

Parametric and Nonparametric Volatility Measurement

Torben G. Andersen*, Tim Bollerslev**, and Francis X. Diebold***

*Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL

**Department of Economics, Duke University, Durham, NC

***Department of Economics, University of Pennsylvania, Philadelphia, PA

Contents

1. Introduction	68
2. Volatility Definitions	69
2.1. Continuous-Time No-Arbitrage Price Processes	69
2.2. Notional, Expected, and Instantaneous Volatility	74
2.3. Volatility Modeling and Measurement	82
3. Parametric Methods	84
3.1. Continuous-Time Models	85
3.2. Discrete-Time Models	92
4. Nonparametric Methods	103
4.1. ARCH Filters and Smoothers	103
4.2. Realized Volatility	109
5. Directions for Future Research	124
Acknowledgments	124
References	124

Abstract

Volatility has been one of the most active areas of research in empirical finance and time series econometrics during the past decade. This chapter provides a unified continuous-time, frictionless, no-arbitrage framework for systematically categorizing the various volatility concepts, measurement procedures, and modeling procedures. We define three different volatility concepts: (i) the notional volatility corresponding to the sample-path return variability over a fixed time interval, (ii) the expected volatility over a fixed time interval, and (iii) the instantaneous volatility corresponding to the strength of the volatility process at a point in time. The parametric procedures rely on explicit functional form assumptions regarding the expected and/or instantaneous volatility. In the discrete-time ARCH class of models, the expectations are formulated in terms of directly observable variables, while the discrete- and continuous-time stochastic volatility models involve latent state variable(s). The nonparametric procedures are generally free from such functional form assumptions and hence afford estimates of notional volatility that are flexible yet consistent (as the sampling frequency of the underlying returns increases). The nonparametric procedures include ARCH filters and smoothers designed to measure

the volatility over infinitesimally short horizons, as well as the recently-popularized realized volatility measures for (nontrivial) fixed-length time intervals.

Keywords: realized volatility; stochastic volatility; quadratic return variation; ARCH filters; GARCH

1. INTRODUCTION

Since Engle's (1982) seminal paper on ARCH models, the econometrics literature has focused considerable attention on time-varying volatility and the development of new tools for volatility measurement, modeling, and forecasting.¹ These advances have, in large part, been motivated by the empirical observation that *financial asset return volatility* is time-varying in a persistent fashion, across assets, asset classes, time periods, and countries.² Asset return volatility, moreover, is central to finance, whether in asset pricing, portfolio allocation, or risk management, and standard financial econometric methods and models take on a very different, *conditional*, flavor when volatility is properly recognized to be time-varying.

The combination of powerful methodological advances and important applications within empirical finance produced explosive growth in the financial econometrics of volatility dynamics, with the econometrics and finance literatures cross-fertilizing each other furiously. Initial developments were tightly parametric, but the recent literature has moved in less parametric, and even fully nonparametric directions. Here, we review and provide a unified framework for interpreting both the parametric and nonparametric approaches.

In Section 2, we define three different volatility concepts: (i) the notional volatility corresponding to the ex-post sample-path return variability over a fixed time interval, (ii) the ex-ante expected volatility over a fixed time interval, and (iii) the instantaneous volatility corresponding to the strength of the volatility process at a point in time.

In Section 3, we survey parametric approaches to volatility modeling, which are based on explicit functional form assumptions regarding the expected and/or instantaneous volatility. In the discrete-time ARCH class of models, the expectations are formulated in terms of directly observable variables, while the discrete- and continuous-time stochastic volatility (SV) models both involve latent state variable(s).

In Section 4, we survey nonparametric approaches to volatility modeling, which are generally free from such functional form assumptions and hence afford estimates of notional volatility that are flexible yet consistent (as the sampling frequency of the underlying returns increases). The nonparametric approaches include ARCH filters and

¹We use the terms volatility and variation interchangeably throughout, with the exact meaning specifically defined in the relevant context.

²See, for example, Bollerslev et al. (1992).

smoothers designed to measure the volatility over infinitesimally short horizons, as well as the recently-popularized realized volatility measures for (nontrivial) fixed-length time intervals.

We conclude in Section 5 by highlighting promising directions for future research.

2. VOLATILITY DEFINITIONS

Here, we introduce a unified framework for defining and classifying different notions of return volatility in a continuous-time no-arbitrage setting. We begin by outlining the minimal set of regularity conditions invoked on the price processes and establish the notation used for the decomposition of returns into an expected, or mean, return and an innovation component. The resulting characterization of the price process is central to the development of our different volatility measures, and we rely on the concepts and notation introduced in this section throughout the chapter.

2.1. Continuous-Time No-Arbitrage Price Processes

Measurement of return volatility requires determination of the component of a given price increment that represents a return innovation as opposed to an expected price movement. In a discrete-time setting, this identification may only be achieved through a direct specification of the conditional mean return, for example through an asset pricing model, as economic principles impose few binding constraints on the price process. However, within a frictionless continuous-time framework, the no-arbitrage requirement quite generally guarantees that, instantaneously, the return innovation is an order of magnitude larger than the mean return. This result is not only critical to the characterization of arbitrage-free continuous-time price processes but it also has important implications for the approach one may use for measurement and modeling of volatility over short return horizons.

We take as given a univariate risky logarithmic price process $p(t)$ defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The price process evolves in continuous time over the interval $[0, T]$, where T is a (finite) integer. The associated natural filtration is denoted $(\mathcal{F}_t)_{t \in [0, T]} \subseteq \mathcal{F}$, where the information set, \mathcal{F}_t , contains the full history (up to time t) of the realized values of the asset price and other relevant (possibly latent) state variables, and is otherwise assumed to satisfy the usual conditions. It is sometimes useful to consider the information set generated by the asset price history alone. We refer to this coarser filtration, consisting of the initial conditions and the history of the asset prices only, by $(F_t)_{t \in [0, T]} \subseteq F \equiv F_T$ so that by definition, $F_t \subseteq \mathcal{F}_t$. Finally, we assume there is an asset guaranteeing an instantaneously risk-free rate of interest although we shall not refer to this rate explicitly. Many more risky assets may, of course, be available, but we explicitly retain a univariate focus for notational simplicity. The extension to the multivariate setting is conceptually straightforward as discussed in specific instances below.

The continuously compounded return over the time interval $[t - h, t]$ is then

$$r(t, h) = p(t) - p(t - h), \quad 0 \leq h \leq t \leq T. \quad (2.1)$$

We also adopt the following short-hand notation for the cumulative return up to time t , i.e., the return over the $[0, t]$ time interval:

$$r(t) \equiv r(t, t) = p(t) - p(0), \quad 0 \leq t \leq T. \quad (2.2)$$

These definitions imply a simple relation between the period-by-period and the cumulative returns that we use repeatedly in the sequel:

$$r(t, h) = r(t) - r(t - h), \quad 0 \leq h \leq t \leq T. \quad (2.3)$$

A maintained assumption throughout is that – almost surely (P) (henceforth denoted (*a.s.*)) – the asset price process remains strictly positive and finite so that $p(t)$ and $r(t)$ are well defined over $[0, T]$ (*a.s.*). It follows that $r(t)$ has only countably (although possibly infinitely) many jump points over $[0, T]$, and we adopt the convention of equating functions that have identical left and right limits everywhere. We also assume that the price and return processes are squared integrable.

Defining $r(t-) \equiv \lim_{\tau \rightarrow t, \tau < t} r(\tau)$ and $r(t+) \equiv \lim_{\tau \rightarrow t, \tau > t} r(\tau)$ uniquely determines the right-continuous, left-limit (*càdlàg*) version of the process, for which $r(t) = r(t+)$ (*a.s.*), and the left-continuous, right-limit (*càglàd*) version, for which $r(t) = r(t-)$ (*a.s.*), for all t in $[0, T]$. In the following, we assume without loss of generality that we are working with the *càdlàg* version of the various return processes.

The jumps in the cumulative price and return process are then

$$\Delta r(t) \equiv r(t) - r(t-), \quad 0 \leq t \leq T. \quad (2.4)$$

Obviously, at continuity points for $r(t)$, we have $\Delta r(t) = 0$. Moreover, given the at most countably infinite number of jumps, we generically have

$$P(\Delta r(t) \neq 0) = 0, \quad (2.5)$$

for an arbitrarily chosen t in $[0, T]$. This does not imply that jumps necessarily are rare, since as already noted, Eq. (2.5) is consistent with there being a (countably) infinite number of jumps over any discrete interval – a phenomenon referred to as an explosion. Jump processes that do not explode are termed *regular*. For regular processes, the anticipated jump frequency is conveniently characterized by the instantaneous jump intensity, i.e., the probability of a jump over the next instant of time, and expressed in units that reflect the expected (and finite) number of jumps per unit time interval.

Henceforth, we invoke the standard assumptions of no arbitrage and finite-expected returns. Within our frictionless setting, these conditions imply that the log-price process must constitute a (special) semimartingale (see Back, 1991; Harrison and Kreps, 1978). This, in turn, affords the following unique canonical return decomposition (e.g., Protter, 1992).

Proposition 1 *Return Decomposition*

Any arbitrage-free logarithmic price process subject to the regularity conditions outlined above may be uniquely represented as

$$r(t) \equiv p(t) - p(0) = \mu(t) + M(t) = \mu(t) + M^c(t) + M^J(t), \quad (2.6)$$

where $\mu(t)$ is a predictable and finite-variation process, $M(t)$ is a local martingale that may be further decomposed into $M^c(t)$, a continuous sample path, infinite-variation local martingale component, and $M^J(t)$, a compensated jump martingale. We may normalize the initial conditions such that all components may be assumed to have initial conditions normalized such that $\mu(0) \equiv M(0) \equiv M^c(0) \equiv M^J(0) \equiv 0$, which implies that $r(t) \equiv p(t)$.

Proposition 1 provides a unique decomposition of the instantaneous return into an expected return component and an (martingale) innovation. Over discrete intervals, the relation becomes slightly more complex. Letting the expected returns over $[t - h, t]$ be denoted by $m(t, h)$, Eq. (2.6) implies

$$m(t, h) \equiv E[r(t, h) | \mathcal{F}_{t-h}] = E[\mu(t, h) | \mathcal{F}_{t-h}], \quad 0 < h \leq t \leq T, \quad (2.7)$$

where

$$\mu(t, h) \equiv \mu(t) - \mu(t - h), \quad 0 < h \leq t \leq T, \quad (2.8)$$

and the return innovation takes the form

$$r(t, h) - m(t, h) = (\mu(t, h) - m(t, h)) + M(t, h), \quad 0 < h \leq t \leq T. \quad (2.9)$$

The first term on the right-hand side of (2.9) signifies that the expected return process, even though it is (locally) predictable, may evolve stochastically over the $[t - h, t]$ interval.³ If $\mu(t, h)$ is predetermined (measurable with respect to \mathcal{F}_{t-h}), and thus known at time $t - h$, then the discrete-time return innovation reduces to $M(t, h) \equiv$

³In other words, even though the conditional mean is *locally* predictable, all return components in the special semimartingale decomposition are generally stochastic: not only volatility but also the jump intensity, the jump size distribution and the conditional mean process may evolve randomly over a finite interval.

$M(t) - M(t - h)$. However, any shift in the expected return process during the interval will generally render the initial term on the right-hand side of (2.9) nonzero and thus contribute to the return innovation over $[t - h, t]$.

Although the discrete-time return innovation incorporates two distinct terms, the martingale component, $M(t, h)$, is generally the dominant contributor to the return variation over short intervals, i.e., for h small. To discuss the intuition behind this result, which we formalize in the following section, it is convenient to decompose the expected return process into a purely continuous, predictable finite-variation part, $\mu^c(t)$, and a purely predictable jump part, $\mu^J(t)$. Because the continuous component, $\mu^c(t)$, is of finite variation, it is locally an order of magnitude smaller than the corresponding contribution from the continuous component of the innovation term, $M^c(t)$. The reason is – loosely speaking – that an asset earning, say a positive expected return over the risk-free rate must have innovations that are an order of magnitude larger than the expected return over infinitesimal intervals. Otherwise, a sustained long position (infinitely, many periods over any interval) in the risky asset will tend to be perfectly diversified due to a law of large numbers, as the martingale part is uncorrelated. Thus, the risk-return relation becomes unbalanced. Only if the innovations are large, preventing the law of large numbers from becoming operative, will this not constitute a violation of the no-arbitrage condition (see Maheswaran and Sims, 1993, for further discussion related to the implications of the classical Harrison and Kreps, 1978, equivalent martingale assumption). The presence of a nontrivial $M^J(t)$ component may similarly serve to eliminate arbitrage and retain a balanced risk-return trade-off relationship.

Analogous considerations apply to the jump component for the expected return process, $\mu^J(t)$, if this factor is present. There cannot be a predictable jump in the mean – i.e., a perfectly anticipated jump in terms of both time and size – unless it is accompanied by large jump innovation risk as well so that $Pr(\Delta M(t) \neq 0) > 0$. Again, intuitively, if there was a known, say, positive jump, then this induces arbitrage (by going long the asset) unless there is offsetting (jump) innovation risk.⁴ Most of the continuous-time asset pricing literature ignores predictable jumps, even if they are logically consistent with the framework. One reason may be that their existence is fragile in the following sense. A fully anticipated jump must be associated with release of new (price relevant) information at a given point in time. However, if there is any uncertainty about the timing of the announcement so that it is only known to occur within a given minute, or even a few seconds, then the timing of the jump is more aptly modeled by a continuous hazard

⁴This point is perhaps most readily understood by analogy to a discrete-time setting. When there is a predictable jump at time t , the instant from $t-$ to t is effectively equivalent to a trading period, say from $t - 1$ to t , within a discrete-time model. In that setting, no asset can earn a positive (or negative) excess return relative to the risk-free rate over $(t - 1, t]$ without bearing genuine risk as this would otherwise imply a trivial arbitrage opportunity. The argument ruling out a predictable price jump without an associated positive probability of a jump innovation is entirely analogous.

function where the jump probability at each point in time is zero, and the predictable jump event is thus eliminated. In addition, even if there were predictable jumps associated with scheduled news releases, the size of the predictable component of the jump is likely much smaller than the size of the associated jump innovation so that the descriptive power lost by ignoring the possibility of predictable jumps is minimal. Thus, rather than modifying the standard setup to allow for the presence of predictable (but empirically negligible) jumps, we follow the tradition in the literature and assume away such jumps.

Although we will not discuss specific model classes at length until later sections, it is useful to briefly consider a simple example to illustrate the somewhat abstract definitions given in the current section.

Example 1 *Stochastic Volatility Jump Diffusion with Nonzero Mean Jumps*

Consider the following continuous-time jump diffusion expressed in stochastic differential equation (SDE) form,

$$dp(t) = (\mu + \beta\sigma^2(t))dt + \sigma(t)dW(t) + \kappa(t)dq(t), \quad 0 \leq t \leq T,$$

where $\sigma(t)$ is a strictly positive continuous sample path process (a.s.), $W(t)$ denotes a standard Brownian motion, $q(t)$ is a counting process with $dq(t) = 1$ corresponding to a jump at time t , and $dq(t) = 0$ otherwise, while the $\kappa(t)$ process gives the sizes of the jumps and is only defined for jump times t for which $dq(t) = 1$. We assume that the jump size distribution has a constant mean of μ_κ and variance of σ_κ^2 . Finally, the jump intensity is assumed constant (and finite) at a rate λ per unit time. In the notation of Proposition 1, we then have the return components,

$$\mu(t) = \mu^c(t) = \mu \cdot t + \beta \int_0^t \sigma^2(s)ds + \lambda \cdot \mu_\kappa \cdot t,$$

$$M^c(t) = \int_0^t \sigma(s)dW(s),$$

$$M^J(t) = \sum_{0 \leq s \leq t} \kappa(s)dq(s) - \lambda \cdot \mu_\kappa \cdot t,$$

where by definition, the last summation consists of all the jumps that occurred over the $[0, T]$ time interval. Notice that the last term of the mean representation captures the expected contribution coming from the jumps, while the corresponding term is subtracted from the jump innovation process to provide a unique (compensated) jump martingale representation for M^J .

