Scand. J. of Economics 106(2), 165–185, 2004
DOI: 10.1111/j.1467-9442.2004.00360.x

Diebold, F.X. (2004),
"The Nobel Prize for Robert F. Engle,"
Scandinavian Journal of Economics, 106, 165-185, 2004.

The Nobel Memorial Prize for Robert F. Engle

*Francis X. Diebold**
University of Pennsylvania, Philadelphia, PA 19104, USA
fdiebold@sas.upenn.edu

I. Prologue

Engle's footsteps range widely. His major contributions include early work on band-spectral regression, development and unification of the theory of model specification tests (particularly Lagrange multiplier tests), clarification of the meaning of econometric exogeneity and its relationship to causality, and his later stunningly influential work on common trend modeling (cointegration) and volatility modeling (ARCH, short for *A*uto*R*egressive *C*onditional *H*eteroskedasticity).[1] More generally, Engle's cumulative work is a fine example of best-practice applied time-series econometrics: he identifies important dynamic economic phenomena, formulates precise and interesting questions about those phenomena, constructs sophisticated yet simple econometric models for measurement and testing, and consistently obtains results of widespread substantive interest in the scientific, policy and financial communities.

Although many of Engle's contributions are fundamental, I focus largely on the two most important: the theory and application of cointegration, and the theory and application of dynamic volatility models. Moreover, I discuss much more extensively Engle's volatility models and their role in financial econometrics, for several reasons. First, Engle's Nobel citation was explicitly "for methods of analyzing economic time series with time-varying volatility (ARCH)", whereas Granger's was for "for methods of analyzing economic time series with common trends (cointegration)". Second, the credit for creating the ARCH model goes exclusively to Engle, whereas the original cointegration idea was Granger's, notwithstanding Engle's powerful and well-known contributions to the development. Third, volatility models are a key part of the financial econometrics theme that defines

* Financial support from the Guggenheim Foundation and the National Science Foundation is gratefully acknowledged. Comments from Tim Bollerslev, David Hendry and Frank Schorfheide improved both the content and exposition, but they are in no way responsible for remaining flaws.

[1] See Engle (1973), Engle (1984), Engle, Hendry and Richard (1983), Engle and Granger (1987), and Engle (1982), respectively.

Engle's broader contributions, whereas cointegration has found wider application in macroeconomics than in finance. Last and not least, David Hendry's insightful companion review of Granger's work, also in this issue of the *Scandinavian Journal of Economics*, discusses the origins and development of cointegration in great depth.

In this brief and selective review, I attempt a description of "what happened and why" (to quote a popular American talk-show host). My approach has been intentionally to avoid writing a long review, as several extensive surveys of the relevant literatures already exist. In addition, I have de-emphasized technical aspects, focusing instead on the intuition, importance and nuances of the work.[2] In Section II, I discuss cointegration in the context of Engle's previous and subsequent work, which facilitates the extraction of interesting and long-running themes, several of which feature prominently in Engle's volatility models. I discuss the basic ARCH volatility model in Section III, and variations, extensions and applications in Section IV. I conclude in Section V.

II. Cointegration

Despite my declared intent to focus on volatility models, I begin by highlighting some aspects of cointegration. It is important to do so, because Engle's work on cointegration (as well as his earlier work, which I also touch upon) reveals central and sometimes subtle themes of the modeling approach that permeates all of his work, including that on volatility modeling. I give only the briefest possible intuitive sketch of the cointegration idea; those hungry for details may consult any of several fine surveys, including Watson (1995) and Johansen (1995).[3]

Basic Structure

Engle and Granger (1987) is the seminal paper on cointegration and perhaps the most cited paper in the history of econometrics, treating specification, representation, estimation and testing. The basic specification is simple. Consider, for example, two nonstationary series, x and y, both of which are integrated of order 1, or I(1), and therefore have no tendency to revert to any

[2] A recent interview with Engle in *Econometric Theory* provides additional insights into the development and impact of his work; see Diebold (2003). A complete list of Engle's publications is available on his web site (http://www.stern.nyu.edu/~rengle/).
[3] The collection edited by Engle and Granger (1991) also nicely illustrates many variations and extensions of the basic theory, as well as the wide variety of applications.

fixed mean (e.g., two random walks). In general, there is no way to form a weighted average of x and y to produce a stationary mean-reverting series integrated of order 0, or I(0), but in the special case where such a weighting does exist, we say that x and y are cointegrated. Suppose, for example, that,

$$\begin{aligned} x_t &= x_{t-1} + \epsilon_{xt} \\ y_t &= x_t + \epsilon_{yt}, \end{aligned} \tag{1}$$

where both ϵ_{xt} and ϵ_{yt} are white noise. Both x and y are I(1), but they are cointegrated, because $y_t - x_t = \epsilon_{yt}$, which is I(0). Cointegration causes variables to move together in the long run, which produces stationarity of the cointegrating combination (in this case $x_t - y_t$) despite the nonstationarity of its components. The long-run co-movements of variables in cointegrated systems are ultimately driven by their dependence on underlying common stochastic trends (in this case x_t).

Cointegration has strong and appealing economic motivation, because both theory and observation suggest that common stochastic trends arise naturally in macroeconomics and finance. In modern macroeconomic models, for example, the equilibrium dynamics of a large number of observed variables are typically driven by a small number of underlying shocks, or "factors" (e.g., technology shocks and preference shocks). This factor structure produces cointegration when the underlying shocks are I(1). The empirical facts accord with the theory: the great macroeconomic ratios (consumption/output, investment/output, etc.) often appear stationary despite the fact that their components clearly are not. Financial asset prices behave similarly: spreads or ratios (e.g., bond yield spreads across different maturities) often appear stationary whereas their components do not.

Cointegrated vector autoregressions (VARs) turn out to be closely related to the error-correction model long associated with the "LSE School" of econometrics, thus unifying models and communities. Error-correction models involve long-run equilibrium relationships augmented with short-run adjustment dynamics, in a way that captures the key idea that the sign and size of the current deviation from equilibrium should impact the future evolution of the system.

