Consider conditional mean and variance

$$\begin{aligned} \mu_t &= E(x_t|\mathcal{F}_{t-1}) = g(\mathcal{F}_{t-1}), \\ \sigma_t^2 &= Var(x_t|\mathcal{F}_{t-1}) = h(\mathcal{F}_{t-1}), \end{aligned}$$

where $g(\cdot)$ and $h(\cdot)$ are well-defined functions with $h(\cdot) > 0$.

- It is common to assume that $\mu_t = \mu$.
- For a linear series, $g(\cdot)$ is a linear function of $\mathcal{F}_{t-1}$ and $h(\cdot) = \sigma_a^2$.
  In statistical literature, they focus on $g(\cdot)$.
  Model x_t as a stationary $ARMA(p,q)$.

$$\begin{aligned} x_t &= \mu_t + a_t, \\ \mu_t &= \phi_0 + \sum_{i=1}^{p} \phi_i x_{t-i} - \sum_{i=1}^{q} \theta_i a_{t-i}. \end{aligned}$$

- σ_t^2 is $Var(a_t|\mathcal{F}_{t-1})$.

Martingales and Random Walks

- A martingale is a stochastic process $\{x_t\}$ with the following properties:
  - $E(|x_t|) < \infty$ for each t;
  - $E(x_t|\mathcal{F}_s) = x_s$, whenever $s \leq t$.
    $\mathcal{F}_s$: the σ-algebra comprising events determined by observations over the interval $[0, t]$
    $\mathcal{F}_s \subset \mathcal{F}_t$ when $s \leq t$
  - Right continuous: $\mathcal{F}_t = \cap_{s>t}\mathcal{F}_s$
- $\mathcal{F}_t = \sigma(x_s; s \leq t)$: the past history of $\{x_t\}_0^t$ itself

$$E(x_t - x_s|\mathcal{F}_s) = 0, \quad s \leq t. \qquad (1)$$

- (1) can be written equivalently as

$$x_t = x_{t-1} + a_t$$

  where a_t is the martingale increment or martingale difference.
  Is it a random walk model?
  - The martingale rules out any dependence of the conditional expectation of $x_t - x_{t-1}$ on the information available at t.

- The random walk rules out not only *any dependence of the conditional expectation of $x_t - x_{t-1}$ on the information available at t* also **dependence involving the higher conditional moments of $x_t - x_{t-1}$**.
- Financial series are known to go through protracted quiet periods and also protracted periods of turbulence.
  This type of behavior could be modelled by a process in which successive conditional variances of $x_t - x_{t-1}$ (but not successive levels) are positive correlated.
  Such a specification would be consistent with a martingale, but not with the more restrictive random walk.

- Submartingale: $E(x_t - x_s | \mathcal{F}_s) \geq 0$, $s \leq t$.
- Supermartingale: $E(x_t - x_s | \mathcal{F}_s) \leq 0$, $s \leq t$.

Non-Linearity

- For random walk,
  a_t is $WN(0, \sigma^2)$. (i.e., stationary, uncorrelated, from a fixed distribution)
  a_t is $SWN(0, \sigma^2)$ (i.e., independent too)

- For martingale differences, a_t can be non-stationary.
- Why do we consider the dependence between conditional variances?
- Financial time series often go through protracted quiet periods interspersed with bursts of turbulence.
- Use non-linear stochastic processes to model such volatility.
- Suppose x_t is generated by the process $\Delta x_t = \eta_t$ with

$$\eta_t = a_t + \beta a_{t-1} a_{t-2},$$

where a_t is $SWN(0, \sigma^2)$.

- Properties of η_t:

$$\begin{aligned} E(a_t) &= 0, \\ V(a_t) &= constant, \\ E(\eta_t \eta_{t-k}) &= E(a_t a_{t-k} + \beta a_{t-1} a_{t-2} a_{t-k} \\ &\quad + \beta a_{t-1} a_{t-k-1} a_{t-k-2} \\ &\quad + \beta^2 a_{t-1} a_{t-2} a_{t-k-1} a_{t-k-2}). \end{aligned}$$

- For all $k \neq 0$, each of the term in the ACF has zero expectation.
  η_t behaves like an independent process.