A final comment is in order. We purposely express the price changes and associated returns in Proposition 1 over a discrete time interval. The concept of an instantaneous

return used in the formulation of continuous-time models, as in Example 1 given above, is mere short-hand notation that is formally defined only through the corresponding integral representation, such as Eq. (2.6). Although this is a technical point, it has an important empirical analogy: real-time price data are not available at every instant, and due to pertinent microstructure features, prices are invariably constrained to lie on a discrete grid, both in the price and time dimension. Hence, there is no real-world counterpart to the notion of a continuous sample path martingale with infinite variation over arbitrarily small time intervals (say, less than a second). It is only feasible to measure return (and volatility) realizations over discrete time intervals. Moreover, sensible measures can typically only be constructed over much longer horizons than given by the minimal interval length for which consecutive trade prices or quotes are recorded. We return to this point later. For now, we simply note that our main conceptualization of volatility in the next section conforms directly with the focus on realizations measured over nontrivial discrete time intervals rather than vanishing, or instantaneous interval lengths.

2.2. Notional, Expected, and Instantaneous Volatility

This section introduces three different volatility concepts, each of which serves to formalize the process of measuring and modeling the strength of the return variation within our frictionless arbitrage-free setting. Two distinct features importantly differentiate the construction of the different measures. First, given a set of actual return observations, how is the *realized* volatility computed? Here, the emphasis is explicitly on *ex-post* measurement of the volatility. Second, decision making often requires forecasts of future return volatility. The focus is then on *ex-ante* expected volatility. The latter concept naturally calls for a model that may be used to map the current information set into a volatility *forecast*. In contrast, the (ex-post) realized volatility may be computed (or approximated) without reference to any specific model, thus rendering the task of volatility measurement essentially a nonparametric procedure.

It is natural first to concentrate on the behavior of the martingale component in the return decomposition (2.6). However, a prerequisite for observing the $M(t)$ process is that we have access to a continuous record of price data. Such data are simply not available, and even for extremely liquid markets, microstructure effects (discrete price grids, bid-ask bounce effects, etc.) prevent us from ever getting really close to a true continuous sample-path realization. Consequently, we focus on measures that represent the (average) volatility over a discrete time interval rather than the instantaneous (point-in-time) volatility.⁵ This, in turn, suggests a natural and general notion of volatility based

⁵Of course, by choosing the interval very small, one may, in principle, approximate the notion of point-in-time volatility, as discussed further below.

on the quadratic variation process for the local martingale component in the unique semimartingale return decomposition.

Specifically, let $X(t)$ denote any (special) semimartingale. The unique *quadratic variation process*, $[X, X]_t$, $t \in [0, T]$, associated with $X(t)$ is formally defined as

$$[X, X]_t \equiv X(t)^2 - 2 \int_0^t X(s-) dX(s), \quad 0 < t \leq T, \quad (2.10)$$

where the stochastic integral of the adapted càglàd process, $X(s-)$, with respect to the càdlàg semimartingale, $X(s)$, is well-defined (e.g., Protter, 1992). It follows directly that the quadratic variation, $[X, X]$, is an increasing stochastic process. Also, jumps in the sample path of the quadratic variation process necessarily occur concurrent with the jumps in the underlying semimartingale process, $\Delta[X, X] = (\Delta X)^2$.

Importantly, if M is a locally square integrable martingale, then the associated $(M^2 - [M, M])$ process is a local martingale,

$$E[M(t, h)^2 - ([M, M]_t - [M, M]_{t-h}) | \mathcal{F}_{t-h}] = 0, \quad 0 < h \leq t \leq T. \quad (2.11)$$

This relation, along with the following well-known result, provides the key to the interpretation of the quadratic variation process as one of our volatility measures.

Proposition 2 *Theory of Quadratic Variation*⁶

Let a sequence of partitions of $[0, T]$, (τ_m) , be given by $0 = \tau_{m,0} \leq \tau_{m,1} \leq \tau_{m,2} \leq \dots \leq \tau_{m,m} = T$ such that $\sup_{j \geq 0} (\tau_{m,j+1} - \tau_{m,j}) \rightarrow 0$ for $m \rightarrow \infty$. Then, for $t \in [0, T]$,

$$\lim_{m \rightarrow \infty} \left\{ \sum_{j \geq 1} (X(t \wedge \tau_{m,j}) - X(t \wedge \tau_{m,j-1}))^2 \right\} \rightarrow [X, X]_t,$$

where $t \wedge \tau \equiv \min(t, \tau)$, and the convergence is uniform in probability.

Intuitively, the proposition says that the quadratic variation process represents the (cumulative) realized sample-path variability of $X(t)$ over the $[0, t]$ time interval. This observation, together with the martingale property of the quadratic variation process in (2.11), immediately points to the following theoretical notion of return variability.

Definition 1 *Notional Volatility*

The Notional Volatility over $[t - h, t]$, $0 < h \leq t \leq T$, is

$$v^2(t, h) \equiv [M, M]_t - [M, M]_{t-h} = [M^c, M^c]_t - [M^c, M^c]_{t-h} + \sum_{t-h < s \leq t} \Delta M^2(s). \quad (2.12)$$

⁶ The theory of quadratic variation generalizes to situations in which the “stopping times” or partitions of $[0, T]$ are random and independent of $X(t)$, and satisfy, with probability one for $m \rightarrow \infty$, $\tau_{m,0} \rightarrow 0$, $\sup_{j \geq 1} \tau_{m,j} \rightarrow T$, and $\sup_{j \geq 0} (\tau_{m,j+1} - \tau_{m,j}) \rightarrow 0$; see, e.g., Chung and Williams (1983).

This same volatility was first highlighted in a series of papers by Andersen et al. (2001b, 2003a) and Barndorff-Nielsen and Shephard (2002a,b). The latter authors term the corresponding concept actual volatility.

Under the maintained assumption of no predictable jumps in the return process and noting that the quadratic variation of any finite-variation process, such as $\mu^c(t)$, is zero, we also have

$$v^2(t, h) \equiv [r, r]_t - [r, r]_{t-h} = [M^c, M^c]_t - [M^c, M^c]_{t-h} + \Sigma_{t-h < s \leq t} \Delta r^2(s). \quad (2.13)$$

Consequently, the notional volatility equals (the increment to) the quadratic variation for the return series. Equation (2.13) and Proposition 2 also suggest that (ex-post) it is possible to approximate the notional volatility arbitrarily well through the accumulation of ever finely sampled high-frequency squared return, and that this approach remains consistent independent of the expected return process. We shall return to a much more detailed analysis of this idea in our discussion of nonparametric ex-post volatility measures in Section 4.

Similarly, from (2.13) and Proposition 2, it is evident that the notional volatility, $v^2(t, h)$, directly captures the sample path variability of the log-price process over the $[t - h, t]$ time interval. In particular, the notional volatility explicitly incorporates the effect of (realized) jumps in the price process: jumps contribute to the realized return variability and forecasts of volatility must account for the potential occurrence of such jumps. It also follows, from the properties of the quadratic variation process, that

$$E[v^2(t, h) | \mathcal{F}_{t-h}] = E[M(t, h)^2 | \mathcal{F}_{t-h}] = E[M^2(t) | \mathcal{F}_{t-h}] - M^2(t - h), \quad 0 < h \leq t \leq T. \quad (2.14)$$

Hence, the *expected notional volatility* represents the expected future (cumulative) squared return innovation. As argued in Section 2.1, this component is typically the dominant determinant of the expected return volatility.

For illustration, consider again the example introduced in Section 2.1. More complicated specifications and issues related to longer horizon returns are considered in Section 3.

Example 2 *Stochastic Volatility Jump Diffusion with Nonzero Mean Jumps (Revisited)*

The log-price process evolves according to

$$dp(t) = (\mu + \beta\sigma^2(t))dt + \sigma(t)dW(t) + \kappa(t)dq(t), \quad 0 \leq t \leq T.$$

The notional volatility is then

$$v^2(t, h) = \int_{t-h}^t \sigma^2(s)ds + \Sigma_{t-h < s \leq t} \kappa^2(s),$$

where again the last sum is to be interpreted as consisting of all the nonzero squared jumps that occurred over the $[t - h, t]$ time interval. The expected notional volatility involves taking the conditional expectation of this expression. Without an explicit model for the volatility process, this cannot be given in closed form. However, for small h , the (expected) notional volatility is typically very close to the value attained if volatility is constant. In particular, to a first-order approximation,

$$E[v^2(t, h) | \mathcal{F}_{t-h}] \approx \sigma^2(t - h) \cdot h + \lambda \cdot h \cdot (\mu_\kappa^2 + \sigma_\kappa^2) = [\sigma^2(t - h) + \lambda(\mu_\kappa^2 + \sigma_\kappa^2)] \cdot h,$$

while

$$m(t, h) \approx [\mu + \beta \cdot \sigma^2(t - h) + \lambda \cdot \mu_\kappa] \cdot h.$$

Thus, the expected notional volatility is of order h , the expected return is of order h (and the variation of the mean return of order h^2), whereas the martingale (return innovation) is of the order $h^{1/2}$, and hence an order of magnitude larger for small h .

It is obvious that, when volatility is stochastic, the ex-post (realized) notional volatility will not correspond to the ex-ante expected volatility. More importantly, Eq. (2.14) implies that even the ex-ante expected notional volatility generally is *not* identical to the usual notion of return volatility as an *ex-ante* characterization of future return variability over a discrete holding period. The fact that the latter quantity is highly relevant for financial decision making motivates the standard discrete-time expected volatility concept as defined below.

Definition 2 *Expected Volatility*

The expected volatility over $[t - h, t]$, $0 < h \leq t \leq T$, is defined by

$$\begin{aligned} \zeta^2(t, h) &\equiv E[\{r(t, h) - E(\mu(t, h) | \mathcal{F}_{t-h})\}^2 | \mathcal{F}_{t-h}] \\ &= E[\{r(t, h) - m(t, h)\}^2 | \mathcal{F}_{t-h}]. \end{aligned} \tag{2.15}$$

If the $\mu(t, h)$ process is not measurable with respect to \mathcal{F}_{t-h} , the expected volatility will typically differ from the expected notional volatility in Eq. (2.14).⁷ Specifically, the future return variability in Eq. (2.15) reflects both genuine return innovations, as in Eq. (2.14), and intraperiod innovations to the conditional mean process. Trivially, of course, for models with an assumed constant mean return, or for one-period-ahead discrete-time volatility forecasts with given conditional mean representation, the two concepts coincide.

⁷Note that (t, h) refers to the $[t - h, t]$ time interval. So that while the notional volatility, $v^2(t, h)$, is only measurable with respect to \mathcal{F}_t , the expected volatility, $\zeta^2(t, h)$, is by definition \mathcal{F}_{t-h} -measurable.

Of course, for a continuous-time model, any volatility forecast over a discrete time interval invariably entails multiperiod considerations (typically a continuum). In particular, following Andersen et al. (2001b), the expected volatility may generally be expressed as

$$\begin{aligned}\zeta^2(t, h) &= E[(r(t, h) - m(t, h))^2 | \mathcal{F}_{t-h}] \\ &= E[v^2(t, h) | \mathcal{F}_{t-h}] + \text{Var}[\mu(t, h) | \mathcal{F}_{t-h}] + 2 \cdot \text{Cov}[M(t, h), \mu(t, h) | \mathcal{F}_{t-h}].\end{aligned}\quad (2.16)$$

The expected volatility therefore involves the expected notional volatility (quadratic variation) as well as two terms induced by future within-forecast-period variation in the conditional mean. The random variation in the mean component is a direct source of future return variation, and any covariation between the return and conditional mean innovations will further impact the return variability. However, under standard conditions and moderate forecast horizons, the dominant factor is indisputably the expected notional volatility, as the innovations to the mean return process generally will be very small relative to the cumulative return innovations. Importantly, this does not rule out asymmetric effects from current return innovations to future return volatility, as in the so-called leverage and volatility feedback effects discussed further below, which work exclusively or primarily through the notional volatility process.

Continuous-time models often portray the volatility process as perpetually evolving. From this perspective, the focus on volatility measurement over a fixed interval length, h , is ultimately arbitrary. A more natural theoretical concept is provided by the expected instantaneous volatility, measured as the current strength of the volatility process expressed per unit of time,

$$\lim_{h \rightarrow 0} \zeta^2(t + h, h)/h = \lim_{h \rightarrow 0} [E\{([M, M]_{t+h} - [M, M]_t)/h\} | \mathcal{F}_t]. \quad (2.17)$$

This is especially true when the underlying logarithmic price path is continuous, i.e., $M^J(t) \equiv 0$, in which case the (scaled) notional and expected instantaneous volatilities coincide,

$$\begin{aligned}\lim_{h \rightarrow 0} \zeta^2(t + h, h)/h &= \lim_{h \rightarrow 0} [E\{([M^c, M^c]_{t+h} - [M^c, M^c]_t)/h\} | \mathcal{F}_t] \\ &= \lim_{h \rightarrow 0} \{([M^c, M^c]_{t+h} - [M^c, M^c]_t)/h\} = \lim_{h \rightarrow 0} v^2(t + h, h)/h.\end{aligned}\quad (2.18)$$

Inspired by these relations, we adopt the following definition of instantaneous volatility.⁸

⁸The definition adapted here implies that σ_t^2 is a càdlàg process. An alternative càglàd definition is sometimes used in the literature.

Definition 3 *Instantaneous Volatility*

The instantaneous volatility at time t , $0 \leq t \leq T$, is

$$\sigma_t^2 \equiv \lim_{h \rightarrow 0} [E\{([M^c, M^c]_{t+h} - [M^c, M^c]_t)/h \mid \mathcal{F}_t\}]. \quad (2.19)$$

This definition is consistent with the terminology commonly used in the literature on continuous-time parametric SV models.⁹ Barndorff-Nielsen and Shephard (2002a,b), in a slightly different setting, refer to the corresponding concept as Spot Volatility.

The continuous-time models in the theoretical asset and derivatives pricing literature frequently assume that the sample paths are continuous, with the corresponding diffusion processes given in the form of SDEs (as in the Example 2 given above), rather than through (abstract) integral representations for continuous sample path semimartingales along the lines of Proposition 1. This does not involve any loss of generality, as illustrated by the following well-known result (e.g., Karatzas and Shreve, 1991; Protter, 1992).

Proposition 3 *Martingale Representation Theorem*

For any univariate, square-integrable, continuous sample path, logarithmic price process, which is not locally riskless, there exists a representation such that for all $0 \leq t \leq T$, a.s.(P),

$$r(t, h) = \mu(t, h) + M(t, h) = \int_{t-h}^t \mu(s) ds + \int_{t-h}^t \sigma(s) dW(s), \quad (2.20)$$

where $\mu(s)$ is an integrable, predictable, and finite-variation stochastic process, $\sigma(s)$ is a strictly positive càdlàg stochastic process satisfying

$$P \left[\int_{t-h}^t \sigma^2(s) ds < \infty \right] = 1,$$

and $W(s)$ is a standard Brownian motion.

The integral representation in (2.20) is equivalent to the standard (short-hand) sde specification for the logarithmic price process,

$$dp(t) = \mu(t) dt + \sigma(t) dW(t), \quad 0 \leq t \leq T. \quad (2.21)$$

Hence, within the class of continuous sample path semimartingale (diffusion) models, there are no consequential restrictions involved in stating the model directly through a SDE. In accordance with Definition 3, the volatility coefficient process in this formulation, $\{\sigma_t^2\}_{t \in [0, T]}$, is usually termed the *instantaneous volatility*, and we have the following

⁹A definition of instantaneous volatility similar to Eq. (2.17) is suggested by Comte and Renault (1998) in their discussion of alternative inference procedures for a continuous-time long-memory volatility model.

direct link between these alternative volatility representations,

$$\sigma_t^2 = \lim_{h \rightarrow 0} \sigma^2(t+h) = \lim_{h \rightarrow 0} \left(\int_t^{t+h} \sigma^2(s) ds / h \right). \quad (2.22)$$

It is also immediately evident that in this situation, the notional volatility (or the increment to the quadratic variation process) equals the so-called *Integrated Volatility*,

$$v^2(t, h) = [M, M]_t - [M, M]_{t-h} = \int_{t-h}^t \sigma^2(s) ds. \quad (2.23)$$

The integrated volatility plays a key role in the SV option pricing literature. Hull and White (1987) document that option prices follow Black–Scholes with the simple modification that the constant volatility is replaced by the expected quadratic return variation over the time to expiry for pure diffusive models without asymmetries between return and volatility innovations.¹⁰ For further discussion of derivatives pricing models and related empirical procedures, see, e.g., Bates (1996b), Garcia et al. (2001), and Garcia et al. (2010).