Consider again a two-variable system, and suppose that in long-run equilibrium y and x are related by $y = bx$, so that the deviation from equilibrium is $z = y - bx$. We acknowledge that the current deviation from equilibrium may impact the future by modeling the change in each variable as a function of lagged changes of all variables (to capture short-run dynamics in the standard VAR tradition), *and* the lagged value of z, the so-called error-correction term (to capture adjustment toward long-run equilibrium). Allowing for one lag of Δx and one lag of Δy, we write

$$\begin{aligned} \Delta x_t &= \alpha_x \Delta x_{t-1} + \beta_x \Delta y_{t-1} + \gamma_x z_{t-1} + u_{xt} \\ \Delta y_t &= \alpha_y \Delta x_{t-1} + \beta_y \Delta y_{t-1} + \gamma_y z_{t-1} + u_{yt}. \end{aligned} \tag{2}$$

So long as one or both of γ_x and γ_y are nonzero, the system is very different from a VAR in first differences; the key distinguishing feature is the inclusion of the error-correction term, so that the current deviation from equilibrium affects the future evolution of the system. As can be readily verified, for example, the cointegrated system (1) has error-correction representation

$$\begin{aligned} \Delta x_t &= u_{xt} \\ \Delta y_t &= -z_{t-1} + u_{yt}, \end{aligned} \tag{3}$$

where $u_{xt} = \epsilon_{xt}$, $u_{yt} = \epsilon_{xt} + \epsilon_{yt}$ and $z_t = y_t - x_t$. This illustrates the central and much more general result alluded to above, introduced by Granger (1981) and refined by Engle and Granger (1987): cointegrated VARs and error-correction models are equivalent—a VAR is cointegrated if and only if it has an error-correction representation.

Cointegration in the Context of Engle's Broader Research Agenda

Engle's work on cointegration lends insight into several central themes that run throughout his research, including his research on volatility modeling. First, it shows a keen interest in differences in economic dynamics across frequencies. That is, one sees a clear distinction between "long run" vs. "short run" dynamics, "permanent" vs. "transitory" shocks, and so on. A cointegrated VAR, for example, is effectively an econometric implementation of the idea of long-run equilibrium supplemented with short-run dynamics, producing higher coherence among variables at low frequencies than at high. Engle's early work also emphasizes this idea of the decomposition of dynamics by frequency, with possibly different models for different frequencies, as made starkly explicit in Engle's band-spectral regression, as in e.g., Engle (1973), and unobserved-components models of trend, seasonal and cyclical dynamics, as in e.g., Engle (1978). The theme continues to the present. For example, Engle and Smith's (1999) "StoPBreak" model allows for endogenous time-variation in the persistence of the effects of shocks. Similarly, Engle and Lee's (1999) unobserved components model for volatility facilitates decomposition of volatility fluctuations into permanent and transitory components. Moreover, the analysis of ultra-high-frequency data pioneered by Engle and Russell (1998) and Engle (2000) is intrinsically concerned with aspects of high-frequency dynamics, such as microstructure noise, intra-day calendar effects, and the arrival of quotes and trades in real time.

Second, awareness of the practical relevance yet econometric intractability of high-dimensional systems, and the economic appeal and econometric convenience of factor structure as a potential solution, permeate Engle's work. Factor structure is economically appealing because macroeconomic and financial theorizing leads naturally to environments with far fewer underlying sources of dynamic variation than variables observed.[4] Factor structure is econometrically convenient because it implies that high-dimensional dynamic modeling of hundreds or even thousands of variables, which in general would be empirically intractable (indeed, absurdly so), can in fact be done quite simply: beneath the superficial appearance of a plethora of series, there are just a few parameters to be estimated. Early on, Engle worked productively with dynamic factor models; see e.g., Engle and Watson (1981) and Engle, Lilien and Watson (1985). Models of cointegration are also very explicit dynamic factor models, in which the common factors are the I(1) common trends. Engle's subsequent work on "common features", as in e.g., Engle and Kozicki (1993) and Engle and Vahid (1993), is a direct outgrowth of models of cointegration and hence factor models.[5] Much of Engle's work on high-dimensional volatility modeling proceeds by invoking factor structure and closely related ideas; see e.g., Bollerslev and Engle (1993).

Third, Engle's work shows a preference for models that are parametric, parsimonious and indeed typically linear, estimated by (Gaussian) maximum likelihood. The themes of unobserved components modeling, factor models, cointegration and common features all support this claim, as does Engle's prowess in state space modeling, Kalman filtering and Gaussian likelihood evaluation via prediction-error decomposition; see e.g., Engle and Watson (1985). Engle's volatility models, although nonlinear because they involve variances, which are expectations of squares, are effectively just linear (and typically conditionally Gaussian) models of squared variables. Hence, when appropriately interpreted, they also fit the linear/Gaussian mold.

Fourth, and intimately related to Engle's preference for parametric, parsimonious models, is the long-standing concern with *forecasting*—a key task in practical policy, financial and business applications—that runs throughout Engle's work. Parsimonious models guard against in-sample overfitting,

[4] For example, as has already been mentioned, macroeconomic models typically involve just a few shocks, which determine a much larger number of variables in equilibrium. Similarly, financial asset pricing models are factor models; the classic CAPM is a one-factor model, and more recent generalizations are three- or four-factor models.

[5] If two variables have property X but there exists linear combination that does not, we say that they have common feature X.

or "data mining", which typically produces models that fit well historically but perform miserably in out-of-sample forecasting. Hence the tightly parametric—and perhaps seemingly naïve—approach favored by Engle is intentional, and not at all naïve. Rather, his models are simultaneously sophisticated and simple, precisely in keeping with Zellner's (1992) version of the "KISS principle" (*K*eep *I*t *S*ophisticatedly *S*imple).

Concern with forecasting has not only influenced Engle's general approach, but also produced a variety of specific forecasting contributions. In the cointegration context, for example, Engle and Yoo (1987) show that as the forecast horizon lengthens, optimal forecasts of cointegrated variables will satisfy the cointegrating relationship(s) arbitrarily closely. In the volatility context, the very essence of Engle's contribution is allowance for forecastable forecast-error variances. From a traditional forecasting viewpoint, this allows confidence intervals on forecasts appropriately to expand and contract as volatility fluctuates. The explosion of interest in Engle's volatility models in financial contexts, moreover, reflects interest not only in improved confidence intervals around point forecasts, but also, and crucially, intrinsic interest in volatility forecastability, which profoundly influences fundamental financial tasks, including asset pricing, portfolio allocation and risk management.

III. ARCH

Good dynamic economic theorists, like all good theorists, naturally want to get the facts straight before theorizing. Time-series tools are tailor-made for providing precise quantification of dynamics; hence the explosive growth of time-series econometric methods and applications in the last forty years. A short list of active sub-fields—many of which were touched upon in Section II—includes vector autoregressive macroeconomic modeling, dynamic factor models, causality, integration and persistence, cointegration, seasonality, unobserved-components models, state-space representations and the Kalman filter, regime-switching models, nonlinear dynamics and nonlinear filtering. But all of those topics concern aspects of linear or nonlinear conditional *mean* dependence. What about conditional *variance* dependence—that is, what about conditional heteroskedasticity?