- The conditional expectation is

$$\hat{\eta}_{t+1} = E(\eta_{t+1} | \eta_t, \eta_{t-1}, \cdots) = \beta a_t a_{t-1}.$$

- x_t is not a martingale because

$$E(x_{t+1} - x_t | \eta_t, \eta_{t-1}, \cdots) = \hat{\eta}_{t+1} \neq 0.$$

Testing the Random Walk Hypothesis

- Autocorrelation tests:
  - Suppose $w_t = \triangle x_t$ is $SWN(0, \sigma^2)$.
  - The sample autocorrelations (standardized by $\sqrt{T}$) calculated from the realization $\{w_t\}_1^T$ will be $N(0, 1)$.
  - Reject the hypothesis if, for example, $\sqrt{T}|r_1| > 1.96$.
  - Portmanteau tests: $Q^*(K)$ and $Q(K)$.
  - Those tests rely on the assumption that the random walk innovation is strict white noise.
    Refer to page 126 for further discussion.
- Calendar effects:
  - Consider autocorrelations associated with specific timing patterns.
  - January effect: Stock returns in this month are exceptionally large.
  - Weekend effect: Monday mean returns are negative rather than positive as for all other weekdays.

- Holiday effect: a much larger mean return for the day before holidays
- Turn-of-the-month effect: the four-day return around the turn of a month is greater than the average total monthly return
- Intramonth effect: the return over the first half of a month is significantly larger than the return over the second half

Stochastic Volatility

- Allow the variance (or the conditional variance) of the process to change either at certain discrete points in time or continuously.
- A stationary process must have a constant variance, certain conditional variances can change.
- For a non-linear stationary process x_t, the variance $Var(x_t)$ is a constant for all t, but the conditional variance $Var(x_t|x_{t-1}, x_{t-2}, \ldots)$ can change from period to period.
- Non-stationary variance or variance dependent on past observations and additional variables
- The models are non-linear, have high **kurtosis**, and positive autocorrelation between *squared* returns.

Stochastic volatility (SV) models

- $\{x_t\}_1^t$: the product process

$$x_t = \mu + \sigma_t U_t \tag{2}$$

where
$E(U_t) = 0$ and $Var(U_t) = 1$ for all t,
$Var(x_t|\sigma_t) = \sigma_t^2$, and
σ_t is a positive random variable.

- $E(x_t) = \mu$,

$$E(x_t - \mu)^2 = E(\sigma_t^2 U_t^2) = E(\sigma_t^2),$$

and autocovariance

$$E(x_t - \mu)(x_{t-k} - \mu) = E(\sigma_t \sigma_{t-k} U_t U_{t-k})$$
$$= E(\sigma_t \sigma_{t-k} U_t) E(U_{t-k}) = 0.$$

- Typically $U_t = (x_t - \mu)/\sigma_t$ is assumed to be normal and independent of σ_t.

- (2) is motivated by the discrete time approximation to the stochastic differential equation

$$\frac{dP}{P} = d(\log(P)) = \mu dt + \sigma dW$$

where $x_t = \triangle \log(P_t)$ and $W(t)$ is standard Brownian motion.
This is the usual diffusion process used to price financial assets in theoretical models of finance.

- In the world of time series analysis, write the above sdf by setting $dt = 1$.
  We then have

$$\log(P_{t+1}) - \log(P_t) = \mu + \sigma(W_{t+1} - W_t).$$

- Although x_t is a white noise, the squared and absolute deviation, $S_t = (x_t - \mu)^2$ and $M_t = |x_t - \mu|$, can be autocorrelated.

$$\begin{aligned} Cov(S_t, S_{t-k}) &= E(\sigma_t^2 \sigma_{t-k}^2) E(U_t^2 U_{t-k}^2) - (E(\sigma_t^2))^2 \\ &= E(\sigma_t^2 \sigma_{t-k}^2) - (E(\sigma_t^2))^2. \end{aligned}$$

- Fact: Almost all sample paths W of Brownian motion are of unbounded variation.
  They are not differentiable.

How do we model σ_t?

- The distribution of σ_t is skewed to the right. Consider a log-normal distribution.
- Define

$$h_t = \log(\sigma_t^2) = \gamma_0 + \gamma_1 h_{t-1} + \eta_t \qquad (3)$$

where $\eta_t \sim NID(0, \sigma_\eta^2)$ and is independent of U_t.
h_t represents the random and uneven flow of new information into financial market.