To further appreciate the different volatility concepts, it is instructive to consider an illustrative example.

Illustration 1 *Continuous-Time GARCH Model*

The three panels in Fig. 2.1 show the time series of artificially simulated logarithmic prices, $p(t)$, one-period returns, $r(t, 1)$, and corresponding instantaneous volatilities, $\sigma(t)$, for $t = 1, 2, \dots, 2,500$.¹¹ Comparing the middle and bottom panel, it is evident that the instantaneous volatility from the model directly dictates the strength of the observed return variation. However, even though the instantaneous volatility is a natural theoretical concept, and we refer to it frequently below, practical volatility measurement invariably takes place over discrete time intervals. The notational h -period volatility was introduced with exactly this consideration in mind. Of course, the difference between the notional and instantaneous volatility will depend upon the persistency of the underlying process and the value of h . In particular, the two volatility concepts formally coincide in the limit

¹⁰More generally, the concept of model-free implied volatility, constructed from the cross-section of out-of-the-money options over different strikes but at the identical maturity, see, e.g., Britten–Jones and Neuberger (2000) and Carr and Madan (1998), allows for extraction of a nonparametric measure of the expected notional volatility under the risk-neutral (pricing or martingale) measure. We discuss these developments briefly later on.

¹¹The data are generated by the continuous-time GARCH model defined in Eq. (4.4) below, $dp(t) = \sigma(t)dW(t)$ and $d\sigma^2(t) = (\omega - \theta\sigma^2(t))dt + (2\alpha)^{1/2}\sigma^2(t)dV(t)$, where $\theta = 0.01005$, $\omega = 0.01005$, and $\alpha = 0.01095$, corresponding to a one-period discrete-time weak-form GARCH(1,1) model with $\alpha_1 = 0.09$, $\beta_1 = 0.9$, and unconditional variance equal to unity (see Andersen and Bollerslev, 1998a, for further details).




Figure 2.1 The first three panels in the figure plot simulated logarithmic prices, $p(t)$, one-period returns, $r(t, 1)$, and instantaneous volatilities, $\sigma_t = \sigma(t)$, for $t = 1, 2, \dots, 2500$. The fourth and fifth panel depict the corresponding scaled h -period notional volatilities, $v(t, h)\sqrt{h}$, and scaled expected volatilities, $\zeta(t, h)\sqrt{h}$, for $h = 22$. The prices and volatilities are generated by a continuous-time GARCH model defined by $dp(t) = \sigma(t)dW(t)$ and $d\sigma^2(t) = (\omega - \theta\sigma^2(t))dt + (2\alpha)^{1/2}\sigma^2(t)dV(t)$. The parameters in the continuous-time model are calibrated to match a one-period weak-form GARCH(1,1) model with $\alpha_1 = 0.09$, $\beta_1 = 0.9$ and unconditional variance equal to unity; see the discussion in Sections 3.1.1 and 4.1, along with Andersen and Bollerslev (1998a), for further details.

for $h \rightarrow 0$, but for large values of h , the two measures are clearly different, and in the limit for $h \rightarrow \infty$, the per period notional volatility, $v^2(t, h)/h$, will generally be constant (provided that it exist). To illustrate, the fourth panel in the figure plots the scaled notional volatility, $v(t, h)\sqrt{h}$, from the same model corresponding to a “month,” or $h = 22$. This series is obviously much smoother than the instantaneous volatility. Finally, the fifth panel in the figure shows the “monthly” expected volatility,¹² where for comparison purposes with the other plots, we have scaled by the forecast horizon, i.e., $\zeta(t, 22)/\sqrt{22}$. Because $\mu(t, h)$ is constant (and equal to zero), the expected volatility equals the expected notional volatility, which explains the apparent similarities in the two shapes. Also, since the underlying volatility process is quite persistent, the (scaled) expected volatility appears fairly similar to the previously depicted instantaneous volatility, even for $h = 22$. Nonetheless, the three volatility measures shown in the figure obviously differ and speak importantly to different aspects of the underlying data-generating process.

In the next section, we further stress the general relationship between the various volatility concepts for alternative parametric volatility models and nonparametric volatility measurements. Specific characterization of the volatility estimates and measurements are postponed to the following sections, where we present a more detailed study for each major class of models.

2.3. Volatility Modeling and Measurement

The approaches for empirically quantifying volatility naturally falls into two separate categories, namely procedures based on estimation of parametric models and more direct nonparametric measurements. Within the parametric volatility classification, alternative models exploit different assumptions regarding the expected volatility, $\zeta^2(t, h)$, through distinct functional forms and the nature of the variables in the information set, \mathcal{F}_{t-h} . In contrast, the data-driven or nonparametric volatility measurements typically quantify the notional volatility, $v^2(t, h)$, directly. Both set of procedures differ importantly in terms of the choice of time interval for which the volatility measure applies, e.g., a discrete interval, $h > 0$, or a point-in-time (instantaneous) measure, obtained as the limiting case for $h \rightarrow 0$.

Within the discrete-time parametric models, the most significant distinction concerns the character of the variables in the information set, \mathcal{F}_{t-h} , which in turn governs the type of estimation and inference techniques that are required for their practical implementation. In the ARCH class of models, the expected volatility, $\zeta^2(t, h)$, is parameterized as a function of past returns only, or F_{t-h} , although other observable variables could easily be included in \mathcal{F}_{t-h} . In contrast, the parameterized expectations in the SV class of

¹²The expected volatility in the continuous-time GARCH model is formally given by $\zeta^2(t, h) = h(\omega/\theta) + \theta^{-1}[\sigma^2(t) - (\omega/\theta)] [1 - \exp(-h\theta)]$.

models explicitly rely on latent state variables. As we move to continuous-time parametric representations of either model, the assumption that all past returns are observable implies that the distinction between the two classes of models effectively vanishes, as the latent volatility state variables may be extracted without error from the frictionless, continuous-time price record. The following definition formalizes these categorizations.

Definition 4 *Parametric Volatility Models*

Discrete-time parametric volatility models *explicitly parameterize the expected volatility, $\zeta^2(t, h)$, $h > 0$, as a nontrivial function of the time $t - h$ information set, \mathcal{F}_{t-h} . In the ARCH class of models, \mathcal{F}_{t-h} depends on past returns and other directly observable variables only. In the SV class of models, \mathcal{F}_{t-h} explicitly incorporates past returns as well as latent state variables. Continuous-time volatility models provide an explicit parameterization of the instantaneous volatility, σ_t^2 , as a (nontrivial) function of the \mathcal{F}_t information set, with additional volatility dynamics possibly introduced through time variation in the process governing jumps in the price path.*

In addition to these three separate model classes, so-called implied volatility approaches also figure prominently in the literature. The implied volatilities are typically based on a parametric model for the returns, as defined above, along with an *asset pricing model* and an augmented information set consisting of options prices and/or term structure variables. Intuitively, if the number of available derivatives prices at time $t - h$ pertaining to the price of the asset at time t included in the augmented information set, \mathcal{F}_{t-h} , exceeds the number of latent state variables in the parametric model for the returns, it is possible to back out a value for $\zeta^2(t, h)$ by inverting the theoretical asset pricing model; see, e.g., Bates (1996b), Renault (1997), and Chapters 9 and 12 in this volume for a discussion of the extensive literature on options implied volatilities and related procedures.^{13,14}

In contrast to the parametric procedures categorized above, the nonparametric volatility measurements are generally void of any specific functional form assumptions about the stochastic process(es) governing the local martingale, $M(t)$, as well as the predictable

¹³Most prominent among these procedures are, of course, the Black–Scholes option implied volatilities based on the assumption of an underlying continuous-time random walk model first analyzed empirically by Latané and Rendleman (1976). More detailed empirical analyses of Black–Scholes-implied volatilities along with generalizations to allow for more realistic price dynamics have been the subject of an enormous literature, an incomplete list of which includes Bakshi et al. (1997), Canina and Figlewski (1993), Chernov and Ghysels (2000), Christensen and Prabhala (1998), Day and Lewis (1992), Duan (1995), Dumas et al. (1998), Fleming (1998), Heston (1993), Heston and Nandi (2000), Hull and White (1987), and Wiggins (1992). The so-called model-free implied volatilities computed from option prices without the use of a particular pricing model have recently been proposed by Britten-Jones and Neuberger (2002), Carr and Madan (1998), and Demeterfi et al. (1999), and analyzed empirically by Bollerslev and Zhou (2006), Bollerslev et al. (2005), Carr and Wu (2006, 2009), Garcia et al. (2001), and Jiang and Tian (2005).

¹⁴The 30-day VIX-implied volatility index of the Chicago Board Options Exchange (CBOE), for which there is an active futures market, are based on S&P500 index options along with the model-free-implied volatility formula of Britten-Jones and Neuberger (2002); see Carr and Wu (2009) for further discussion. A similar construct underlies the VXN index for the NASDAQ-100 and the new VDAX for the DAX index on the Deutsche Termin Börse (DTB). Earlier versions of these indexes were based on weighted averages of Black–Scholes-implied volatilities; see Fleming et al. (1995) and Whaley (1993) for further discussion of these historically first volatility indexes.

and finite-variation process, $\mu(t)$, in the unique return decomposition. These procedures also differ importantly from the parametric models in their focus on providing measures of the notional volatility, $v^2(t, h)$, rather than the expected volatility, $\zeta^2(t, h)$. In addition, the nonparametric procedures generally restrict the measurements to be functions of the coarser filtration, F_t , generated by the return on the asset only. In parallel to the parametric measures, the nonparametric procedures may be further differentiated depending upon whether they let $h \rightarrow 0$ and thus provide measures of instantaneous volatility, or whether they explicitly operate with a strictly positive $h > 0$ resulting in realized volatility measures over a discrete nontrivial time interval.

Definition 5 *Nonparametric Volatility Measurement*

Nonparametric volatility measurement utilizes the ex-post returns, or F_τ , in extracting measures of the notional volatility. ARCH filters and smoothers are designed to measure the instantaneous volatility, σ_t^2 . The filters only use information up to time $\tau = t$, while the smoothers are based on $\tau > t$. Realized volatility measures directly quantify the notional volatility, $v^2(t, h)$, over (nontrivial) fixed-length time intervals, $h > 0$.

Within the class of instantaneous volatility measures, the ARCH filters first formally developed by Nelson (1992) (see also the collection of papers in Rossi, 1996) rely exclusively on the past return record, typically through a weighted rolling regression, while the smoothers, or two-sided filters, from Nelson (1996b) exploit (ex-post) future prices. Realized volatility approaches may similarly be categorized according to whether the measurement of $v^2(t, h)$ exploits only price observations within the interval $[t - h, t]$ itself or filtering/smoothing techniques are used to also incorporate return observations outside of $[t - h, t]$. An important advantage of exploiting only interval-specific information is it produces asymptotically unbiased measures, and therefore approximately serially uncorrelated measurement errors, under quite general conditions. A potential drawback is that useful information from adjacent intervals is ignored. Consistency of both ARCH filters and smoothers and realized volatility procedures generally require the length of the underlying sampling interval for the returns within $[t - h, t]$ approaches zero (even for the ARCH filters and smoothers where h itself is shrinking). We next turn to a more detailed discussion of these different procedures for modeling and measuring volatility within the context of the general setup in Section 2.1.

3. PARAMETRIC METHODS

Parametric volatility models and their implementation constitute one of the cornerstones of modern empirical asset pricing, and a large econometrics and statistics literature has been devoted to the development and theoretical foundation of differently parameterized volatility models. A thorough review of this literature is beyond the scope of this chapter; see, e.g., the existing surveys by Andersen et al. (2006a), Bollerslev et al. (1992),

Bollerslev et al. (1994), Ghysels et al. (1996), and Shephard (1996). Instead, we shall merely highlight the most important ideas as they relate specifically to the volatility measurement for the different model classes.

3.1. Continuous-Time Models

Much of the theoretical asset pricing literature is cast in continuous time. Within this tradition, the sample path of the price process is also commonly assumed to be continuous. This approach is convenient because the representation in Proposition 1 then ensures that, locally, the mean and variance are of the same order. Consequently, the framework effectively involves a dynamic mean–variance trade-off, which typically allows for a tractable analysis of asset pricing and portfolio choice problems. On the other hand, we usually do not observe a record of continuously evolving asset prices, and all but the very simplest specifications tend to imply intractable conditional return distributions for the corresponding discretely observed returns. This issue has historically inhibited empirical work on estimation and inference for realistic continuous-time asset price processes, although a burst of research activity in this area over the last few years has allowed important headway to be made. As a result, the parametric approach to continuous-time modeling is beginning to have a practical impact on return volatility modeling. We will not discuss estimation and inference techniques for this class of model in any detail, however, but rather outline the conceptual issues that distinguish this approach from the discrete-time modeling approach discussed above and the nonparametric volatility measurement discussed subsequently. Other chapters in this handbook offer extensive coverage of parametric and semi(non)parametric estimation techniques for diffusion processes (e.g., Ait-Sahalia et al., 2010; Bandi and Phillips, 2010; Bibby et al., 2010; Gallant and Tauchen, 2010; Jacod, 2010; Johannes and Polson, 2010)

The continuous-time parametric models are directly compatible with the no-arbitrage framework outlined in Section 2, so the specific volatility concepts carry over without modification. However, the specifications of the models traditionally adapted in the literature differ from the general semimartingale representation, and instead rely (implicitly) on Proposition 3 in expressing the models (in short-hand format) as SDEs driven by underlying Brownian motions and, in the case of discontinuities, Poisson jump processes.

3.1.1. Continuous Sample Path Diffusions

The number of alternative continuous-time specifications for asset returns used in the literature is much too large for a comprehensive review to be included here. For illustrative purposes, we simply consider the relevant volatility concepts implied by a few standard formulations.

The simplest possible case is provided by the time-invariant diffusion,

$$dp(t) = \mu dt + \sigma dW(t), \quad 0 \leq t \leq T, \quad (3.1)$$

which underlies the Black–Scholes option pricing formula. Obviously, this process has a deterministic mean return so the expected return volatility trivially equals the expected notional volatility. Moreover, because the volatility is also constant, the expected notional volatility is identical to the notional volatility. Formally, we thus have for the Black–Scholes setting,

$$\zeta^2(t, h) = E[(r(t, h) - m(t, h))^2 | \mathcal{F}_{t-h}] = v^2(t, h) = \int_{t-h}^t \sigma^2(s) ds = \sigma^2 \cdot h.$$

As discussed further in Section 4.1 below, this model is also straightforward to estimate from discretely sampled data by, e.g., maximum likelihood, as the returns are i.i.d. and normally distributed. Of course, the model is overwhelmingly rejected for moderately frequently sampled data (say, daily, weekly, or monthly), as it fails to accommodate the well-documented strong intertemporal volatility dependencies.

For some price series (notably real commodity prices and exchange rates), it is often sensible to postulate a stationary logarithmic price process. Popular models for such series – inspired by the interest rate literature – include the Ornstein–Uhlenbeck (OU) processes and the square-root, or Cox, Ingersoll and Ross (1985) (CIR), processes. These models take the general form

$$dp(t) = \phi(\mu - p(t))dt + \sigma(t)dW(t), \quad 0 \leq t \leq T. \quad (3.2)$$

The drift specification ensures mean reversion in the process, given appropriate regularity conditions and a well-behaved diffusion (volatility) coefficient process. Letting $\sigma(s) \equiv \sigma$ results in the standard OU model, while having $\sigma(s) \equiv \sigma p^\gamma(s)$ produces a constant elasticity of variance (CEV) model, with the CIR model as a special case for $\gamma = 1/2$. The CEV class of models was first proposed in the asset pricing literature by Cox and Ross (1976), and further popularized for interest rates by Chan et al. (1992). The attraction of the specific OU and CIR formulations stems primarily from the tractable distributions for discretely observed data, and from the accompanying closed-form solutions for many related asset and derivatives pricing problems. Explicit solutions for the expected volatility and the expected notional volatility may be derived from existing results in the literature. One immediate observation is that these two volatility concepts now differ as the return innovations will impact the mean process randomly over the forecast horizon. Nonetheless, the expected notional volatility will remain the dominant component in empirically realistic situations.