Historically, economists viewed heteroskedasticity as largely a cross-sectional, as opposed to a time-series, phenomenon. Hence economists are almost always introduced to heteroskedasticity in cross-sectional contexts, such as when the variance of a cross-sectional regression disturbance depends on one or more of the regressors. A classic example is Engel curve estimation, in which the variance of the disturbances in expenditure equations tends to grow with income. It turns out, however, that heteroskedasticity is pervasive in the time-series context of financial asset returns, in which volatility clustering (that is, contiguous periods of high or low volatility) features prominently.[6] Moreover, it also turns out that acknowledgment and modeling of return volatility dynamics is crucially important for financial economics. Unfortunately, cross-sectional heteroskedasticity models are ill-suited to dynamic environments, and specification and econometric treatment of an appropriate dynamic volatility model involves many subtleties and is far from a trivial extension.

Engle's ARCH model and subsequent volatility modeling research program provided a workable and elegant solution, solving many problems and stimulating a huge amount of related research that advanced not only the econometrics of dynamic volatility and correlation modeling, but also forecasting, asset pricing, portfolio allocation, risk management, market microstructure modeling, duration modeling and ultra-high-frequency data analysis. Hence any list of the last forty years' crucial time-series econometric developments must also include the dynamic volatility models pioneered by Engle. Indeed, the originality, importance and influence of Engle's ARCH models put them at the top of the list, tied for first place with Granger's cointegration models.

The Basic ARCH Model

Engle (1982), another of the most widely cited papers in the history of econometrics, started it all. Surveys of the massive ensuing literature include Bollerslev, Chou and Kroner (1992), Gourieroux (1992), Bera and Higgins (1993), Bollerslev, Engle and Nelson (1994), Diebold and Lopez (1996), Andersen, Bollerslev and Diebold (2004), and Diebold (2004). The collection edited by Engle (1995) nicely illustrates many variations and extensions of the basic theory, as well as a wide variety of applications.

Engle's basic contribution cuts to the core of modern time-series analysis. To understand it, recall that Wold's (1938) celebrated decomposition theorem establishes that any linearly indeterministic covariance stationary process $\{x_t\}$ can be written as a one-sided moving average of (weak) white noise innovations. In an obvious notation,

$$y_t = \mu + B(L)\,\epsilon_t \tag{4}$$

$$\epsilon_t \sim WN(0, \sigma^2), \tag{5}$$

[6] In what follows, I use interchangeably the terms volatility clustering, volatility dynamics, volatility persistence, volatility forecastability and volatility predictability.

where $B(L) = \Sigma_{i=0}^{\infty} b_i L^i$, $\Sigma_{i=0}^{\infty} b_i^2 < \infty$ and $b_0 = 1$. The uncorrelated innovation sequence $\{\epsilon_t\}$ need not be Gaussian and therefore need not be independent. But in the Gaussian case, the lack of correlation does of course imply independence, and for some four decades after Wold, the profession implicitly adopted a Gaussian perspective, focusing almost exclusively on linear models of conditional mean dependence with i.i.d. innovations.

To appreciate the limitations of the traditional linear/i.i.d. model, consider a Wold representation with i.i.d., as opposed to merely uncorrelated, innovations. The *unconditional* mean and variance are

$$E[y_t] = \mu \quad \text{and} \quad E\left[(y_t - \mu)^2\right] = \sigma_\epsilon^2 \sum_{i=0}^{\infty} b_i^2,$$

both of which are constant, as must be the case under covariance stationarity. Now consider the *conditional* mean and variance. The conditional mean naturally adapts to the conditioning information and is given by

$$E[y_t \,|\, \Omega_{t-1}] = \mu + \sum_{i=1}^{\infty} b_i \, \epsilon_{t-i}, \tag{6}$$

where $\Omega_{t-1} = \{\epsilon_{t-1}, \epsilon_{t-2}, \ldots\}$. One might hope that the conditional variance would adapt to the conditioning information as well, but unfortunately it does not. Instead, it is constant:

$$E\left[(y_t - E[y_t \,|\, \Omega_{t-1}])^2 \,|\, \Omega_{t-1}\right] = \sigma^2. \tag{7}$$

This asymmetric treatment of conditional mean and conditional variance dynamics—allowance for flexible linear conditional mean dynamics but no allowance for conditional variance dynamics—proves crippling in financial contexts.

Now let us introduce Engle's (1982) ARCH process by considering a richer Wold representation, with ARCH innovations:

$$y_t = \mu + B(L)\, \epsilon_t \tag{8}$$

$$\epsilon_t \,|\, \Omega_{t-1} \sim N(0, h_t) \tag{9}$$

$$h_t = \omega + \gamma(L)\, \epsilon_t^2, \tag{10}$$

where $\omega > 0$, $\gamma(L) = \Sigma_{i=1}^{\infty} \gamma_i L^i$, $\gamma_i \geqslant 0$ and $\Sigma_{i=1}^{\infty} \gamma_i < 1$, which satisfies Wold's theorem (in particular, the innovation ϵ_t is serially uncorrelated) while nevertheless explicitly incorporating dynamic conditional heteroskedasticity. Both the unconditional mean and variance are constant, again as required for covariance stationarity. However, both the conditional mean *and* the conditional variance are time-varying:

$$E[y_t \mid \Omega_{t-1}] = \mu + \sum_{i=1}^{\infty} b_i \, \epsilon_{t-i} \tag{11}$$

$$E\left[(y_t - E[y_t \mid \Omega_{t-1}])^2 \mid \Omega_{t-1}\right] = \omega + \sum_{i=1}^{\infty} \gamma_i \, \epsilon_{t-i}^2. \tag{12}$$

This model desirably treats conditional mean and variance dynamics in a symmetric fashion, by allowing for movement in each.