- $x_t = \mu + U_t \exp(h_t/2)$, where U_t is always stationary.
x_t will be stationary if and only if h_t is.
Or, $|\gamma_1| < 1$.
- Moments of x_t or S_t:
For even r,

$$\begin{aligned} E(x_t - \mu)^r &= E(U_t^r) E\left(\exp(r h_t/2)\right) \\ &= \frac{r!}{2^{r/2} r/2!} \exp\left(\frac{r}{2}\mu_h + \frac{r}{2}\frac{\sigma_h^2}{2}\right) \end{aligned}$$

where $\mu_h = E(h_t) = \gamma_0/(1-\gamma_1)$ and $\sigma_h^2 = V(h_t) = \sigma_\eta^2/(1-\gamma_1^2)$.
All odd moments are zero.

- kurtosis:

$$\frac{E(x_t-\mu)^4}{[E(x_t-\mu)^2]^2} = 3\exp(\sigma_h^2) > 3$$

  The process has fatter tails than a normal distribution.

- autocorrelation: Refer to page 129.

- Taking logarithms of (2) yields

$$\begin{aligned}\log(S_t) &= h_t + \log(U_t^2) \\ &= \mu_h + \frac{\eta_t}{1-\gamma_1 B} + \log(U_t^2) \\ \log(S_t) &\sim ARMA(1,1)\end{aligned}$$

  with non-normal innovations.

- The main difficulty with using SV models is that they are rather difficult to estimate.

ARCH Processes

- In (3), σ_t was dependent upon the information set $\{\eta_t, \sigma_{t-1}, \sigma_{t-2}, \ldots\}$.
- Now, consider the case that σ_t are a function of past values of x_t,

$$\sigma_t^2 = h(x_{t-1}, x_{t-2}, \ldots).$$

- ARCH(1) process: Engle (1982)
- First-order autoregressive conditional heteroskedastic process:
  Write ϵ_t as $\sigma_t U_t$ where $\{U_t\}$ is a sequence of iid r.v. with mean 0 and variance 1.

$$\sigma_t^2 = h(x_{t-1}) = \alpha_0 + \alpha_1 \epsilon_{t-1}^2, \qquad (4)$$

  where $\alpha_0, \alpha_1 > 0$.
  - The (mean-corrected) asset return x_t is serially uncorrelated but dependent.
  - The dependence of x_t can be described by a simple quadratic function.
  - Large deviations of x_{t-1} from the mean μ then cause a large variance for the next day.

Large returns tend to be followed by another large return

- Distribution of U_t: standard normal, standardized Student-t, or generalized error distribution.
- When $U_t \sim NID(0,1)$ and independent of σ_t,

$$x_t = \mu + U_t \sigma_t$$

is white noise and conditionally normal, i.e.

$$x_t | x_{t-1}, x_{t-2}, \cdots \sim NID(\mu, \sigma_t^2)$$

so that

$$Var(x_t | x_{t-1}) = \alpha_0 + \alpha_1 (x_{t-1} - \mu)^2.$$

  - $E(x_t) = E[E(x_t | \mathcal{F}_{t-1})] = \mu$
  - Unconditional variance:

$$\begin{aligned} Var(x_t) &= E[E(U_t^2 \sigma_t^2 | \mathcal{F}_{t-1})] \\ &= \alpha_0 + \alpha_1 E(x_{t-1} - \mu)^2. \end{aligned}$$

  - Because x_t is a stationary process, we have $Var(x_t) = \alpha_0/(1-\alpha_1)$ if $\alpha_1 < 1$.
  - It possesses constant variance yet changing conditional variance.

– When $0 < \alpha_1^2 < 1/3$, the fourth moment is finite.

$$E(U_t)^2 = 3[Var(U_t)]^2 \times \frac{1-\alpha_1^2}{1-3\alpha_1^2}.$$

– The fourth moment of U_t is greater than that of a normal random variable when $\alpha_1 \neq 0$.
This implies that the U_t process is heavy-tailed and it is capable of producing clusters of outliers.