Example 3 *Ornstein–Uhlenbeck (OU) Processes*

To illustrate, consider the simple OU process,

$$dp(t) = -\phi p(t)dt + \sigma dW(t), \quad (3.3)$$

where for simplicity, we fix $\mu \equiv 0$. Also, for simplicity and without loss of generality, consider the return $r(h, h) \equiv r(h)$ over the $[0, h]$ time interval, or $t - h = 0$. The associated martingale component, $M(h) = \sigma \int_0^h dW(s) = \sigma \cdot W(h)$, then implies that the notional volatility equals $v^2(h, h) = \sigma^2 \cdot h$. Furthermore, the explicit solution to the OU SDE takes the form

$$\begin{aligned} r(h) &= p(0)(\exp(-\phi h) - 1) + \sigma \int_0^h \exp(-\phi(h-s)) dW(s) \\ &= \{p(0)(\exp(-\phi h) - 1) + \sigma \int_0^h [\exp(-\phi(h-s)) - 1] dW(s)\} + \sigma \int_0^h dW(s) \quad (3.4) \\ &= \mu(h) + M(h). \end{aligned}$$

Notice, the predictable component, corresponding to the first parenthesis in the second equation, only depends on the martingale innovation process through a weighted average of past realizations, as the current realization of $W(h)$ receives zero weight from $[\exp(-\phi(h-s)) - 1]$ at time $s = h$. Moreover, by (conditional) normality of the OU process, the expected volatility, $\zeta^2(h, h) = E[(r(h) - m(h))^2 | \mathcal{F}_0]$, may be expressed as

$$\begin{aligned} \zeta^2(h, h) &= [(1 - \exp(-2\phi h))/(2\phi)]\sigma^2 \approx \sigma^2 \cdot h - \phi h^2 \sigma^2 + (2/3) \cdot \sigma^2 \phi^2 h^3 \\ &= E[v^2(h, h) | \mathcal{F}_0] - \phi h^2 \sigma^2 + (2/3) \cdot \sigma^2 \phi^2 h^3. \end{aligned}$$

Since $\phi > 0$, the expected volatility is thus locally smaller than the expected notional volatility. This occurs because of the mean-reverting drift coefficient. Large return innovations will tend to be partially undone over the forecast horizon. However, to first order in h , expected volatility equals expected notional volatility, confirming the crucial role of the latter concept.¹⁵ Further, in reference to Eq. (2.16) in Section 2, it is possible to show that

$$\begin{aligned} \text{Var}[\mu(h) | \mathcal{F}_0] &= [h + \{1 - \exp(-2\phi h)/(2\phi)\} - 2\{1 - \exp(-\phi h)/\phi\}] \cdot \sigma^2 \\ &\approx (\phi^2/3)h^3 \sigma^2, \end{aligned}$$

and

$$\text{Cov}[M(h), \mu(h) | \mathcal{F}_0] = [\{1 - \exp(-\phi h)/\phi\} - h] \cdot \sigma^2 \approx -(\phi/2)h^2 \sigma^2 + (\phi^2/6)h^3 \sigma^2.$$

¹⁵A simple numerical example illustrates the orders of magnitude. The OU process is typically estimated, or calibrated, to capture slowly evolving long-run swings in the logarithmic price process (or interest rate) away from the unconditional mean. Such movements induce a relatively small degree of predictability in the short-term asset returns, but long-term mean reversion, as manifest by a small mean reversion parameter for data calibrated to an annual frequency, say $\phi = 0.1$. At the daily frequency, or $h = 1/250$, clearly $\phi \cdot h^2 \approx 0$ so that the difference between the expected volatility and the notional volatility is negligible. Even at the quarterly frequency, or $h = 1/4$, the deviation is a modest 2.5%. Of course, this number is somewhat sensitive to the assumed strength of the mean reversion.

*Obviously, the contribution of the variation in the drift process is generally (locally) negligible, so that the main contribution to the expected volatility (beyond the notional volatility) stems from the covariance between return innovations and the future path of the mean process. The negative correlation between these components lowers the overall expected volatility (albeit the effect typically is small).*¹⁶

Unfortunately, the entire class of one-factor models covered by Eq. (3.2) falter dramatically when confronted with actual price or return data. In order to obtain more satisfactory empirical fits, the literature has moved towards multi-factor parametric formulations. A natural approach is to let the volatility process be governed by an independent source of random variation, leading to a (genuine) continuous-time SV model. An influential specification is given by the square-root volatility model popularized by Heston (1993) corresponding to $\delta = 1/2$ in the CEV diffusion,

$$d\sigma^2(t) = (\omega - \theta\sigma^2(t))dt + \zeta(\sigma^2(t))^2dV(t), \quad 0 \leq t \leq T, \quad (3.5)$$

where the standard Brownian motion process, $V(t)$, may be correlated with the $W(t)$ process driving the returns, thus introducing an asymmetric return–volatility relation into the asset price dynamics. This model is particularly attractive as it allows for closed-form solutions for option prices. An extensive analysis of multivariate square-root (or affine) processes in modeling term-structure dynamics is provided in Dai and Singleton (2000) (see also Piazzesi, 2010 in this handbook).

Alternatively, the diffusive volatility may be assumed proportional to $\sigma^2(t)$ as in the continuous-time GARCH model of Nelson (1990a),

$$d\sigma^2(t) = (\omega - \theta\sigma^2(t))dt + \zeta\sigma^2(t)dV(t), \quad 0 \leq t \leq T. \quad (3.6)$$

We will return to a more detailed discussion of this specific model in Section 4.1 below on ARCH filters and smoothers. This is also the model used in generating the different volatility sample paths depicted in Fig. 2.1.

Another popular choice is to represent the logarithmic volatility process by an OU diffusion process. As discussed further in Section 3.2.2 below, this formulation corresponds to an (approximate) discrete-time lognormal stochastic autoregressive volatility (SARV)(1) model. In either case, the relation between the expected volatility and the expected notional volatility may be found from the general formula in Eq. (2.16). If the two Wiener innovation processes are correlated, all three terms become operative,

¹⁶Although these calculations are specific to the OU process, the orders of magnitude are indicative of the relative importance of the components governing the expected volatility. In fact, the OU process displays a very strong covariance between the return innovations and the expected returns process, suggesting that this example, if anything, overstates the typical contribution of the terms beyond the (expected) notional volatility in determining the expected volatility for many asset classes.

although the expected notional volatility (expected quadratic variation) continues to dominate empirically.

The SV diffusions above are considerably harder to estimate from discretely observed data than the classical one-factor models of the OU or CIR variety. Intuitively, because of the latent information structure, any inference procedure must either rely on a (potentially noisy) proxy for the latent volatility or integrate out the latent stochastic variable(s) from the model. However, recent progress has made relatively efficient inference possible through a variety of simulation-based procedures such as Efficient Method of Moments (EMM) or Markov Chain Monte Carlo (MCMC) methods. More detailed discussions of these procedures are available in other chapters in the handbook.

Also, optimal measurements of the latent instantaneous volatility process may, in principle, be obtained by standard nonlinear filtering and smoothing procedures (e.g., Kitagawa, 1987), although the direct implementation of these procedures in the present context typically involves prohibitively expensive high-dimensional integration. Important advances to circumvent these problems allowing for the practical numerical calculation and extraction of latent volatility measurements include the particle filters in Pitt and Shephard (1999) and the reprojection approach advocated by Gallant and Tauchen (1998). Again, we refer to other chapters in this handbook for a more detailed treatment of these procedures. We will, however, return to a discussion of specialized continuous-time filtering methods in Section 4.1 below.

Meanwhile, the mounting empirical evidence obtained from the estimation of the continuous-time SV models discussed immediately above clearly suggest that while the models do provide major improvements over the traditional one-factor models in which $\sigma(t)$ is assumed to depend directly $p(t)$ only, the models continue to be decidedly rejected (see, e.g., Andersen et al., 2002; Andersen and Lund, 1997; Bollerslev and Zhou, 2002; Eraker, 2004; Eraker et al., 2003; Gallant and Tauchen, 1997).

These failures have prompted a number of authors to add additional parametrically specified diffusion factors (e.g., Chernov et al., 2003). In light of the general representation in Eq. (2.20) in Proposition 3, it is evident that such multifactor models simply provide an alternative way of specifying the return dynamics that ultimately may be reduced to a single-factor representation for the univariate process. The advantage is that the system may be defined through a sum of different factors, each following a simple dynamic process rather than a single factor with a more complex specification. For example, one may approximate (apparent) long-range dependencies in the volatility process through a sum of multiple distinct AR(1) factors (e.g., Gallant et al., 1999).¹⁷

¹⁷As shown by Chen et al. (2003), nonlinear functions of a continuous-time Markovian process may exhibit long-memory type dependencies in the form of an unbounded spectrum at frequency zero. Alternative diffusive long-memory type formulations have also been considered by Comte and Renault (1996, 1998).

The ability to produce a simple parametric representation is extremely convenient, if not critical, for economic interpretation and implementation of tractable estimation strategies through standard (simulation based) likelihood and method of moments techniques. To illustrate, consider the general k -factor model,

$$r(t) \equiv p(t) - p(0) = \int_{t-h}^t \mu(s) ds + \sum_{j=1, \dots, k} \int_{t-h}^t \sigma_j(s) dW_j(s),$$

where the $\sigma_j(t)$ refer to the j th volatility factor and $W(t) = (W_1(t), \dots, W_k(t))$ denotes a k -dimensional vector process of independent standard Brownian motions. The notional volatility then follows straightforwardly as the sum of the integrated constituent components,

$$v^2(t, h) = \int_{t-h}^t \sigma^2(s) ds = \int_{t-h}^t \left\{ \sum_{j=1, \dots, k} \sigma_j^2(s) \right\} ds = \sum_{j=1, \dots, k} \int_{t-h}^t \sigma_j^2(s) ds.$$

As such, none of the general principles change, but the requisite calculations for, say, the different terms in Eq. (2.16) may certainly become more involved.

It is arguably premature to judge the empirical performance of the parametric multifactor continuous sample path (pure diffusion) volatility models for asset returns, as this work truly is in its infancy. It is clear, nonetheless, that such models serve as alternatives as well as complements for the parametric jump-diffusion models that we turn to next.

3.1.2. Jump Diffusions and Lévy-Driven Processes

At the highest sampling frequencies, there is compelling evidence of the existence of jumps in asset price processes. Specifically, the arrival of important news such as macro-economic announcements (at the aggregate level) or earnings reports (at the firm level) typically induce a discrete jump associated with an immediate revaluation of the asset; see, e.g., Andersen and Bollerslev (1998b), Andersen et al. (2003b), and Johannes (2004) for direct parametric modeling of jumps along with an analysis of their economic import. Likewise, much evidence from the implied volatility literature – which extracts information about market expectations concerning the future return distribution directly from option prices – point toward the importance of incorporating discrete jump probabilities into the analysis of the return dynamics; see, e.g., Bates (1996a) and Bakshi et al. (1997) for earlier work along these lines.

In the same way that the Brownian motion constitutes the basic building block of continuous-time martingales, the standard Poisson jump process serves as the basic building block for pure (compensated) jump martingales (e.g., Merton, 1982). Thus, one may

accommodate the relevant jump features in an arbitrage-free continuous-time logarithmic price process by adding a Poisson jump component with appropriate time variation in the jump intensity and/or in the jump distribution (as in Example 1 and 2 in Section 2 above). In line with this reasoning, let $q(t)$ denote a Poisson point process, with $dq(t) = 1$ indicating a jump at time t , and $dq(t) = 0$ otherwise, and (possibly time-varying) jump intensity denoted $\lambda(t)$.¹⁸ Also, let the random jump size be denoted by $\kappa(t)$, where the process is only defined for $dq(t) = 1$. We then have the general representation

$$\begin{aligned} r(t, h) &= \mu(t, h) + M(t, h) \\ &= \int_{t-h}^t \mu(s) ds + \int_{t-h}^t \sigma(s) dW(s) + \sum_{t-h \leq s \leq t} \kappa(s) \cdot dq(s). \end{aligned} \quad (3.7)$$

The associated notional volatility process explicitly incorporates the jumps,

$$\begin{aligned} v^2(t, h) &\equiv [M, M]_t - [M, M]_{t-h} = [M^c, M^c]_t - [M^c, M^c]_{t-h} + \sum_{t-h \leq s \leq t} \Delta M^2(s) \\ &= \int_{t-h}^t \sigma^2(s) ds + \sum_{t-h \leq s \leq t} \kappa^2(s) \cdot dq(s). \end{aligned} \quad (3.8)$$

The computation of the corresponding expressions for the expected notional volatility and the expected volatility will depend on the specific parametric formulation.

To illustrate, consider the simple jump diffusion in Merton (1976) with constant mean and diffusion volatility coefficients as well as i.i.d. jumps; i.e., $\sigma(t) \equiv \sigma$, $\lambda(t) \equiv \lambda$. Also, denote the mean and the variance of the jump distribution by μ_κ and σ_κ^2 , respectively. The notional volatility is now a stochastic variable, reflecting the random occurrence of jumps. Moreover, only if the mean jump size is zero, $\mu_\kappa = 0$, the expected notional volatility will coincide with the expected volatility (the latter does not contain the squared mean term, μ_κ^2 , below),

$$E[v^2(t, h) | \mathcal{F}_{t-h}] = \sigma^2 \cdot h + E[\sum_{t-h \leq s \leq t} \Delta M^2(s) | \mathcal{F}_{t-h}] = \sigma^2 \cdot h + \lambda \cdot h \cdot (\mu_\kappa^2 + \sigma_\kappa^2).$$

Also, considering the corresponding normalized expected volatility for $h \rightarrow \infty$, it follows from Eq. (2.16) that

$$\lim_{h \rightarrow \infty} \zeta^2(t, h)/h = \sigma^2 + \lambda \cdot (\mu_\kappa^2 + \sigma_\kappa^2),$$

¹⁸Formally, $P[q(t) - q(t-h) = 0] = 1 - \int_0^h \lambda(t-h+s) ds + o(h)$, $P[q(t) - q(t-h) = 1] = \int_0^h \lambda(t-h+s) ds + o(h)$, and $P[q(t) - q(t-h) \geq 2] = o(h)$.

which differs from the instantaneous volatility associated with the continuous martingale component, as defined in Eq. (2.19), and here restricted to be constant, $\sigma^2(t) \equiv \sigma^2$.

In this context, it is also important to recognize that the assumption on $\sigma(t)$, and the corresponding continuous sample path martingale representation in Proposition 3, does not rule out jumps in the $\sigma(t)$ process. However, the presence of jumps in either $\sigma(t)$ and/or $r(t, h)$ invalidates the standard consistency arguments underlying the nonparametric ARCH filters and smoothers discussed in Section 4.1 below,¹⁹ while the so-called realized volatility measures in Section 4.2 generally remains consistent for the notional volatility, even in the presence of jumps.

As for the multifactor parametric diffusion representations, the empirical evidence on jump diffusion models for asset returns is still inconclusive. Early work estimated (overly) simple representations in line with the time-invariant jump diffusion discussed above (e.g., Akgiray and Booth, 1986; Ball and Torous, 1985; Jarrow and Rosenfeld, 1984; Press, 1967), but these models are clearly at odds with the data. More realistic models have recently been explored by, e.g., Andersen et al. (2002), Duffie et al. (2000), Eraker (2004), Eraker et al. (2003), and Pan (2002). Although these formulations improve dramatically on the fit of traditional univariate diffusion and standard SV representations, a general consensus about the relative performance of the various alternative specifications remains elusive at this (early) point.