Interestingly, the serial correlation in volatility afforded by the ARCH conditional variance function (10) had been observed empirically, decades earlier. In his classic analysis of returns on speculative markets, Mandelbrot (1963, p. 418) was very clear, noting that "large changes tend to be followed by large changes—of either sign—and small changes tend to be followed by small changes". But Mandelbrot emphasized the unconditional nonnormality of returns, rather than their volatility clustering. For the next twenty years, volatility clustering generated little interest; instead, the literature focused primarily on the nonnormality, and in particular the fat tails, of unconditional return distributions.[7]

In a huge conceptual advance, Engle's work changed the situation radically, by proposing a rigorous, probabilistic, likelihood-based framework for volatility estimation and forecasting. Moreover, Engle was not content with purely theoretical contributions. Beginning with the original application to UK inflation in Engle (1982), he actively pursued empirical applications in the macroeconometric context of volatility dynamics in the price level, including Engle (1983) and Engle, Granger and Robins (1986). Crucially, when work such as Diebold (1988) and Bollerslev, Engle and Wooldridge (1988) revealed very strong and pervasive ARCH effects in foreign exchange and equity markets (in contrast to aggregate inflation and other macroeconomic variables, for which it had by then become clear that the evidence was mixed), Engle quickly switched from macroeconomic to financial economic applications. He recognized, moreover, that many parts of financial economics would change fundamentally in the presence of

[7] See, for example, Fama (1965), Blattberg and Gonnedes (1974), and references therein.

forecastable return volatility, and he followed through for twenty years, building major parts of the new field of financial econometrics.[8]

The appeal of the linear/ARCH framework (8)–(10) is essentially three-fold. First, it reflects an appreciation of the fact that the linear/i.i.d. framework embraces only a small part of Wold's vision, because it neglects the possibility of nonlinear dependence in ϵ.

Second, it reflects two implicit judgments, which turn out to be correct in many economic and financial contexts. The first is that the linear conditional mean dynamics associated with $B(L)$ (or ARMA approximations to $B(L)$) are likely to be roughly adequate approximations to the true conditional mean dynamics, in which case nonlinear dependence in ϵ_t would come primarily from higher-ordered conditional moment dependence rather than from neglected nonlinear conditional mean dependence. The second implicit judgment is that the key neglected conditional moment is the conditional variance.

Third, it uses a natural and powerfully simple strategy for modeling conditional variance dynamics, which precisely parallels the Wold representation for the conditional mean, namely modeling the conditional variance—the conditionally expected value of ϵ_t^2—as a linear function of ϵ_{t-1}^2, ϵ_{t-2}^2, and so on.

Structure and Properties of the ARCH Process

Many features of the ARCH process are novel relative to more traditional setups. For example, and most obviously, the fact that the conditional variance, h_t, is now a serially correlated random variable has many implications. For example, unlike the traditional linear/i.i.d. model, in which the conditional prediction error variance depends only on the forecast horizon, it now depends on the conditioning information set, and this dependence can be exploited to produce accurate volatility forecasts, which have numerous uses in asset pricing, portfolio allocation and risk management.

The conditional variance (12) of an ARCH process is a deterministic function of the conditioning information. This is entirely natural and appropriate, just as, for example, the conditional mean (11) is also a deterministic function of conditioning information. One can, however, entertain the possibility that a separate stochastic "volatility shock" is also operative, which leads directly to the class of so-called stochastic volatility models, which itself has been the focus of a huge research effort since the mid-1980s, motivated by and spurred onward by Engle's ARCH research.[9]

[8] For interesting commentary on the furious pace of development, both retrospective and prospective, see Engle (2002b).
[9] See Taylor (1986), Ghysels, Harvey and Renault (1996) and Shephard (2004).

In addition to the fact that the ARCH conditional variance is a deterministic function of past observations, the ARCH process is specified directly in terms of the conditional density, (9). Hence, for any given set of parameter values, the Gaussian likelihood can be evaluated trivially via a prediction-error decomposition—that is, by recognizing that the likelihood, which is the joint density of the observed data, is simply the product of the conditional densities (9)—and then maximized numerically. Hence, estimation of ARCH models is in principle straightforward, in sharp contrast to the estimation of stochastic volatility models, the estimation of which is notoriously challenging, due to the latency of the separate volatility shock.

It is instructive to view ARCH from the vantage point of volatility proxy construction. One obvious proxy for the true underlying volatility, h_t, is simply ϵ_t^2. Define the corresponding measurement error as $\nu_t = \epsilon_t^2 - h_t$. Unfortunately, the variance of ν_t (the noise) is generally large relative to the variance of h_t (the signal). One would like to find a way to reduce the noise, perhaps by smoothing, which is precisely what the ARCH model does: it replaces the noisy volatility proxy, ϵ_t^2, with a much less noisy proxy, a weighted average of the history of ϵ_t^2, $\gamma(L)\,\epsilon_t^2$. Indeed, the ARCH model is effectively an autoregressive model for ϵ_t^2, and the ARCH conditional variance function is simply the conditional mean function of the autoregressive model for ϵ_t^2. To see this, substitute $h_t = \epsilon_t^2 - \nu_t$ into the ARCH conditional variance function (10), which produces the autoregressive representation for ϵ_t^2,

$$\epsilon_t^2 = \omega + \gamma(L)\,\epsilon_t^2 + \nu_t. \tag{13}$$

Hence the name "ARCH".

It can also be shown that the implied unconditional distribution of even a conditionally Gaussian ARCH process is leptokurtic (and, of course, conditionally fat-tailed ARCH processes are even more unconditionally leptokurtic), and that the leptokurtosis vanishes under temporal aggregation.[10] The three key properties of ARCH processes—volatility clustering, fat-tailed unconditional distributions and reduction of those fat tails under temporal

[10] The unconditional leptokurtosis of ARCH processes follows from the persistence in conditional variance, which produces the clusters of "low volatility"and "high volatility" episodes associated with observations in the center and in the tails of the unconditional distribution. The reduction of leptokurtosis under temporal aggregation follows from the facts that (1) a low-frequency change is simply the sum of the corresponding high-frequency changes; for example, an annual change is the sum of the internal quarterly changes, each of which is the sum of its internal monthly changes, and so forth, and (2) Gaussian central limit theorems hold for sums of covariance stationary ARCH processes, as shown by Diebold (1988).

aggregation—match closely the properties of actual high-frequency financial asset returns. Hence the ARCH process provides an appealing unification of the three key properties—previously believed to be disparate—of financial asset return dynamics and distributions.

IV. Variations, Extensions and Applications

The ARCH idea is both simple and very powerful, and the huge and vibrant literature that it stimulated continues its brisk pace of advance. Here, I attempt to provide a feel for the flavor and nuances of that literature. Variations and extensions of the basic ARCH model are nicely intertwined with applications, so I discuss them simultaneously.

GARCH

Engle's (1982) breakthrough launched a huge research program, with both theoretical and empirical fronts. On the theoretical front, steady progress was made on important tasks such as refining the associated asymptotic theory, including quasi-maximum likelihood estimation and testing, characterization of moment structure, calculation of diffusion limits, and so on. Monte Carlo explorations of small-sample properties of various estimators and tests also proved informative, as in Engle, Hendry and Trumble (1985).