- Using (4), the series of $S_t = (x_t - \mu)^2$ satisfy

$$E(S_t|S_{t-1}) = \alpha_0 + \alpha_1 S_{t-1}.$$

a stationary AR(1) process

- ARCH(q) process:

$$\begin{aligned}\sigma_t &= h(x_{t-1}, \ldots, x_{t-q}) \\ &= \left(\alpha_0 + \sum_{i=1}^{q} \alpha_i (x_{t-i} - \mu)^2\right)^{1/2},\end{aligned}$$

where α_0 and $\alpha_i \geq 0$, $1 \leq i \leq q$.
S_t: an AR(q) process

- The process is weakly stationary if all the roots of the characteristic equation associated with the ARCH parameters, $\alpha(B)$, lie outside the unit circle, i.e., if $\Sigma_{i=1}^{q} \alpha_i < 1$.

- Unconditional variance:

$$\alpha_0/(1 - \sum_{i=1}^{q} \alpha_i)$$

- Conditional variance:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2$$

or

$$\epsilon_t^2 = \alpha_0 + \alpha(B)\epsilon_{t-i}^2 + v_t.$$

- Weakness of ARCH models:
  - This model treats *positive* and *negative* returns in the same manner, because it depends on the square of the previous returns.
    In practice, it is well-known that for financial time series the prices respond differently to positive and negative returns.
  - The ARCH model is rather restrictive. For the ARCH(1) model of α_1^2 must be between 0 and 1/3. For higher-order ARCH models, the constraint is even stronger.

– ARCH models often over-predict the volatility, because they respond slowly to isolated large stocks to the return series.

Building ARCH Models

- Step 1: Remove the linear dependence of the return series and test for ARCH effects.
  - Mean Effect: Build an ARIMA model for the observed time series to remove any serial correlations in the data.
  - For most asset return series, this step amounts to remove the sample mean from the data if the sample mean is significantly different from zero.
  - Define $\epsilon_t = x_t - \mu_t$.
  - Examine the squared series ϵ_t^2 to check for conditional heteroscedasticity.
- Step 2: Order determination
  If conditional heteroscedasticity is detected, we use the PACF of ϵ_t^2 to determine the ARCH order.
- Step 3: Estimation
  - Conditional MLE
  - Software: S-plus, RATS
- Step 4: The fitted ARCH model is carefully examined and refined if necessary.

skewness, kurtosis, standardized residuals, and etc

Likelihood Function and ARCH Estimation

- Note that
$$\sigma_t = \left(\alpha_0 + \sum_{i=1}^{q} \alpha_i (x_{t-i} - \mu)^2\right)^{1/2}.$$
σ_t is a function of $x_{t-i} - \mu$ $(1 \le i \le q)$ and $q+1$ parameters α_i $(0 \le i \le q)$.

- Denote by ω the set of parameters μ, α_0, α_1, ..., α_q.

- The likelihood function for T observed returns is
$$L(x_1, x_2, \ldots, x_T|\omega)$$
$$= f(x_1|\omega) f(x_2|I_1, \omega) \cdots f(x_T|I_{T-1}, \omega).$$
Here $f(x_t|I_{t-1}, \omega)$ denotes the conditional density of x_t given the previous observations $I_{t-1} = \{x_1, x_2, \ldots, x_{t-1}\}$ and the parameter vector ω.

- Under the normality assumption, for $t > q$,
$$\begin{aligned} f(x_t|I_{t-1}, \omega) &= f(x_t|\sigma_t) \\ &= (\sqrt{2\pi}\sigma_t)^{-1} \exp\left[-\frac{1}{2}(x_t - \mu)^2/\sigma_t^2\right]. \end{aligned}$$

Or, the likelihood function of is

$$\prod_{t=q+1}^{T} \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left[-\frac{(x_t-\mu)^2}{2\sigma_t^2}\right] \times f(x_1, \cdots, x_q|\omega).$$

- The conditional maximum likelihood estimate ω for observations $q+1$ to T, maximizes

$$L_q(\omega) = \prod_{t=q+1}^{n} f(x_t|I_{t-1}, \omega).$$

The log likelihood function becomes

$$-\sum_{t=q+1}^{T} \left[\frac{1}{2}\ln(\sigma_t^2) + \frac{1}{2}\frac{a_t^2}{\sigma_t^2}\right],$$

where $\sigma_t^2 = \alpha_0^2 + \alpha_1 a_{t-1}^2 + \cdots + \alpha_q a_{t-m}^2$ can be evaluated recursively.

The GARCH Model

- The ARCH model often requires many parameters to adequately described the evolution of volatility of an asset return.
  For the monthly return series of S&P 500 index, an ARCH(9) model is needed for the volatility series.