Another recent proposal is to retain the continuous sample path strategy for the asset returns, but model the volatility process as a non-Gaussian OU process driven by pure upward Lévy jumps (e.g., Barndorff-Nielsen and Shephard, 2001). A primary motivation for this approach is to retain analytic tractability of the temporal aggregation process involved in the construction of volatility forecasts within a reasonably descriptive continuous-time setting.²⁰ Technically, the jumps in the volatility process introduces no new conceptual theoretical issues, as the Lévy processes are semimartingales, and as such the general apparatus for diffusion processes discussed above applies directly. Lévy-driven long-memory type formulations have also been proposed by Anh et al. (2002) and Brockwell and Marquardt (2005). The empirical implementation of these approaches are still in their infancy, but the preliminary results are intriguing.

3.2. Discrete-Time Models

Even if trading and pricing are naturally thought of as evolving in continuous time within the frictionless no-arbitrage setting outlined in Section 2, it is often more convenient to work directly with parametric models for the associated discrete-time returns. Such an approach is naturally motivated by situations in which prices are only observed

¹⁹As discussed further below, the consistence of the ARCH filters may still be established on a case-by-case basis for certain jump processes.

²⁰Although the markets are formally incomplete in this situation, a corresponding analytical option pricing formula based on the minimal-entropy martingale measure have been developed by Nicolato and Venardos (2003).

at regular fixed time intervals (daily closing prices, end-of-the-month prices). Alternatively, if trading is only feasible at given discrete points in time, the relevant return distribution is fully described by the conditional discrete-time dynamics. Either perspective allows us to embed the discrete-time ARCH and SV models in our basic continuous-time setting. Hence, for the remainder of this section, we assume that prices are only observed (and trades only possible) at discrete and equally spaced points in time, $t = 0, h, 2 \cdot h, \dots, T - h, T$, where by assumption T is proportional to h .

The discrete-time models, at a minimum, assume that the correct specification of the one-step-ahead conditional mean and variances is known up to a low-dimensional parameter vector. That is, the models (parsimoniously) parameterize the first two conditional return moments,²¹

$$m(t, h) = E[r(t, h) | \mathcal{F}_{t-h}] = E[\mu(t, h) | \mathcal{F}_{t-h}] = \mu(t, h) \quad (3.9)$$

$$\zeta^2(t, h) = E[(r(t, h) - m(t, h))^2 | \mathcal{F}_{t-h}], \quad (3.10)$$

where $m(t, h)$ and $\mu(t, h)$ coincide because the one-step-ahead conditional mean is predictable. Of course, in contrast to the continuous sample path diffusion models corresponding to $h \rightarrow 0$, which may be defined completely through the instantaneous drift and volatility coefficients, the first two conditional moments of the one-period returns do not fully characterize the dynamic return distribution.

The restriction of only observing prices at equidistant points in time is readily interpreted, within the continuous-time setting, as a pure jump process with known jump times but random jump sizes. In the notation of the previous section,

$$\Delta M(t) = r(t, h) - \mu(t, h) = r(t, h) - m(t, h), \quad t = h, \dots, T. \quad (3.11)$$

As such, it follows directly from the definition of the notional volatility over $[t - h, t]$ that

$$v^2(t, h) = [M, M]_t - [M, M]_{t-h} = \Delta M^2(t). \quad (3.12)$$

Moreover, the expected notional volatility over $[t - h, t]$ simply equals the conditional one-period-ahead variance as specified by the model,

$$\zeta^2(t, h) = E[v^2(t, h) | \mathcal{F}_{t-h}] = E[\Delta M^2(t) | \mathcal{F}_{t-h}]. \quad (3.13)$$

²¹The information set, \mathcal{F}_{t-h} , is (implicitly) restricted to the corresponding discrete-time realizations of the process along with any other discrete-time (possibly latent) state variables.

This same result is generally not true for multiperiod forecasts, or volatilities over longer horizons, $[t - k \cdot h, t]$ where $k > 1$. In this situation, any variation in the conditional mean process within the forecast horizon will contribute to the return variation, so the expected notional volatility typically is not equal to the expected volatility. However, as discussed further below, the contribution from the variation in the conditional mean will usually only be of second-order importance unless the forecast horizon is very long.

To more explicitly clarify the relationship between the notional volatility and total return variability within the multiperiod setting, recall the generic return decomposition for a discrete-time pure-jump process in Proposition 1,

$$r(t) = \mu(t) + M^J(t), \quad t = 0, h, 2 \cdot h, \dots, T - h, T, \quad (3.14)$$

where

$$\mu(t) = \Sigma_{\tau=1, \dots, t/h} E(r(\tau \cdot h, h) | \mathcal{F}_{(\tau-1) \cdot h}) = \Sigma_{\tau=1, \dots, t/h} \mu(\tau \cdot h, h), \quad (3.15)$$

$$M^J(t) = \Sigma_{\tau=1, \dots, t/h} \Delta M(\tau) = \Sigma_{\tau=1, \dots, t/h} (r(\tau \cdot h, h) - \mu(\tau \cdot h, h)) \quad (3.16)$$

Now, for any integer $k > 0$,

$$\mu(t, k \cdot h) = \mu(t) - \mu(t - k \cdot h) = \Sigma_{\tau=1, \dots, k} \mu(t - (k - \tau) \cdot h, h).$$

This $\mu(t, k \cdot h)$ term represents the cumulative conditional one-period-ahead expected returns and not the conditional multistep-ahead expected return. Specifically, for $k > 1$, the term $\mu(t, k \cdot h)$ is generally not equal to $E(r(t, k \cdot h) | \mathcal{F}_{t-k \cdot h}) = m(t, k \cdot h)$, even though

$$E[\mu(t, k \cdot h) | \mathcal{F}_{t-k \cdot h}] = m(t, k \cdot h). \quad (3.17)$$

Hence, we obtain the decomposition of the k -period-expected volatility over $[t - k \cdot h, t]$,

$$\begin{aligned} \zeta^2(t, k \cdot h) &= E[(r(t, k \cdot h) - m(t, k \cdot h))^2 | \mathcal{F}_{t-k \cdot h}] \\ &= E[M^2(t, k \cdot h) + (\mu(t, k \cdot h) - m(t, k \cdot h))^2 + 2 \cdot M(t, k \cdot h) \cdot \mu(t, k \cdot h) | \mathcal{F}_{t-k \cdot h}] \\ &= E([M, M]_t - [M, M]_{t-k \cdot h} | \mathcal{F}_{t-k \cdot h}) + \text{Var}[\mu(t, k \cdot h) | \mathcal{F}_{t-k \cdot h}] \\ &\quad + 2 \cdot \text{Cov}[M(t, k \cdot h), \mu(t, k \cdot h) | \mathcal{F}_{t-k \cdot h}]. \end{aligned} \quad (3.18)$$

Trivially, as noted above, for the one-period-ahead forecasts, or $k = 1$, there cannot be any within-period variability in the conditional mean process, so the last two terms in (3.18)

vanish, and the expected volatility equals the expected notional volatility. However, for multiple-period forecasts, the stochastic evolution of the conditional mean within the interval contributes to the overall return variability, both through the variation in the conditional mean itself and through the covariance between the return innovations and future (within forecast horizon) changes in the conditional mean return. However, the period-by-period conditional mean is generally much smaller than the volatility, and the shifts in the conditional mean are smaller yet. Hence, the expected notional volatility, or expected quadratic variation, remains the dominant component for the multiperiod return variability in empirically realistic situations.²²

Further, notice that the so-called *Leverage Effect* (e.g., Black, 1976) impacts only the expected notional volatility (expected quadratic variation) and none of the other terms. The hypothesis stipulates a (negative) correlation between the return innovations, ΔM , and the size of future return innovations, $(\Delta M)^2$, essentially predicting a left-skewed distribution for the return innovations. With no impact on the conditional mean, only the quadratic variation process is affected. The closely related *volatility feedback effect* (e.g., Campbell and Hentschel, 1992) has an impact through the covariance term, but it remains limited by the size of the shifts in the conditional mean. Again, the hypothesis essentially implies a leftward skew in the return innovation distribution. Intuitively, given a positive volatility risk premium, large negative return innovations are magnified, whereas large positive return innovations are dampened due to the increase in the expected future return required to compensate for a positive and persistent shock to future volatility. Hence, technically, the volatility feedback not only tends to raise the expected volatility directly but it also induces a negative correlation between the return innovations and the future expected mean returns.

Returning to the basic discrete-time setup, the conditional moments in Eqs. (3.9) and (3.10) allow for relatively easy and consistent statistical inference concerning the unknown parameters by a standard generalized method of moments (GMM) estimator (Hansen, 1982), or for SV and latent state variable(s), a Simulated Method of Moments (SMM) type estimator (Duffie and Singleton, 1993). Of course, simple method-of-moments estimators with ill-chosen moment conditions may behave poorly, both asymptotically and in finite samples (Andersen et al., 1999), and much of the literature on discrete-time volatility models has been concerned with the development of more efficient estimation procedures under auxiliary assumptions. In particular, assuming that the standardized innovations, $(r(t, h) - \mu(t, h))/\zeta(t, h)$, belong to a specific parametric family of distributions, maximum likelihood estimation (MLE) and corresponding

²²Of course, the exact terms involved in the multiple-period volatility forecasts will depend upon the specific functional form and the underlying distributional assumptions. Their practical computation may not be trivial, or even feasible in closed form, necessitating the use of numerical simulation techniques (see, e.g., Geweke, 1989). We shall not be concerned with these more computationally oriented aspects of the problem in this chapter.

Gaussian quasi-MLE (QMLE) procedures (Bollerslev and Wooldridge, 1992) are both conceptually straightforward to implement for the ARCH class of models, while more complicated procedures are generally required for discrete-time SV models.

Next, we briefly review some of the popular discrete-time parametric volatility models. The key distinguishing features for each class of models consist of the functional form for the conditional moments in Eqs. (3.9) and (3.10), the variables in the information set \mathcal{F}_{t-h} , along with any additional distributional assumptions. The performance of the different models, such as the fit to the data and precision of forecasts, as well as the ease of computing parameter estimates and the various terms in the volatility forecast expressions, depends importantly on these features.

3.2.1. ARCH Models

The ARCH class of models was first introduced in the seminal paper by Engle (1982). It has since enjoyed unprecedented empirical success along with a myriad of extensions and further theoretical developments. Indeed, most of our empirical knowledge to date concerning the temporal dependencies in financial market volatility have arguably been gleaned from estimation and inference with ARCH type models. Several surveys of this burgeoning literature already exist (an incomplete list of which includes, Andersen and Bollerslev, 1998c; Andersen et al., 2006a; Bollerslev et al., 1992, 1994; Diebold and Lopez, 1995; Engle and Kroner, 1995; Engle, 2004; Engle and Patton, 2001), and we will not attempt yet another comprehensive review. However, it is useful to briefly summarize the key developments and model formulations within the current framework.

The ARCH class of models differ from the discrete-time SV models discussed below, in that the parameterized conditional expectations in Eqs. (3.9) and (3.10) depend exclusively on directly observable variables. This assumption greatly facilitates statistical inference vis-a-vis SV models, and the widespread empirical use of ARCH style models, in part, stems from the ease with which traditional (quasi-) maximum likelihood-based procedures may be implemented.

Any time series model in which the conditional variance depends nontrivially on the time $t - h$ observable information set is now commonly referred to as an ARCH model. This terminology is explained by the particular parametric formulation first adapted by Engle (1982). Specifically, in the so-called ARCH(p) model, $\zeta^2(t, h)$ is parameterized as an autoregressive distributed lag of p -squared innovations,

$$\zeta^2(t, h) = \omega + \sum_{j=1, \dots, p} \alpha_j \cdot (r(t - j \cdot h, h) - \mu(t - j \cdot h, h))^2 \equiv \omega + \alpha(L, h)(r(t, h) - \mu(t, h))^2, \quad (3.19)$$

where $\omega > 0$ and $\alpha_j \geq 0$ to ensure positivity of $\zeta^2(t, h)$ (a.s.), and the $\alpha(L, h)$ lag polynomial is defined by $\alpha_1 L^h + \alpha_2 L^{2h} + \dots + \alpha_p L^{ph}$. Meanwhile, a more parsimonious characterization of the intertemporal volatility dependencies is often obtained by the

generalized ARCH, or GARCH(p, q), model (Bollerslev, 1986),

$$\begin{aligned}\zeta^2(t, h) &= \omega + \sum_{j=1, \dots, p} \alpha_j \cdot (r(t - j \cdot h, h) - \mu(t - j \cdot h, h))^2 + \sum_{i=1, \dots, q} \beta_i \cdot \zeta^2(t - j \cdot h, h) \\ &= \omega + \alpha(L, h)(r(t, h) - \mu(t, h))^2 + \beta(L, h)\zeta^2(t, h).\end{aligned}\quad (3.20)$$

For the popular GARCH(1,1) model, the parameter restrictions $\omega > 0$, $\alpha_1 \geq 0$, and $\beta_1 \geq 0$ obviously guarantees positivity of $\zeta^2(t, h)$. Corresponding conditions for the general case are presented in Nelson and Cao (1992). Rearranging the terms, the GARCH(p, q) model is readily interpreted as an ARMA model for $[r(t, h) - \mu(t, h)]^2$ in which the autoregressive and moving average polynomials are given by $[\alpha(L, h) + \beta(L, h)]$ and $[1 - \beta(L, h)]$, respectively.²³ Hence, provided that all the roots of the characteristic equation, $\alpha(x, h) + \beta(x, h) = 1$, have norm greater than one, the model is covariance stationary, and the unconditional h -period (one-period) variance equals $E[\zeta^2(t, h)] = \omega(1 - \alpha(1, h) + \beta(1, h))^{-1}$. Weaker conditions for strict stationarity have been derived by Nelson (1990b) and Bougerol and Picard (1992), while higher order moment conditions have been developed by Ling and McAleer (2002) among others.

The leverage effect, briefly discussed earlier, stipulates a negative correlation between current return innovations and future expected conditional variances. The GJR–GARCH model (Glosten et al., 1993), in which the α_j coefficients in $\alpha(L, h)$ in Eq. (3.20) depend on the sign of the corresponding return innovations, $r(t - j \cdot h, h) - \mu(t - j \cdot h, h)$, was specifically designed to accommodate such asymmetries. A similar motivation underlies the EGARCH model in Nelson (1991). Defining the standardized innovations,

$$z(t, h) \equiv (r(t, h) - \mu(t, h))/\zeta(t, h), \quad (3.21)$$

the EGARCH(p, q) model takes the form

$$\begin{aligned}\log[\zeta^2(t, h)] &= \omega + \alpha(L, h) \{ \theta \cdot z(t, h) + \gamma \cdot [|z(t, h)| - E(|z(t, h)|)] \} \\ &\quad + \beta(L, h) \log[\zeta^2(t, h)],\end{aligned}\quad (3.22)$$

where as before $\alpha(L, h)$ and $\beta(L, h)$ denote p th- and q th- order lag polynomials, respectively. Obviously, for $\theta < 0$, the model predicts a negative relation between current returns and future conditional variances. The logtransform complicates the calculation of (unbiased) multistep conditional variance forecasts but conveniently avoids having to impose nonnegativity constraints on the parameters. The EGARCH model also requires a specific distributional assumption for $z(t, h)$.

²³Note that the assumption of a finite second-order moment in the ARMA representation corresponds to finite fourth-order unconditional moments of the returns.

Alternatively, as discussed above, asymmetries in the return–volatility relationship may also be attributed to the so-called volatility feedback effect. This feature is captured by the ARCH-in-Mean type formulation (Engle et al., 1987), in which the functional form for the conditional mean, $\mu(t, h)$, depends explicitly on the conditional variance of the process, $\zeta^2(t, h)$. Which of these competing specifications is best able to capture the empirically observed asymmetry in equity return volatility has been the subject of several empirical studies (e.g., Bekaert and Wu, 2000; Campbell and Hentschel, 1992).