On the empirical front, applications were initially hindered by the fact that estimation of the ARCH(∞) model (10) is of course infeasible in practice, and simple low-ordered truncation of $\gamma(L)$ invariably appeared inadequate in financial market contexts, in which volatility tends to be highly persistent. Hence, early empirical studies often adopted high-ordered ARCH(p) models, truncated at some point for tractability rather than realism, often with constraints on the lag weights such as linear decay, again for tractability rather than realism. Bollerslev's (1986) GARCH model quickly remedied this unfortunate situation.

To understand and appreciate GARCH, recall that ARMA processes can be viewed as approximations to infinite-ordered autoregressive processes, in which infinite-ordered autoregressive lag operator polynomials are approximated by ratios of low-ordered lag operator polynomials. This is useful when modeling conditional mean dynamics in economic time series if the required autoregressive lag operator polynomial turns out to be very high-ordered, which sometimes happens. The accuracy of such rational approximations, which motivates the use of ARMA processes, led to Bollerslev's (1986) brilliant parallel proposal for ARCH processes: approximate the infinite-ordered ARCH autoregressive lag operator polynomial $\gamma(L)$ by the ratio of two low-ordered lag operator polynomials, $\alpha(L)$ of order p, say, and $\beta(L)$ of order q. Doing so produces the GARCH(p, q) process,[11]

$$h_t = \omega + \alpha(L)\,\epsilon_t^2 + \beta(L)\,h_t, \tag{14}$$

where $\alpha(L) = \Sigma_{i=1}^{p}\,\alpha_i\,L^i, \beta(L) = \Sigma_{i=1}^{q}\,\beta_i\,L^i, p, q < \infty$, $\omega > 0$ and $\alpha(1) + \beta(1) < 1$. GARCH is to ARCH (for conditional variance dynamics) as ARMA is to AR (for conditional mean dynamics). Indeed, it can be shown that if a series ϵ_t is GARCH(p, q), then ϵ_t^2 is ARMA((max(p, q), p). Interestingly, the simple GARCH(1, 1) model,

$$h_t = \omega + \alpha\,\epsilon_{t-1}^2 + \beta\,h_{t-1}, \tag{15}$$

is so often empirically adequate that it has achieved something of a canonical status.

Asymmetric Response and the Leverage Effect

Numerous alternative functional forms for the conditional variance function (14) have been suggested, several of which allow for an asymmetric volatility response to the lagged squared innovations, depending on their sign, so that the effect of a negative innovation on volatility may differ from that of a positive innovation. This allowance for asymmetric response proves useful for modeling the "leverage effect" often observed in stock returns, in which a negative innovation boosts volatility by more than a positive innovation of the same absolute magnitude. Engle and Ng (1993) go much farther, with a contribution that has subsequently found use in numerous contexts. Instead of allowing for a single threshold, with differing effects on each side of zero, they allow for piecewise linear "news impact curves" with many kinks, rather than just one kink at zero. Visual examination of such curves often proves revealing and is now a standard empirical tool.

Non-Gaussian Conditional Densities

It is sometimes desirable to allow for non-Gaussian conditional densities in the GARCH model, because it is sometimes found that the Gaussian GARCH model does not explain all of the leptokurtosis in asset returns.

[11] GARCH stands for generalized ARCH. It is actually a specialization, not a generalization, of the ARCH(∞) model (10), although it is of course a generalization of a finite-ordered ARCH process. The point is that although it is a specialization of (10), it is tremendously useful empirically, whereas (10) is not.

That is, although conditionally Gaussian GARCH processes are unconditionally fat-tailed, their implied unconditional densities are sometimes not fat-tailed enough to match the unconditional distribution of actual high-frequency asset returns. Hence, various proposals have been made for allowing for non-Gaussian conditional densities. One of the most interesting is by Engle and González-Rivera (1991), who propose a semiparametric methodology in which the conditional variance function is parametrically specified and consistently estimated by quasi-ML in the usual fashion, after which the conditional density is then estimated nonparametrically from the standardized returns.

Very High Volatility Persistence

GARCH volatility persistence is governed by the dominant root of $1 - \alpha(z) - \beta(z)$, and roots near unity are common in high-frequency financial data. Thus, motivated by the 1980s research boom in conditional-mean unit-root dynamics, Bollerslev and Engle (1986) proposed and explored a special case of the GARCH model, the so-called integrated GARCH (IGARCH) model. Roughly, IGARCH is to GARCH as ARIMA is to ARMA, although there are some interesting twists (e.g., IGARCH processes remain strictly stationary).

Although conditional variance dynamics are often empirically found to be highly persistent, it is difficult to ascertain whether they are actually integrated.[12] However, circumstantial evidence against IGARCH arises from consideration of temporal aggregation. Due to the infinite unconditional second moment of IGARCH processes, temporal aggregation does not produce unconditional normality, whereas actual series displaying GARCH effects seem to approach normality when temporally aggregated.

Thus began the search for richer models of highly persistent volatility, which actively continues to this day. Prominently featured in the ensuing literature are models of long memory volatility, which are covariance stationary but display persistence stronger than that achievable within the covariance stationary GARCH class. One route to long memory is through superposition of conventional short-memory volatility models—that is, through multi-factor volatility models—as emphasized for example in Barndorff-Nielsen and Shephard (2001). Interestingly, the permanent-transitory component model of Engle and Lee (1999)—written many years before it was eventually published—is roughly in that spirit.

Another route to long memory is completely nonparametric estimation of volatility from the high-frequency return data that is increasingly available.

[12] This difficulty obviously parallels the literature on unit roots and trend breaks; see Stock (1994) for a fine overview.

This approach, which in principle can produce arbitrarily accurate volatility and correlation estimates, is the focus of the rapidly growing "realized volatility" literature, surveyed by Andersen, *et al.* (2004). The methods of this new literature—and its emphasis on long memory—very much stand on the shoulders of Engle's path-breaking work, in this case Ding, Granger and Engle (1993), who showed clear evidence of long memory in daily absolute returns.

Multivariate Models, Dimensionality Reduction and Factor Structure

From the beginning, it was widely recognized that extension of the GARCH model to the multivariate case was crucial, due to the centrality of correlation structures in financial applications. It took ten years of hard work for Engle and Kroner (1995) to provide a fully general treatment, with what is now known as the BEKK model, in reference to Baba, Engle, Kraft and Kroner (Baba and Kraft made important early contributions but eventually withdrew to pursue other interests). The BEKK model is useful for predicting covariance matrices in low-dimensional situations. In all but the lowest-dimensional situations, however, unrestricted volatility models quickly become unworkable, due to the huge number of parameters that must be estimated by numerical optimization. Not surprisingly, then, various strategies for facilitating high-dimensional volatility modeling emerged at the top of the research agenda.