- For a log return series x_t, the conditional mean μ_t can be adequately described by an ARMA model. Let $\epsilon_t = x_t - \mu_t$ be the mean-corrected log return.

- Generalized ARCH (GARCH(p,q)) process: Bollerslev (1986, 1988);

$$\begin{aligned} \epsilon_t &= \sigma_t U_t, \\ \sigma_t^2 &= \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2, \end{aligned}$$

  where $\{U_t\}$ is a sequence of iid random variables with mean 0 and variance 1, $\alpha_i \geq 0$, $\beta_j \geq 0$, and $\Sigma_{i=1}^{\max(p,q)}(\alpha_i + \beta_i) < 1$.
  If $p = 0$, it reduces to a pure ARCH(q) model.

- Let $\eta_t = \epsilon_t^2 - \sigma_t^2$. We get the following equivalent form:

$$\begin{aligned}\epsilon_t^2 &= \alpha_0 + \sum_{i=1}^{\max(p,q)} (\alpha_i + \beta_i)\epsilon_{t-i}^2 \\ &\quad + \eta_t - \sum_{j=1}^{q} \beta_j \eta_{t-j}. \qquad (5)\end{aligned}$$

It is an ARMA form for the squared series ϵ_t^2.

A GARCH model can be regarded as an application of the ARMA idea to the squared series ϵ_t^2.

- Using the unconditional mean of an ARMA model, we have

$$E(\epsilon_t^2) = \frac{\alpha_0}{\Sigma_{i=1}^{\max(p,q)}(\alpha_i + \beta_i)}$$

provided that the denominator of the above fraction is positive.

- The process is weakly stationary if and only if the roots of $\alpha(B) + \beta(B)$ lie outside the unit circle, i.e., $\alpha(1) + \beta(1) < 1$.

- A popular model for financial time series: GARCH$(1, 1)$ process

$$\sigma_t^2 = \alpha_0 + \alpha_1\epsilon_{t-1}^2 + \beta_1\sigma_{t-1}^2.$$

To be well-defined, $0 \leq \alpha_1, \beta_1 \leq 1$ and $\alpha_1 + \beta_1 < 1$.

- Volatility clustering: A large ϵ_{t-1}^2 or σ_{t-1}^2 gives rise to a large σ_t^2.

- Heavy tail: If $1 - 2\alpha_1^2 - (\alpha_1 + \beta_1)^2 > 0$, then

$$\frac{E(\epsilon_t^4)}{[E(\epsilon_t^2)]^2} = \frac{3[1 - (\alpha_1 + \beta)^2]}{1 - (\alpha_1 + \beta_1)^2 - 2\alpha_1^2} > 3.$$

- ϵ_t and σ_t^2 are strictly stationary if and only if

$$E(\log(\beta_1 + \alpha_1 U_t^2)) < 0.$$

Generalized ARCH

- I(ntergrated)GARCH(p, q):
  - If the AR polynomial of the GARCH representation has a unit root, we then have an IGARCH mode.
  - IGARCH models are unit-root GARCH models.
  - The key feature of IGARCH models is that the impact of past squared shocks η_{t-i} ($i > 0$) on ϵ_t^2 is persistent.
  - $\alpha(1) + \beta(1) = 1$
    Here I refers to integrated. See page 134.
  - Consider IGARCH$(1, 1)$ model.

    $$\sigma_t^2 = \alpha_0 + \beta_1 \sigma_{t-1}^2 + (1 - \beta_1)\epsilon_{t-1}^2,$$

    where $0 < \beta_1 < 1$.
- E(xponential)GARCH model:
  - Allow for asymmetric effects between positive and negative asset returns.
  - Nelson (1991)

    $$\log \sigma_t^2 = \alpha_0 + \alpha_1 f(\epsilon_{t-1}/\sigma_{t-1}) + \beta_1 \log \sigma_{t-1}^2$$

where

$$f(\epsilon_{t-1}/\sigma_{t-1}) = \theta_1\epsilon_{t-1}/\sigma_{t-1} + (|\epsilon_{t-1}/\sigma_{t-1}| - E|\epsilon_{t-1}/\sigma_{t-1}|).$$

- The asymmetry allows volatility to respond more rapidly to falls in a market than to corresponding rises. See page 137.

- Long memory volatility processes: The FIGARCH model.
  - FI refers to fractionally integrated.
  - See (4.10) in page 139.