Another important empirical finding concerns the strong degree of volatility persistence estimated with most daily and weekly financial rates of return. This is manifested by the autoregressive polynomials describing the variance dynamics in the GARCH(p, q) formulations, $1 - \alpha(x, h) - \beta(x, h)$, and the EGARCH formulations, $1 - \beta(x, h)$, having (their largest) roots very close to unity. The IGARCH model of Engle and Bollerslev (1986) directly imposes this condition; i.e., $\alpha(1, h) + \beta(1, h) = 1$. However, the imposition of a unit root in the conditional variance arguably exaggerates the true dynamic dependencies, and several alternative long-memory, or fractionally integrated, ARCH type formulations have recently been estimated and analyzed more formally in the literature (e.g., Baillie et al., 1996; Bollerslev and Mikkelsen, 1996; Ding et al., 1993; Giraitis et al., 2000, 2004, 2005; Robinson, 1991, 2001; Zumbach, 2004). Possible explanations for the apparent long-memory dependencies based on the aggregation of multiple volatility components and/or stochastic regime-switching models have been explored by Andersen and Bollerslev (1997), Diebold and Inoue (2001), and Liu (2000) among others (see also the related component model in Engle and Lee, 1999). This remains a very active area of current research.

Our focus in this chapter has been almost exclusively univariate. Nonetheless, most interesting questions in asset pricing finance and risk management call for a multivariate framework involving not just conditional variances but also time-varying conditional covariances. From a conceptual view point, the extension of the univariate ARCH class of models to a multivariate setting presents few new issues. However, conditions to ensure that the parameterized conditional covariance matrices are positive definite (a.s.) and involve only a manageable (small) number of parameters are both important considerations from a practical perspective. In the diagonal GARCH model of Bollerslev et al. (1988), the conditional variances and covariances are parameterized as univariate GARCH(p, q) processes; i.e., the ij th element in the conditional covariance matrix depends on a distributed lag of past values of the same element and the cross products of the corresponding innovations. The related BEKK GARCH formulation (Engle and Kroner, 1995) guarantees that the covariance matrices are positive definite. The constant conditional correlation model in Bollerslev (1990) is empirically among the most frequently applied multivariate ARCH models. This model has recently been extended to incorporate parsimoniously parameterized time-varying conditional correlations by Engle (2002) and Tse and Tsui (2002). Other multivariate formulations

allowing for relatively easy implementation in large dimensions include the R-GARCH model in Gallant and Tauchen (2000), the flexible GARCH model of Ledoit et al. (2003), the regime-switching dynamic correlation model of Pelletier (2006), the sequential conditional correlation model of Palandri (2006), and the matrix EGARCH model of Kawakatsu (2006).

Meanwhile, most industry applications entailing large-scale covariance matrix measurements rely on J.P. Morgan's RiskMetrics (Morgan, 1997). The RiskMetrics procedure is based on exponential smoothing, and as such corresponds directly to a diagonal IGARCH(1,1) model in which all the intercepts in the conditional covariance matrix are fixed at zero and identical values of α and $\beta \equiv 1 - \alpha$ are used across all assets. The use of the same smoothing parameter ($\beta = 0.94$ with daily data) obviously facilitates the implementation and automatically guarantees that the covariance matrix measurements are positive definite. Nonetheless, when viewed as a data-generating process as opposed to a filter, the RiskMetrics procedure is formally degenerate (Nelson, 1990b).

One major theoretical drawback to the GARCH class of models concerns their lack of closed-form aggregation. This is true both intertemporally and cross-sectionally. For example, if daily asset returns follow a univariate GARCH(p, q) model, the corresponding weekly returns are not GARCH(p, q). Similarly, if a collection of asset returns follow a multivariate GARCH(p, q) model, (nontrivial) portfolio returns are not GARCH(p, q). The Weak GARCH class of models was explicitly introduced by Drost and Nijman (1993) and Nijman and Sentana (1996) to address this issue. In a weak GARCH model, $\zeta^2(t, h)$ has the interpretation of a parameterized linear projection for the squared innovation. In contrast to the conditional expectations underlying the standard ARCH formulations, the linear projections are closed under temporal (and in the multivariate case cross-sectional) aggregation. However, the linear projections do not easily translate into the volatility concepts in Section 2 and, as emphasized by Meddahi and Renault (1996, 2004), asset pricing relationships are based on conditional expectations as opposed to linear projections. Thus, even though the difference between the linear projections and the true conditional expectations may be numerically small in empirical realistic situations, this limits the applicability and the formal interpretation of the weak GARCH class of models. The discrete-time square-root SARV (SR-SARV) models provide an alternative formulation that circumvent these problems. We next turn to a discussion of this and other discrete-time SV models.

3.2.2. Stochastic Volatility Models

The SV models differ from the ARCH class of models in that the information set, \mathcal{F}_{t-h} , underlying the conditional expectations in Eqs. (3.9) and (3.10) is not directly measurable with respect to the time $t - h$ observable filtration. This is typically the result of the inclusion of two separate stochastic innovations: one innovation term relating the conditional mean of the process to the actually observed return and a second

innovation relating the *latent* volatility process to its conditional mean. This type of formulation is typically motivated by the mixture-of-distributions hypothesis (MDH) and the idea of a latent information arrival process. The MDH was originally put forth by Clark (1973) as a way of conceptualizing the distributional characteristics of speculative returns, and the basic hypothesis has subsequently been extended and analyzed empirically by Epps and Epps (1976), Taylor (1982), Tauchen and Pitts (1983), Andersen (1996), Andersen and Bollerslev (1997), Ané and Geman (2000), among many others, to allow for more realistic temporal dependencies in the underlying latent information arrival process(es). We shall return to a discussion of these ideas in Section 4.2 below. The actual parameterizations of the most popular discrete-time SV models are often rationalized through the discretization of specific continuous-time SV models. We do not provide an exhaustive review of the pertinent discrete-time SV class of models here but simply refer to the excellent surveys offered in Taylor (1994), Shephard (1996), and Ghysels et al. (1996).

In parallel to the GARCH and EGARCH class of models discussed above, most of the parametric SV models used in the literature are based on an autoregressive formulation for a continuous function of the (now) latent volatility process,

$$f[\zeta^2(t, h)] = \omega + \beta(L, h)f[\zeta^2(t, h)] + u(t, h), \quad (3.23)$$

where $\beta(L, h)$ denotes a p th-order distributed lag polynomial, and $u(t, h)$ is a martingale difference sequence; i.e., $E[u(t, h) | \mathcal{F}_{t-h}] = 0$. This class of models is commonly referred to as a SARV(p), model. Intuitively, it is the innovation term, $u(t, h)$, which distinguishes the SV from the ARCH class of models.²⁴ Of course, analogous to the GARCH class of models discussed above, for the SARV(p) model in (3.23) to be well defined, $\zeta^2(t, h)$ must be positive (a.s.). Depending on the functional form for $f(\cdot)$, this restricts the admissible parameters in $\beta(L, h)$ and/or the support of $u(t, h)$. Most of the models estimated in the literature have included only a single lag in the $\beta(L, h)$ polynomial. Conditions to ensure ergodicity and stationarity for the general SARV(1) model are presented in Andersen (1994). The two leading cases are given by the lognormal stochastic autoregressive volatility model in which $f(x) \equiv \log(x)$ and $u(t, h)$ is assumed to be Gaussian, and the square-root,²⁵ or SR-SARV, model corresponding to $f(x) \equiv x$.

The lognormal SV model was first analyzed by Taylor (1982) and subsequently popularized in influential papers by Harvey et al. (1994) and Jacquier et al. (1994). The logarithmic volatility model arises naturally from the standard return formulation, $r(t, h) = \mu(t, h) + \zeta(t, h) \cdot z(t, h)$, in which $z(t, h)$ is an i.i.d. mean zero, unit variance,

²⁴Formal conditions under which the $u(t, h)$ term in the autoregressive formulation cannot be integrated out of the conditional expectations in (3.9) and (3.10), resulting in a genuine SV model, are presented in Andersen (1992).

²⁵This terminology derives from Andersen (1994), who parameterizes an AR(1) model for $f^{-1}[\zeta(t, h)]$. Similarly, the lognormal SV model is sometimes referred to as an exponential SARV model.

white noise process. Rearranging the terms, squaring both sides, and taking logarithms, it follows that

$$y(t, h) \equiv \log[r(t, h) - \mu(t, h)]^2 = \log[\zeta^2(t, h)] + \log[z(t, h)^2]. \quad (3.24)$$

Assuming the mean to be known, this may be interpreted as the measurement equation in a state space representation of the model, with corresponding transition equation defined by the parametric model for $\log[\zeta^2(t, h)]$. In particular, for the lognormal SARV(1) model,

$$\log[\zeta^2(t, h)] = \omega + \beta \cdot \log[\zeta^2(t - h, h)] + u(t, h). \quad (3.25)$$

In this situation, filtered and smoothed measurements of the latent $\log[\zeta^2(t, h)]$ volatility process are readily available by linear Kalman filtering which, as pointed out by Nelson (1988) and Harvey et al. (1994), in turn allows for relatively easy to compute Gaussian QMLE parameter estimates. Of course, the innovations in the measurement equation will generally not be Gaussian, so the Kalman may result in poor measurements of the latent volatility state variable and correspondingly highly inefficient parameter estimates.²⁶

The lognormal SARV(1) formulation may also be justified as a discrete-time approximation to the OU diffusion for the logarithmic instantaneous volatility,

$$d \log(\sigma^2(t)) = -\beta(\log(\sigma^2(t)) - \alpha)dt + \psi dV(t), \quad (3.26)$$

referred to in Section 3.1.1 above. In particular, by a standard Euler scheme, the discrete-time version of the model in (3.26) takes the form

$$\begin{aligned} \log[\zeta^2(t, h)] &= \log[\zeta^2(t - h, h)] - h \cdot \beta \cdot \{\log[\zeta^2(t - h, h)] - \alpha\} \\ &\quad + h^{1/2} \cdot \psi \cdot [V(t, h) - V(t - h, h)]. \end{aligned}$$

The actual parameterization is, of course, different from the model in Eq. (3.25), but the structure corresponds exactly to that of the lognormal SARV(1) model. Interestingly, the continuous-time OU process in Eq. (3.26) also has the interpretation of being the diffusion limit of the discrete-time EGARCH model, in the sense that a sequence of appropriately parameterized EGARCH(1,1) models (as discussed in Section 4.1 below) converges weakly to this model as the length of the sampling interval, h , approaches zero.

Although the latent logarithmic volatility in (3.25) takes the form of an AR(1) model, this translates into an ARMA(1,1) correlation structure for the demeaned logarithmic returns in (3.24), $y(t, h)$. Moreover, following Taylor (1986) and Harvey (1998), this same approximate correlation structure is present for any positive power transform of

²⁶This motivates the extension of the QMLE procedure for the lognormal SV model to a non-Gaussian state space in Kim et al. (1998).

the squared returns; i.e., $\exp[\gamma(t, h)]^c$ where $c > 0$. Hence, the shape of the autocorrelation for the squared returns from the log-normal SARV(1) model mimics that of the empirically popular GARCH(1,1) model.

The second leading class of SV models is given by the SR-SARV(p) model. Following Meddahi and Renault (2004), the SR-SARV(p) model for $\zeta^2(t, h)$ is naturally defined by the marginalization of a p -dimensional latent VAR(1) process. As emphasized by Meddahi and Renault (1996, 2004), this class of models has the advantage of being closed under temporal (and in the multivariate setting cross sectional) aggregation. To appreciate this result, suppose that the true underlying continuous-time volatility is determined by the CEV diffusion,

$$d\sigma^2(t) = (\omega - \theta \cdot \sigma^2(t))dt + \sqrt{2} \cdot \alpha \cdot (\sigma^2(t))^\delta dV(t), \quad (3.27)$$

where $\delta \geq 1/2$ to ensure that the process for $\sigma^2(t)$ is stationary and nonnegative. Since $\zeta^2(t, h)$ is an affine function of $\sigma^2(t)$, it follows that for $\delta \leq 1$, the *exact* discretization of the process must adhere to the basic SR-SARV(1) model structure,

$$\zeta^2(t, h) = \omega + \beta \cdot \zeta^2(t - h, h) + u(t, h), \quad (3.28)$$

where $E[u(t, h) | \mathcal{F}_{t-h}] = 0$. Of course, the h -period time interval is arbitrary so that the expected volatility for the temporally aggregated process, $\zeta^2(t \cdot k, h \cdot k)$, where $k > 1$ and $t = 0, 1, 2, \dots$, must be governed the same AR(1) model structure. As discussed further in Section 4.1 below, the CEV model in (3.27) with $\delta = 1$ may also be interpreted as the diffusion limit of the GARCH(1,1) model.

The discrete-time AR(1) formulations in (3.25) and (3.28) are, of course, somewhat restrictive. In parallel to the developments within the parametric ARCH class of models discussed above, long-memory, or fractionally integrated SV models, better suited at capturing the apparent long-run dependencies in the volatility, have been estimated by Breidt et al. (1998) and Harvey (1998).

Direct extensions of the univariate discrete-time SV models discussed above to a multivariate setting was first explored by King et al. (1994), while earlier work by Diebold and Nerlove (1989) used a related univariate latent ARCH factor structure in parameterizing time-varying conditional covariances. More flexible large-dimensional systems have recently been proposed by Chib et al. (2006).

The SV diffusions (both univariate and multivariate) discussed above are considerably harder to estimate from discretely observed data than the classical one-factor models of the OU or CIR variety, as the inference in essence involve the same complications that plague the estimation of continuous-time SV models. Intuitively, because of the latent information structure, any inference procedure must either rely on a (potentially noisy) proxy for the latent volatility or integrate out the latent stochastic variable(s) from the

model. Again, we refer to other chapters in the handbook for a more detailed discussion of the different procedures designed for doing so,²⁷ as well as the aforementioned surveys by Ghysels et al. (1996) and Shephard (1996). Importantly, from the perspective of volatility measurements, as a by-product of the estimation, many of these procedures result in (approximately optimal) filtered and/or smoothed measurements of the functional latent volatility process, $f[\zeta^2(t, h)]$, conditional on the underlying parametric model and the observable information.

4. NONPARAMETRIC METHODS

The data-driven, or nonparametric volatility measurements afford direct empirical appraisals of the notional volatility, $v^2(t, h)$, without any specific functional form assumptions. The most obvious such measure is, of course, given by the ex-post squared return spanning the $[t - h, t]$ time interval. However, even though the (demeaned) squared return generally provides an unbiased estimator for $v^2(t, h)$, it is also a very noisy estimate. The nonparametric measurements more generally achieve consistency by measuring the volatility as (weighted) sample averages of increasingly finer sampled squared (or absolute) returns over (and possibly outside) the $[t - h, t]$ interval. This immediately raises important issues of efficiency, rates of convergence, and the (asymptotic) distributions for the measurement errors associated with different weighting schemes. At a more fundamental level, however, the nonparametric procedures differ importantly in their assumptions about the length of the time interval h . The instantaneous volatility filters, or ARCH filters and smoothers, discussed next, are based on the assumption of ever more observations over ever finer time intervals (a double limit theory), while the realized volatility measures build on the idea of an increasing number of observations over fixed length time intervals (a single limit theory).²⁸

4.1. ARCH Filters and Smoothers

Parametric ARCH models were designed to parsimoniously model the expected volatility as an explicit function of discretely observed returns; i.e., a parameterized conditional expectation, $\zeta^2(t, h) = E[(r(t, h) - m(t, h))^2 | F_{t-h}]$, where $h > 0$ and F_{t-h} denotes the information set generated by the past returns $r(t - h, h), r(t - 2h, h), \dots$. However, as observed by Nelson (1992), these same discrete-time parametric models may alternatively

²⁷In addition to the general inference procedures, some noteworthy procedures explicitly developed for the estimation of discrete-time SV models include the Bayesian MCMC method in Jacquier et al. (1994, 2004) and Kim et al. (1998), the simulated maximum likelihood technique of Danielsson (1994), the Monte Carlo maximum likelihood approach of Sandmann and Koopman (1998), and the direct MLE through recursive numerical integration in Fridman and Harris (1998).