Initial ad hoc attempts at dimensionality reduction such as the "diagonal model" of Bollerslev, Engle and Wooldridge (1988), although major advances at the time, ultimately proved less than fully satisfying.[13] What was needed was *economic* guidance for dimensionality reduction. Factor structure, which is at the heart of modern financial asset pricing, delivered the goods, as formalized in the "factor GARCH" model of Bollerslev and Engle (1993). Factor GARCH is effectively a model of cointegration in variance or, more generally, a model of a common feature (persistence) in variance, as for example in Engel and Susmel's (1993) application to common volatility in international equity markets. Note the beautiful continuity and coherence in the sequence from integration to cointegration, to common features, to integrated GARCH, to cointegration in variance, to co-persistence in variance. Progress continues to be made, as for example in Engle (2002a, 2002b).

[13] In a diagonal model, variances depend only on own lagged variances and own lagged squared innovations, and covariances depend only on own lagged covariances and own lagged cross-products.

Time-varying Risk Premia, Options Pricing, and the Value of Volatility Forecasts

When financial asset return volatility varies, one naturally suspects that time-varying risk premia may be operative, and a huge literature has used GARCH methods to quantify and test hypotheses concerning risk premia in stock, bond, foreign exchange and commodities markets. Bollerslev, Engle and Wooldridge (1988) provide an early and important stock market application, implementing a conditional capital asset pricing model with time-varying covariance structure, and ultimately time-varying market betas. In a parallel series of path-breaking bond market applications, Engle models time-varying risk premia in the term structure of bond yields using the Engle, Lilien and Robins (1987) ARCH-in-mean model, which allows the conditional variance of a yield also to affect its conditional mean. Successful modeling of term premia promotes successful forecasting of future spot interest rates, as in Engle, Ng and Rothschild (1990), despite the fact that forward interest rates do not provide good forecasts (precisely because of the time-varying risk premia).[14] Moreover, GARCH perspectives also produce very general approaches to empirical asset pricing, as with the empirical pricing kernels of Engle and Rosenberg (2002).

GARCH methods prove useful not only for spot asset pricing, but also for derivative asset pricing; indeed, the impressive cross-fertilization between GARCH econometrics and asset pricing is even more pronounced in options contexts. The key object in option pricing is a forecast of the underlying asset's volatility path over the course of the option's life. This suggests that GARCH models should be useful for improved pricing of options and related derivative products, and subsequent research supports that conjecture. One strand of literature compares the returns on options-based trading strategies implemented using GARCH volatility forecasts to those from strategies implemented using alternative volatility forecasts; see e.g., Engle, Hong, Kane and Noh (1993) and Engle, Kane and Noh (1994). The results generally indicate superior risk-adjusted returns from the GARCH-based strategies, implying superiority of the GARCH volatility forecasts under a highly relevant loss function.

Another strand of the literature examines not the GARCH pricing of options portfolios, but rather the complementary issue of GARCH *hedging* of options portfolios, with similar success. For example, Engle and Rosenberg (1995, 2000) show that delta–gamma hedging of options (using GARCH gamma) outperforms traditional delta hedging, both in simple situations involving a single maturity of fluctuating volatility, and in more complicated situations involving complex fluctuations of the entire volatility term structure.

[14] See also Engle and Rothschild (1992).

Although the discussion thus far has focused on using GARCH econometrics for superior results in financial markets, one can reverse direction, using financial markets for potentially improved GARCH econometrics, an idea that has also met with success. For example, Engle and Mustafa (1992) develop an estimation procedure for GARCH models in the spirit of indirect inference; they proceed by finding the GARCH parameters that make implied model-based options prices maximally close to observed market options prices.

Financial Market Microstructure

A central concern in international finance is the cross-market transmission of shocks. The volatility models pioneered by Engle facilitate the study of cross-market transmission of *volatility* shocks. In an important series of papers that memorably phrase the question in terms of "heat waves vs. meteor showers", Engle investigates whether volatility bursts tend to stay in one place as the earth turns (like a heat wave) or whether they tend to move across the earth from market to market (like the appearance of a meteor shower), documenting strong volatility spillovers across international equity markets; see e.g., Engle, Ito and Lin (1990, 1992, 1994) and Engle and Susmel (1994).

Focus on cross-market volatility interactions in the world economy leads naturally to consideration of cross-asset interactions in a single market, such as individual stocks on an exchange, which, particularly when one approaches the transactions level, raises a host of important questions in the general area of financial market microstructure. In addition, when one considers that transactions data are now quite commonly available due to improvements in electronic markets, data processing and data storage, it is easy to understand the recent explosion of interest in the analysis of ultra-high-frequency return data and its links to microstructural concerns. Not surprisingly, Engle has also led this charge, producing important tools for analyzing irregularly spaced transactions data, with emphasis on modeling and forecasting intra-quote and intra-trade durations, quote–trade linkages, price impacts of trades and market depth.

Engle and Russell (1998) provide a key contribution, the autoregressive conditional duration (ACD) model.[15,16] The ACD model assumes Poisson transactions arrivals, in which the Poisson intensity rate varies in an autoregressive fashion. Maximum likelihood estimation of the model is

[15] See also Engle and Russell (1997), Engle (2000) and Engle and Russell (2004).

[16] Subsequent and important work includes extensions to modeling the bivariate point process of trades and quotes, as in Engle and Lunde (2003) and Engle and Patton (2004), and related analyses of aspects of liquidity, including market depth and the price impacts of trades, as in Engle and Lange (2001) and Dufour and Engle (2000).

readily achieved by prediction-error decomposition of the likelihood, precisely as with the GARCH model. As subsequently became clear, moreover, the ACD and GARCH models are related in much more than the mere mechanics of estimation. In particular, the persistence of the time-varying ACD transaction arrival intensity produces transactions clustering, which corresponds to volatility clustering in calendar time. Roughly speaking, then, ACD duration clustering in transaction time corresponds to ARCH volatility clustering in calendar time!

V. Epilogue

In the past forty years, time-series econometrics developed from infancy to relative maturity. A significant part of that development is due to Robert F. Engle, whose rigorous yet immensely practical work is distinguished by exceptional creativity in empirical dynamic modeling, particularly volatility modeling in financial contexts. Engle's contributions spawned a huge amount of additional work—indeed they spawned new literatures, both theoretical and applied—with contributions not only by leading established econometricians worldwide, but also by a large army of Engle's graduate students, a list of whom now reads like a who's who in the younger generation of time-series econometricians. Engle's contributions form a key pillar of a new discipline, financial econometrics, which is now taught routinely in economics departments and business schools worldwide, and they have been spectacularly successful in bridging academics and industry, influencing modern asset pricing, portfolio allocation and risk management.

The Nobel Memorial Prize for Engle and Granger is a long-awaited and well-deserved prize for time-series econometrics in general—appropriately paralleling the earlier Heckman–McFadden prize for microeconometrics—and for two of its most important contributions in particular: models of volatility (ARCH) and common trends (cointegration). Of crucial importance is the fact that the Engle–Granger prize, like the Heckman–McFadden prize, is for economic research that is not only scientifically path-breaking, but also of great and immediate practical relevance. It is no exaggeration to assert that millions of people, worldwide, have been made better off by the work of Robert F. Engle and Clive W. J. Granger.

References

Andersen, T. G., Bollerslev, T. and Diebold, F. X. (2004), Parametric and Nonparametric Volatility Measurement, forthcoming in L. P. Hansen and Y. Ait-Sahalia (eds.), *Handbook of Financial Econometrics*, North-Holland, Amsterdam.

Barndorff-Nielsen, O. E. and Shephard, N. (2001), Non-Gaussian Ornstein–Uhlenbeck-based Models and Some of Their Uses in Financial Economics (with discussion), *Journal of the Royal Statistical Society B 63*, 167–241.

Bera, A. K. and Higgins, M. L. (1993), ARCH Models: Properties, Estimation and Testing, *Journal of Economic Surveys 7*, 305–362.

Blattberg, R. and Gonnedes, N. (1974), A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices, *Journal of Business 47*, 244–280.

Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics 31*, 307–327.

Bollerslev, T., Chou, R. Y. and Kroner, K. F. (1992), ARCH Modeling in Finance: A Selective Review of the Theory and Empirical Evidence, *Journal of Econometrics 52*, 5–59.

Bollerslev, T. and Engle, R. F. (1986), Modeling the Persistence of Conditional Variances (with discussion), *Econometric Reviews 5*, 1–50.

Bollerslev, T. and Engle, R. F. (1993), Common Persistence in Conditional Variances, *Econometrica 61*, 166–187.

Bollerslev, T., Engle, R. F. and Nelson, D. B. (1994), ARCH Models, in R. F. Engle and D. McFadden (eds.), *Handbook of Econometrics*, *Volume IV*, North-Holland, Amsterdam.

Bollerslev, T., Engle, R. F. and Wooldridge, J. M. (1988), A Capital Asset Pricing Model with Time Varying Covariances, *Journal of Political Economy 95*, 116–131.

Diebold, F. X. (1988), *Empirical Modeling of Exchange Rate Dynamics*, Springer-Verlag, New York.

Diebold, F. X. (2003), The *ET* Interview: Professor Robert F. Engle, *Econometric Theory 19*, 1159–1193.

Diebold, F. X. (2004), *Measuring and Forecasting Financial Market Volatilities and Correlations*, W. W. Norton, New York, forthcoming.

Diebold, F. X. and J. A. Lopez (1996), Modeling Volatility Dynamics, in K. V. Hoover (ed.), *Macroeconometrics: Developments, Tensions and Prospects*, Kluwer Academic Press, Boston, MA, 427–472.

Ding, Z., Granger, C. W. J. and Engle, R. F. (1993), A Long Memory Property of Stock Market Returns and a New Model, *Journal of Empirical Finance 1*, 83–106.

Dufour, A. and Engle, R. F. (2000), Time and the Price Impact of a Trade, *Journal of Finance 55*, 2467–2498.

Engle, R. F. (1973), Band Spectrum Regression, *International Economic Review 15*, 1–11.

Engle, R. F. (1978), Estimating Structural Models of Seasonality, in A. Zellner (ed.), *Seasonal Analysis of Economic Time Series*, US Department of Commerce, Bureau of Census, Washington, DC.

Engle, R. F. (1982), Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation, *Econometrica 50*, 987–1008.

Engle, R. F. (1983), Estimates of the Variance of US Inflation Based on the ARCH Model, *Journal of Money, Credit and Banking 15*, 286–301.

Engle, R. F. (1984), Wald, Likelihood Ratio and Lagrange Multiplier Tests in Econometrics, in Z. Griliches and M. Intrilligator (eds.), *Handbook of Econometrics*, *Volume 2*, North Holland, Amsterdam, 775–826.

Engle, R. F., (ed.) (1995), *ARCH: Selected Readings*, Oxford University Press, Oxford.

Engle, R. F. (2000), The Econometrics of Ultra High Frequency Data, *Econometrica 68*, 1–22.

Engle, R. F. (2002a), Dynamic Conditional Correlation: A Simple Class of Generalized Autoregressive Conditional Heteroskedasticity Models, *Journal of Business and Economic Statistics 20*, 339–350.

Engle, R. F. (2002b), New Frontiers for ARCH Models, *Journal of Applied Econometrics 17*, 425–446.

Engle, R. F. and González-Rivera, G. (1991), Semiparametric ARCH Models, *Journal of Business and Economic Statistics 9*, 345–359.

Engle, R. F. and Granger, C. W. J. (1987), Co-integration and Error Correction: Representation, Estimation and Testing, *Econometrica 55*, 251–276.

Engle, R. F. and Granger, C. W. J., (eds.) (1991), *Long Run Economic Relations: Readings in Cointegration*, Oxford University Press, Oxford.

Engle, R. F., Granger, C. W. J. and Robins, R. (1986), Wholesale and Retail Prices: Bivariate Modeling with Forecastable Variances, in D. Belsley and E. Kuh (eds.), *Model Reliability*, MIT Press, Cambridge, MA.

Engle, R. F., Hendry, D. F. and Richard, J.-F. (1983), Exogeneity, *Econometrica 51*, 277–304.

Engle, R. F., Hendry, D. F. and Trumble, D. (1985), Small-sample Properties of ARCH Estimators and Tests, *Canadian Journal of Economics 18*, 66–93.

Engle, R. F., Hong, C.-H., Kane, A. and Noh, J. (1993), Arbitrage Valuation of Variance Forecasts with Simulated Options, in D. Chance and R. Tripp (eds.), *Advances in Futures and Options Research*, JIA Press, Greenwich, CT.

Engle, R. F., Ito, T. and Lin, W. (1990), Meteor Showers or Heat Waves? Hetroskedastic Intra-daily Volatility in the Foreign Exchange Market, *Econometrica 58*, 525–542.

Engle, R. F., Ito, T. and Lin, W. (1992), Where Does the Meteor Shower Come From? The Role of Stochastic Policy and Coordination, *Journal of International Economics 32*, 221–240.