²⁸Of course, the discrete-time parametric ARCH and SV models discussed in Section 3.1 may also be given the interpretation of fixed length interval filters for extracting $v^2(t, h), h > 0$. However, these type of filters are difficult to characterize and formally justify outside the realm of a specific parametric framework.

be given a nonparametric interpretation as filters designed to extract information about the (latent) instantaneous volatility. In particular, assuming that the sample path of the price and the corresponding instantaneous volatility processes are both continuous, then, although formally misspecified at all discrete sampling frequencies, $h > 0$, an appropriately parameterized sequence of ARCH models, or expected (scaled) volatilities $\zeta^2(t, h)/h$, will consistently (for $h \rightarrow \infty$) estimate the instantaneous volatility, σ_t^2 , at each point in time.

To grasp the intuition behind this powerful result, consider the simple continuous-time random walk model in Eq. (3.1), previously studied by Merton (1980) in this context,

$$dp(t) = \mu dt + \sigma dW(t), \quad 0 \leq t \leq T.$$

Suppose that observations are only available at $n + 1$ equally spaced points over the $[t - h, t]$ time interval, where $0 \leq h < t \leq T$; i.e., $t - h, t - h + (h/n), \dots, t - h + (n - 1) \cdot (h/n), t$. By the definition of the process, the corresponding sequence of $i = 1, 2, \dots, n$, discrete (h/n) -period returns,

$$r(t - h + i \cdot (h/n), h/n) \equiv p(t - h + i \cdot (h/n)) - p(t - h + (i - 1) \cdot (h/n)),$$

is then i.i.d. normally distributed with mean $\mu \cdot (h/n)$ and variance $\sigma^2 \cdot (h/n)$. Hence, the MLE of the drift is simply given by the sample mean of the (scaled) returns,

$$\hat{\mu}_n \equiv n^{-1} \cdot \sum_{i=1, \dots, n} (h/n)^{-1} \cdot r(t - h + i \cdot (h/n), h/n) \equiv r(t, h)/h.$$

It follows immediately that

$$E(\hat{\mu}_n) = \mu.$$

This fixed-interval, or in-fill asymptotic, estimator for the drift only depends on h and not n . The sampling frequency is irrelevant, only the span of the data matters. Thus, although $\hat{\mu}_n$ is an unbiased estimator for μ , it is not consistent as $n \rightarrow \infty$.

Consider now the (unadjusted) estimator for σ^2 defined by the sum of the (scaled) squared returns,

$$\begin{aligned} \hat{\sigma}_n^2 &\equiv n^{-1} \cdot \sum_{i=1, \dots, n} (h/n)^{-1} \cdot r(t - h + i \cdot (h/n), h/n)^2 \\ &= h^{-1} \cdot \sum_{i=1, \dots, n} r(t - h + i \cdot (h/n), h/n)^2. \end{aligned}$$

Because

$$E[r(t - h + i \cdot (h/n), h/n)^2] = \sigma^2 \cdot (h/n) + \mu^2 \cdot (h/n)^2,$$

it follows readily that

$$E(\hat{\sigma}_n^2) = \sigma^2 + \mu^2 \cdot (h/n).$$

Hence, the drift induces only a second-order bias, or $O(n^{-1})$ term, in the estimation of σ^2 for $n \rightarrow \infty$. Moreover, this estimator for the diffusion coefficient is consistent as $n \rightarrow \infty$. To see this, note that

$$\begin{aligned} E[r(t-h+i \cdot (h/n), h/n)^3] &= 3 \cdot \mu \cdot \sigma^2 \cdot (h/n)^2 + \mu^3 \cdot (h/n)^3, \\ E[r(t-h+i \cdot (h/n), h/n)^4] &= 3 \cdot \sigma^4 \cdot (h/n)^2 + 6 \cdot \mu^2 \cdot \sigma^2 \cdot (h/n)^3 + \mu^4 \cdot (h/n)^4, \end{aligned}$$

which along with the second moment given above, and the fact that the returns are i.i.d., implies that

$$\text{Var}(\hat{\sigma}_n^2) = 2 \cdot \sigma^4 \cdot n^{-1} + 4 \cdot \mu^2 \cdot \sigma^2 \cdot n^{-2} \cdot h.$$

Hence by a standard law of large numbers,

$$\text{plim}_{n \rightarrow \infty} \hat{\sigma}_n^2 = \sigma^2.$$

The consistency result for the sample variance estimator for the time-invariant diffusion hinges on the true volatility being constant over $[t-h, t]$. Increasing the number of (scaled) squared return observations over the interval then produces an increasing number of unbiased and uncorrelated measures of σ^2 , and simply averaging these yields a consistent estimator.

This basic idea may, given appropriate regularity conditions, be extended to the general class of continuous sample path diffusions considered in Proposition 3 and Eq. (2.21),

$$dp(t) = \mu(t)dt + \sigma(t)dW(t), \quad 0 \leq t \leq T,$$

under the additional assumption that the sample path for the $\sigma(t)$ process also is continuous. The main difference between this general model and the time-invariant diffusion $\sigma(t) \equiv \sigma$ analyzed in detail above, is that the length of the sampling interval, h , now also must shrink to zero as the sampling intensity within the interval, n , increases.

At an intuitive level, by the assumed sample path continuity, the temporal variation in $\sigma^2(t)$ is readily bounded by restricting the length of the time interval, h , over which the variation is measured,

$$\forall \xi > 0, \exists h > 0 : \sup_{t-h \leq \tau \leq t} |\sigma^2(\tau) - \sigma^2(t)| < \xi, \quad (\text{a.s.})$$

Using this result and refining the arguments above, it is possible to show that the analogous time t (unadjusted) sample variance estimator,

$$\hat{\sigma}_{n,h}(t) \equiv n^{-1} \cdot \sum_{i=1, \dots, n} (h/n)^{-1} \cdot r(t-h+i \cdot (h/n), h/n)^2, \quad (4.1)$$

consistently estimates the instantaneous volatility, provided that $h \rightarrow 0$ and $n \rightarrow \infty$ at the proper rates,

$$\text{plim}_{n \rightarrow \infty, h \rightarrow 0} \hat{\sigma}_{n,h}^2(t) = \sigma^2(t),$$

where the convergence is pointwise in probability.

The trade-off between the length of the sampling interval, $h \rightarrow 0$, and the number of observations, $n \rightarrow \infty$, is analogous to the usual bias-variance trade-off encountered in nonparametric kernel estimation. Similarly, the sample variance estimator in Eq. (4.1) corresponds to a flat kernel scheme and the efficiency of this estimator may generally be improved by using a weighted one- or two-sided average of squared returns. That is the motivation behind the ARCH filters and smoothers developed in a series of papers by Nelson (1992, 1996a,b), Nelson and Foster (1994, 1995), and Nelson and Schwartz (1992) [see also the discussion in Drost and Werker (1996); Duan (1997); Fornari and Mele (2001), and Mele and Fornari (2000)].²⁹

To illustrate, consider the GARCH(1,1) filter for the $(1/n)$ -period returns defined in Nelson (1992),

$$\begin{aligned} \hat{\sigma}_n^2(t) &= \omega_n + \alpha_n \cdot r(t, 1/n)^2 + \beta_n \cdot \hat{\sigma}_n^2(t - 1/n) \\ &= \omega_n \cdot (1 - \beta_n)^{-1} + \sum_{i=0, \dots, \infty} \alpha_n \cdot \beta_n^i \cdot r(t - i/n, 1/n)^2, \end{aligned} \quad (4.2)$$

where

$$\omega_n = \omega/n, \quad \alpha_n = \alpha \cdot (1/n)^{1/2}, \quad \beta_n = 1 - \alpha \cdot (1/n)^{1/2} - \theta/n,$$

and where $\omega > 0, \alpha > 0, \theta > 0$, corresponding to $p = q = 1$, $\zeta^2(t, 1/n) \equiv \hat{\sigma}_n^2(t)$, and $\mu(t, 1/n) \equiv 0$ in Eqs. (3.9–3.10) above. This filter again achieves consistency as $n \rightarrow \infty$ for $\sigma^2(t)$,

$$\text{plim}_{n \rightarrow \infty} \hat{\sigma}_n^2(t) = \sigma^2(t),$$

and, as before, the convergence is pointwise in probability. Note, these arguments explicitly rule out jumps, or discontinuities, in either the drift or diffusion coefficients³⁰ so that the sample path for the instantaneous volatility process, $\sigma^2(t)$, is continuous and coincides with that of the expected (scaled) instantaneous volatility, $\lim_{h \rightarrow 0} v^2(t, h)/h$.

²⁹For earlier work on nonparametric diffusion estimation based on stronger assumptions and different asymptotic arguments; see, e.g., Banon (1978), Dohnal (1987), Genon-Catalot et al. (1992), and Florens-Zmirou (1993).

³⁰The consistency of ARCH filters may still be established on a case-by-case basis for certain jump processes. For instance, the Lévy-driven OU SV model of Barndorff-Nielsen and Shephard (2001) permits an ARMA(1,1) representation so that the GARCH(1,1) filter remains consistent for this particular jump model; see Meddahi and Renault (2004) for further discussion along these lines.

Heuristically, the GARCH(1,1) filter works analogously to the sample variance estimator in Eq. (4.1) by using an (infinite weighted) average of increasingly finer sampled squared returns ever closer to time t . However, the continuous record, or in-fill, asymptotics of ever more observations per interval, $n \rightarrow \infty$, over ever smaller time intervals, $h \rightarrow \infty$, is here achieved by a single asymptotic device dictating both the return sampling frequency and the simultaneous down-weighting of the more distant squared return observations, $r(t - i/n, 1/n)^2$, for large values of i . Note that the parameter configuration in (4.2) underlying this result implies that

$$\lim_{n \rightarrow \infty} (\alpha_n + \beta_n) = \lim_{n \rightarrow \infty} (1 - \theta/n) = 1$$

so that the sequence of GARCH(1,1) filters approaches an IGARCH model, as discussed in Section 3.2.1, in the limit.

Besides providing a consistent volatility filter, such sequences of GARCH models have other interesting and useful properties. For example, if the standardized returns,

$$z(t, 1/n) \equiv n^{1/2} \cdot r(t, 1/n) / \hat{\sigma}_n(t - 1/n), \quad (4.3)$$

are i.i.d. normally distributed then, as shown by Nelson (1990a), the sequence of GARCH(1,1) models defined implicitly by (4.2) converges weakly to the continuous-time GARCH model previously defined in Eq. (3.6),³¹

$$\begin{aligned} dp(t) &= \sigma^2(t) dW(t), \\ d\sigma^2(t) &= (\omega - \theta \cdot \sigma^2(t)) dt + (2 \cdot \alpha)^{1/2} \sigma^2(t) dV(t), \end{aligned} \quad (4.4)$$

where the two Wiener processes are uncorrelated, $\text{Corr}(dW(t), dV(t)) = 0$. Of course, it remains true, that when interpreted as a filter, the sequence of GARCH(1,1) models in (4.2) underlying this diffusion limit consistently extracts the instantaneous volatility, $\sigma^2(t)$, for *any* continuous sample path diffusion.

Many other appropriately parameterized ARCH models share this important property. Specifically, consider the sequence of EGARCH(0,1) models defined by

$$\begin{aligned} \log(\hat{\sigma}_n^2(t)) &= \omega_n + \beta_n \cdot \log(\hat{\sigma}_n^2(t - 1/n)) + \theta_n \cdot z(t, 1/n) + \gamma_n \cdot [|z(t, 1/n)| - (2/\pi)^{1/2}] \\ \omega_n &= \alpha \cdot \beta/n, \beta_n = 1 - \beta/n, \quad \theta_n = \rho \cdot \psi \cdot (1/n)^{1/2}, \\ \gamma_n &= \psi \cdot (1 - \rho^2) \cdot (1 - (2/\pi))^{-1/2} \cdot (1/n)^{1/2}, \end{aligned} \quad (4.5)$$

where $\beta > 0$, $\psi > 0$, and the standardized innovations are defined as in Eq. (4.3). Interpreted as a sequence of filters, this similarly provides consistent estimates (as $n \rightarrow \infty$) of

³¹This particular CEV diffusion process for the instantaneous volatility has also previously been analyzed by Wong (1964).

the instantaneous volatility at each point in time for any continuous sample path diffusion of the general form in Eq. (2.21). In parallel to the consistent GARCH(1,1) filter,

$$\lim_{n \rightarrow \infty} \beta_n = 1$$

so that the root in the autoregressive polynomial dictating the exponential decay in the weights associated with the past absolute standardized returns approaches unity. Under the additional assumption of i.i.d. normally distributed standardized returns, the sequence of EGARCH(0,1) models defined by Eq. (4.5) converges weakly to the OU diffusion for $\log(\sigma^2(t))$,

$$\begin{aligned} dp(t) &= \sigma(t)dW(t) \\ d \log(\sigma^2(t)) &= -\beta [\log(\sigma^2(t)) - \alpha]dt + \psi dV(t), \end{aligned} \quad (4.6)$$

where the instantaneous correlation between the two Wiener processes is determined by the leverage parameter ρ ; i.e., $\text{Corr}(dW(t), dV(t)) = \rho dt$.

Because many candidate ARCH models may serve as consistent filters for the instantaneous volatility, this naturally raises the question of efficiency. The asymptotic distribution theory for the filter errors developed by Nelson and Foster (1994) and Nelson (1996a) allows for a formal analysis of this issue. Intuitively, in the diffusion limit (with continuous sample paths), the process is completely characterized by the first two conditional moments, and the optimal ARCH filter matches both of these. These results for continuous-time SDEs carry over to the design of optimal ARCH filters for the type of SDEs used in the formulation of the discrete-time SV models discussed in Section 3.2.2. In this situation, if the conditional distribution of the innovations are sufficiently fat-tailed, estimating $\sigma^2(t)$ by squaring a distributed lag of past absolute returns, as originally proposed by Taylor (1986) and Schwert (1989), may be more efficient than using a distributed lag of past squared returns. A detailed discussion of these results is beyond the scope of this chapter. However, it is worth noting that the comparisons in Nelson and Foster (1994) related to the diffusion in Eq. (4.6) show that asymptotically (for $n \rightarrow \infty$) the efficiency loss in extracting $\log(\sigma^2(t))$ based on the lognormal SARV(1) model in Eq. (3.26) coupled with the (suboptimal) linear Kalman filter can be substantial relative to the (asymptotically) optimal ARCH filter [which essentially looks like the EGARCH filter defined in Eq. (4.5)]. Of course, this still entails an efficiency loss relative to the optimal nonlinear extraction filter (e.g., Kitagawa, 1987), but as noted above, the numerical integration involved in the implementation of such filters is computationally much more demanding than the simple recursions underlying the filtered volatility estimates from ARCH models.

The ARCH filters explicitly restrict the information set used in the extraction of $\sigma^2(t)$ to past and current returns only; i.e., F_t . Asymptotic (for $n \rightarrow \infty$) optimal ARCH

smoothers involving both lagged and future returns have been developed by Nelson (1996b). The basic idea behind the construction of optimal ARCH smoothers exploit principles similar to those involved in the extension of the Kalman filter to a Kalman smoother (e.g., Anderson and Moore, 1979). It is noteworthy that in contrast to the optimal ARCH filters, the resulting optimal ARCH smoothers do not necessarily match the first two conditional moments of the true distribution. An alternative asymptotic distribution theory for analyzing smoothed volatility measurements is provided by the rolling regression approach in Foster and Nelson (1996). We return to a discussion of some of these results in the following section.

4.2. Realized Volatility

The use of historical, ex-post sample variances computed from higher frequency return data as lower frequency volatility measures has many precedents within the empirical finance literature. For example, Poterba and Summers (1986), French et al. (1987), Pagan and Schwert (1990), and Schwert (1989) rely on monthly sample variances computed from daily returns, Dybvig (1993) uses the cumulative sample variance obtained from daily Treasury yields as a diagnostic, noting its link to the square-bracket process from the theory of semimartingales, while Schwert (1990), Hsieh (1991), and Taylor and Xu (1997) exploit intraday data to produce daily sample return variance measures.³² In spite of the intuitive appeal of using sample variance estimators over fixed horizons as simple non-parametric volatility measures, they appear hard to justify theoretically if volatility truly is time varying. However, by connecting the sample variances, termed *realized volatility* in financial economics, to the theory of quadratic variation, it is possible to more formally justify and assess the properties of such measures. Moreover, this approach to volatility measurement has inspired promising and ongoing new research into volatility modeling based on general distributional assumptions. The formal definition is straightforward.