Engle, R. F., Ito, T. and Lin, W. (1994), Do Bulls and Bears Move Across Borders? International Transmission of Stock Returns and Volatility, *Review of Financial Studies 7*, 507–538.

Engle, R. F., Kane, A. and Noh, J. (1994), Forecasting Volatility and Option Prices of the S&P 500 Index, *Journal of Derivatives 2*, 17–30.

Engle, R. F. and Kozicki, S. (1993), Testing for Common Features, *Journal of Business and Economic Statistics 11*, 369–380.

Engle, R. F. and Kroner, K. F. (1995), Multivariate Simultaneous Generalized ARCH, *Econometric Theory 11*, 122–150.

Engle, R. F. and Lange, J. (2001), Predicting VNET: A Model of the Dynamics of Market Depth, *Journal of Financial Markets 4*, 113–142.

Engle, R. F. and Lee, G. J. (1999), A Permanent and Transitory Component Model of Stock Return Volatility, in R. F. Engle and H. White (eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger*, Oxford University Press, Oxford, 475–497.

Engle, R. F., Lilien, D. M. and Robins, R. P. (1987), Estimating Time-varying Risk Premia in the Term Structure: The ARCH-M Model, *Econometrica 55*, 391–408.

Engle, R. F., Lilien, D. and Watson, M. W. (1985), A DYMIMIC Model of Housing Price Determination, *Journal of Econometrics 28*, 307–326.

Engle, R. F. and Lunde, A. (2003), Trades and Quotes: A Bivariate Point Process, *Journal of Financial Econometrics 1*, 159–188.

Engle, R. F. and Mustafa, C. (1992), Implied ARCH Models from Options Prices, *Journal of Econometrics 52*, 289–311.

Engle, R. F. and Ng, V. K. (1993), Measuring and Testing the Impact of News on Volatility, *Journal of Finance 48*, 1749–1778.

Engle, R. F., Ng, V. K. and Rothschild, M. (1990), Asset Pricing with a Factor ARCH Covariance Structure: Empirical Estimates for Treasury Bills, *Journal of Econometrics 45*, 213–237.

Engle, R. F. and Patton, A. (2004), Impacts of Trades in an Error-correction Model of Quote Prices, *Journal of Financial Markets 7*, 1–25.

Engle, R. F. and Rosenberg, J. V. (1995), GARCH Gammas, *Journal of Derivatives 2*, 47–59.

Engle, R. F. and Rosenberg, J. V. (2000), Testing the Volatility Term Structure Using Option Hedging Criteria, *Journal of Derivatives 8*, 10–28.

Engle, R. F. and Rosenberg, J. V. (2002), Empirical Pricing Kernels, *Journal of Financial Economics 64*, 341–372.

Engle, R. F. and Rothschild, M., (eds.) (1992), ARCH Models in Finance, *Journal of Econometrics 52*, 245–266.

Engle, R. F. and Russell, J. R. (1997), Forecasting the Frequency of Changes in Quoted Foreign Exchange Prices with the Autoregressive Conditional Duration Model, *Journal of Empirical Finance 12*, 187–212.

Engle, R. F. and Russell, J. R. (1998) Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data, *Econometrica 66*, 1127–1162.

Engle, R. F. and Russell, J. R. (2004), Analysis of High Frequency and Transaction Data, forthcoming in L. P. Hansen and Y. Ait-Sahalia (eds.), *Handbook of Financial Econometrics*, North-Holland, Amsterdam.

Engle, R. F. and Smith, A. (1999), Stochastic Permanent Breaks, *Review of Economics and Statistics 81*, 553–574.

Engle, R. F. and Susmel, R. (1993), Common Volatility in International Equity Markets, *Journal of Business and Economic Statistics 11*, 167–176.

Engle, R. F. and Susmel, R. (1994), Hourly Volatility Spillovers Between International Equity Markets, *Journal of International Money and Finance 13*, 3–25.

Engle, R. F. and Vahid, F. (1993), Common Trends and Common Cycles, *Journal of Applied Econometrics 8*, 341–360.

Engle, R. F. and Watson, M. W. (1981), A One-factor Multivariate Time Series Model of Metropolitan Wage Rates, *Journal of the American Statistical Association 76*, 774–781.

Engle, R. F. and Watson, M. W. (1985), The Kalman Filter: Applications to Forecasting and Rational Expectations Models, in T. Bewley (ed.), *Advances in Econometrics: The Fifth World Congress of the Econometric Society*, *Volume 1*, Cambridge University Press, Cambridge, 245–283.

Engle, R. F. and Yoo, B. S. (1987), Forecasting and Testing in Co-integrated Systems, *Journal of Econometrics 35*, 143–159.

Fama, E. F. (1965), The Behavior of Stock Market Prices, *Journal of Business 38*, 34–105.

Ghysels, E., Harvey, A. and Renault, E. (1996), Stochastic Volatility, in G. S. Maddala (ed.), *Handbook of Statistics*. *Volume*. *14*, *Statistical Methods in Finance*, North-Holland, Amsterdam, 119–191.

Gourieroux, C. (1992), *Modeles ARCH et Applications Financieres*, Economica, Paris. (English version, *ARCH Models and Financial Applications,* 1997, Springer-Verlag, New York.)

Granger, C. W. J. (1981), Some Properties of Time Series Data and Their Use in Econometric Model Specification, *Journal of Econometrics 16*, 121–130.

Johansen, S. (1995), *Likelihood-based Inference in Cointegrated Vector Autoregressive Models*, Oxford University Press, Oxford.

Mandelbrot, B. (1963), The Variation of Certain Speculative Prices, *Journal of Business 36*, 394–419.

Shephard, N. (ed.) (2004), *Stochastic Volatility: Selected Readings*, Oxford University Press, Oxford.

Stock, J. H. (1994), Unit Roots, Structural Breaks and Trends, in R. F. Engle and D. McFadden (eds.), *Handbook of Econometrics*, *Volume 4*, North-Holland, Amsterdam, 2740–2843.

Taylor, S. J. (1986), *Modelling Financial Time Series*, John Wiley, Chichester.

Watson, M. W. (1995), Vector Autoregressions and Cointegration, in R. F. Engle and D. McFadden (eds.), *Handbook of Econometrics*, *Volume 4*, North-Holland, Amsterdam, 2843–2915.

Wold, H. O. (1938), *The Analysis of Stationary Time Series*, Almquist & Wicksell, Uppsala.

Zellner, A. (1992), Statistics, Science and Public Policy, *Journal of the American Statistical Association 87*, 1–6.