Definition 6 *Realized Volatility*

The realized volatility over $[t - h, t]$, for $0 < h \leq t \leq T$, is defined by

$$v^2(t, h; n) \equiv \sum_{i=1, \dots, n} r(t - h + (i/n) \cdot h, h/n)^2. \quad (4.7)$$

The realized volatility is simply the second (uncentered) sample moment of the return process over a fixed interval of length h , scaled by the number of observations n (corresponding to the sampling frequency $1/n$) so that it provides a volatility measure calibrated to the h -period measurement interval. Although the definition is stated in terms of equally

³²The work of the Olsen & Associates group in Zürich, Switzerland, as highlighted in the book by Dacorogna et al. (2001), has also been extremely influential in promoting the use of high-frequency intraday price data for more effectively measuring and modeling financial market volatility.

spaced observations, most results discussed below carry over to situations in which the realized volatility is based on the sum of unevenly but increasingly finely sampled squared returns.

The realized volatility measure is closely related to, but different from, the theoretical volatility concepts introduced in Section 2. For example, if the mean return is zero, $\mu(t) \equiv 0$, the realized volatility represents the ex-post sample variance computed from n discretely sampled (h/n) -period returns over $[t - h, t]$. In this case, the realized volatility is (ex-ante) unbiased for the expected volatility, $\zeta^2(t, h)$. Formally, we have the following slight extension of Eq. (2.14) (see, e.g., Protter, 1992, Corollary 3 of Theorem 27, Chapter 2).

Proposition 4 *Realized Volatility as an Unbiased Volatility Estimator*

If the return process is square-integrable and $\mu(t) \equiv 0$, then for any value of $n \geq 1$ and $h > 0$,

$$\zeta^2(t, h) = E[v^2(t, h)|\mathcal{F}_{t-h}] = E[M^2(t, h)|\mathcal{F}_{t-h}] = E[v^2(t, h; n)|\mathcal{F}_{t-h}]. \quad (4.8)$$

As such, the *ex-post* realized volatility is an unbiased estimator of *ex-ante* expected volatility. Of course, the zero mean assumption is highly restrictive but, as we discuss later, the result remains *approximately* true for a stochastically evolving mean return process over relevant horizons under weak auxiliary conditions, as long as the underlying returns are sampled at sufficiently high frequencies.

Another link to our previous discussion is provided by the theory of rolling sample variance estimators within the continuous sample path (diffusion) setting, as formally developed by Foster and Nelson (1996).³³ This theory implies that the realized volatility based on increasingly many return observations over finer and finer time intervals is consistent for the corresponding instantaneous volatility. That is, for $h \rightarrow 0$ and $n \rightarrow \infty$ (at proper rates),

$$\text{plim}_{n \rightarrow \infty, h \rightarrow 0} v^2(t, h; n)/h = \text{plim}_{h \rightarrow 0} v^2(t, h)/h = \sigma^2(t).$$

Although this result is of theoretical interest, it is less robust and less useful in practice. One constraint is that the theory excludes jumps in both the return and volatility processes. More importantly, from a practical perspective, the result hinges on the length of the time interval going to zero and the number of observations going to infinity (over the vanishing interval) simultaneously. This construction is hard to mimic in any relevant sense. Market microstructure features invariably limit the number of (effectively) uncorrelated return observations, so even for highly liquid markets, it is not possible to measure returns (or volatilities) instantaneously. We discuss these practical issues in more detail below.

³³See also the related simulation-based evidence in Andreou and Ghysels (2002).

The rolling regression procedures and associated ARCH filters and smoothers for the instantaneous volatilities are also usually based on long (weighted) averages of the returns. Adjacent instantaneous volatility measures will therefore involve overlapping return observations. This renders formal statistical analysis of the time-series properties of any such derived volatility series complex. The realized volatility approach explicitly seeks to avoid such difficulties by fixing $h > 0$ and interpreting $v^2(t, h; n)$ as a measure of the overall volatility for the $[t - h, t]$ time interval. We turn now toward a general discussion of this approach.

The theoretical properties of realized volatility have been discussed from different perspectives in a number of recent studies including Andersen and Bollerslev (1998a), Andersen et al. (2001b, 2003a), and Barndorff-Nielsen and Shephard (2001, 2002a,b). A simple yet fundamental result follows directly by combining the theory of quadratic variation in Proposition 2 with the Definitions 1 and 6.

Proposition 5 *Consistency of Realized Volatility*

The realized volatility provides a consistent nonparametric measure of the notional volatility,

$$\text{plim}_{n \rightarrow \infty} v^2(t, h; n) = v^2(t, h), \quad 0 < h \leq t \leq T, \quad (4.9)$$

where the convergence is uniform in probability.

The notional volatility plays a crucial role in the return dynamics. From the relation between expected notional volatility and expected volatility in Eq. (2.16), the ex-ante expected notional volatility is also the critical determinant of expected volatility. Any empirical measures of (ex-ante expected) notional volatility based on (2.16) will necessarily depend on the assumed parametric model structure. Proposition 5 implies that, in the limit for increasingly finely sampled returns, or $n \rightarrow \infty$, realized volatility is a consistent (nonparametric) estimator of the (realized) notional volatility over any fixed-length time interval, $h > 0$.

Illustration 2 *Continuous-Time GARCH Model (Revisited)*

The first panel in Fig. 2.2 plots the simulated sample path for the one-period notional volatility, $v(t, h)$, $t = 1, 2, \dots, 2500$, for the same continuous-time GARCH model depicted in Fig. 2.1. To illustrate the consistence of the realized volatility (as $n \rightarrow \infty$) for the notional volatility, the last four panels in Fig. 2.2 plot the time series of realized volatilities, $v(t, h; n)$, for n equal to 1, 3, 24, and 288, respectively. The squared returns ($n = 1$) shown in the second panel obviously provide very noise measures of the notional volatilities. While it is possible to pick out the general shape, the plot is extremely erratic, and it would be hard to accurately assess the true value of $v(t, h)$ on a period-by-period basis. Squaring and summing three within period returns, as in the third panel, clearly helps in reducing the noise. Moving one step further in constructing the realized volatilities from $n = 24$ returns, corresponding to an hourly sampling frequency in a 24-h market,

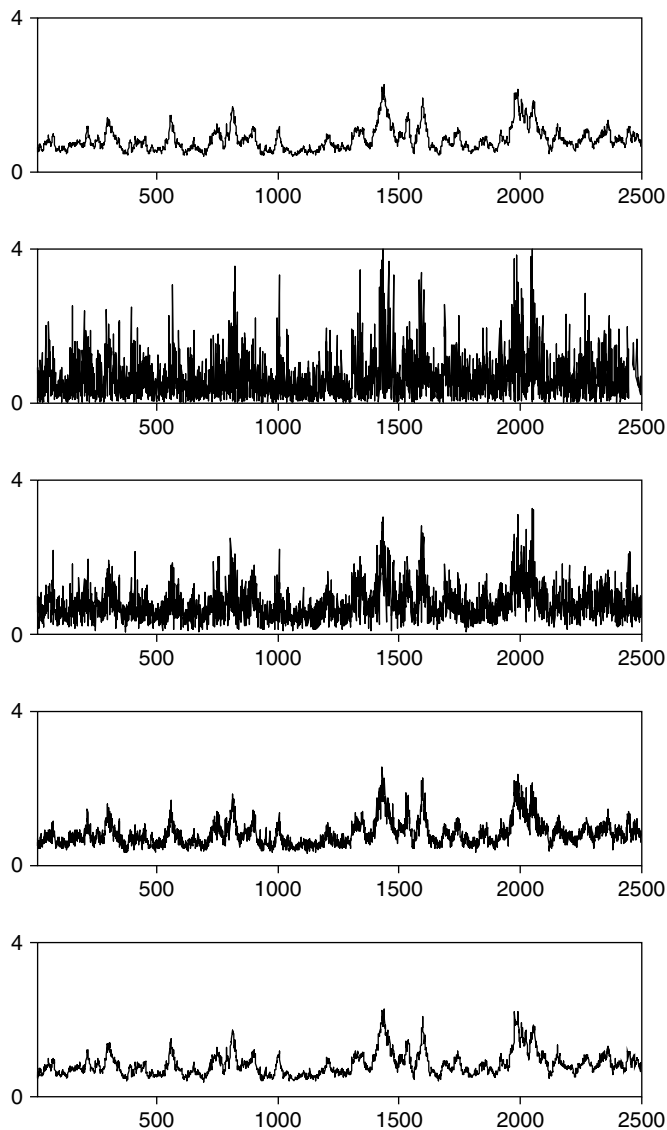


Figure 2.2 The first panel in the figure plots the one-period notional volatility, $v(t, 1)$, $t = 1, 2, \dots, 2500$, from the same continuous-time GARCH model depicted in Fig. 2.1. The remaining four panels show the corresponding realized volatilities, $v(t, 1, n)$, for n equal to 1, 3, 24, and 288, respectively.

or 20-min returns, in a market operating eight hours a day, results in further dramatic improvements. Finally, the final panel for $n = 288$, or five-minute returns in a twenty-four hour market, is almost indistinguishable from the time series of notional volatilities in the top panel.

It is natural to combine the unbiasedness property of realized volatility in Proposition 4 and the consistency result in Proposition 5 to think of ex-post realized volatility measures, in general, as approximately unbiased estimators, and the ex-ante expected values of the realized volatility measures as consistent estimators for the ex-ante expected notional volatility. That is, subject to a uniform integrability condition, as formally discussed in Andersen et al. (2003a),³⁴

$$\text{plim}_{n \rightarrow \infty} E[v^2(t, h; n) | \mathcal{F}_{t-h}] = E[v^2(t, h) | \mathcal{F}_{t-h}], \quad 0 < h \leq t \leq T. \quad (4.10)$$

Importantly, as explained in more detail below, this result and the ability to compute conditional expectations of the notional volatility from the realized volatility in turn allow for the construction of easy-to-implement reduced form volatility forecasting models.

Still, the above consistency results leaves important considerations regarding the size of potential error terms and any finite-sample biases unanswered. We discuss the issue of the measurement errors involved in using realized volatilities for volatility measurement and modeling in more detail below. Meanwhile, it is instructive first to consider a decomposition of the realized volatility measure into the separate terms associated with the potential sources of error and bias. For that purpose, we apply the canonical decomposition to each return component of the realized volatility definition in Eq. (4.7) and simplify notation, so for $i = 1, \dots, n$,

$$\begin{aligned} r(t-h + (i/n) \cdot h, h/n) &= \mu(t-h + (i/n) \cdot h, h/n) + M(t-h + (i/n) \cdot h, h/n) \\ &\equiv \mu_i + M_i. \end{aligned} \quad (4.11)$$

In the frictionless arbitrage-free setting, the return on a risky asset over time intervals of length (h/n) has a martingale innovation of order $(h/n)^{1/2}$, while the corresponding mean component is at most of order (h/n) . In particular, exploiting the notation introduced in Eq. (4.11), we have

$$\begin{aligned} v^2(t, h; n) &= \sum_{i=1, \dots, n} [\mu_i^2 + 2 \cdot \mu_i M_i + M_i^2] \\ &= v^2(t, h) + O_p(n^{-1}) + O_p(n^{-1/2}) + [\sum_{i=1, \dots, n} M_i^2 - v^2(t, h)]. \end{aligned} \quad (4.12)$$

It is apparent that the realized volatility may differ from the notional volatility for two distinct reasons. First, the second and third terms on the right-hand side of the last

³⁴The assumption of a bounded return process provides a simple sufficient condition for this convergence in mean; see, e.g., Hoffmann-Jørgensen (1994), sections 3.22–3.25. For example, one may imagine a bound on the return that prevents a small investment in the asset from ever producing a return that exceeds a (large) multiple of the expected value of all resources available in the worldwide economy. Nonetheless, this result is not true for all admissible price processes covered by Proposition 1; see, e.g., Barndorff-Nielsen and Shephard (2002b) for a counter example.

equation in (4.12) reflect the mean returns, which only truly vanishes in the limit for $n \rightarrow \infty$. However, the expected return over short intervals (large n) are necessarily small, so the contribution from these terms will be empirically negligible. This conclusion is only reinforced by noting that the component of the largest order, $O_p(n^{-1/2})$, represents a covariance term that is limited by the size of the innovations to the expected return over the (h/n) time interval, which typically will be very small. Second, the last term on the right-hand side of (4.12) has the interpretation of a measurement error term, as Proposition 4 shows that the cumulative-squared martingale innovations provide unbiased estimators for the corresponding notional volatility (quadratic variation). Hence, this term has a zero expected value. Nonetheless, for any given value of n , it induces a measurement error that is unrelated to the mean return. This component is the source of empirically relevant deviations between realized volatility and (realized) notional volatility.

The actual size and exact distribution of the errors obviously depend on the particular return process and must be analyzed on a case-by-case basis. Barndorff-Nielsen and Shephard (2002a) provide specific evidence for the OU specification with a background-driving Lévy process. Similarly, Meddahi (2002) presents explicit expressions for the different terms in Eq. (4.12) for the class of eigenfunction SV models and goes on to numerically compare the size of the unconditional variance of the measurement error to the unconditional variance of the notional (integrated) volatility for some of the continuous-time diffusions of the general form in Eq. (2.20) that have been estimated in the existing literature.

The preceding discussion implies that realized volatility is approximately (apart from minor biases induced by the mean component) unbiased for the corresponding notional volatility. Importantly, it also follows from the local martingale property in (2.11) and the decomposition in (4.12), that the associated measurement errors are approximately uncorrelated, i.e.,

$$E[(v^2(t+j, h; n) - v^2(t+j, h)) \cdot (v^2(t, h; n) - v^2(t, h))] = O_p(n^{-1}), \quad (4.13)$$

where $j \neq 0$. Again, the $O(n^{-1})$ term is identically equal to zero in the case of constant mean returns and is otherwise likely to be small in empirically realistic situations. This confirms that realized volatilities provide meaningful and theoretically well-founded volatility measurements. Moreover, they constitute natural and convenient inputs into modeling and inference procedures concerning the expected notional volatility and, by extension, the expected return volatility.³⁵

³⁵This result also underlies the simple GMM estimation procedure for parametric continuous-time SV models in Barndorff-Nielsen and Shephard (2002a) and Bollerslev and Zhou (2002) based on matching sample moments of realized volatility with corresponding model implied moments for notional (integrated) volatility.

It is worth reiterating that the fixed h , large n asymptotics, or realized volatility asymptotics, underlying the results discussed above, is pivotal in practice. In particular, in spite of the theoretical desirability of letting interval size, h , shrink indefinitely as an increasing number of high-frequency return observations is used within each (vanishing) interval (as in the ARCH filters and smoothers discussed in Section 4.1), this idea is difficult (impossible) to mimic in practice. The number of data points, n , that adhere (approximately) to the underlying no-arbitrage semimartingale property over short time intervals is severely limited by various market microstructure frictions. This invariably puts an effective (asset and/or market specific) lower bound on the highest sampling frequency that is applicable in empirical work, say $1/n > 1/N$. We return to this important practical consideration below.

In summary, the realized volatility approach exploiting intraday return observations allow for directly observable return volatility measures that are consistent, approximately unbiased, and have uncorrelated measurement errors. It is natural to exploit these properties by building a time series model directly for the observed realized volatility measures through standard ARMA style modeling. Importantly, such procedures sidestep the complex task of providing an appropriate model for the intraday volatility patterns while still exploiting the inherent information in the high-frequency data for lower frequency volatility movements. Of course, the use of nonparametric volatility measurements invariably entails a loss in statistical efficiency relative to the use of a fully (and by assumption correctly) specified parametric volatility model. We comment further on this issue, and the practical merits of reduced form realized volatility modeling below.

The imposition of additional restrictions on the return process allows for important additional insight into the size and asymptotic distribution of the realized volatility errors. In particular, consider the class of continuous sample path diffusions, characterized by the sde,

$$dp(t) = \mu(t)dt + \sigma(t)dW(t), \quad 0 \leq t \leq T, \quad (4.14)$$

where $\mu(t)$ is predictable and of finite variation (c.f., Proposition 3 and the corresponding SDE in Eq. (2.21)). Extending the infeasible distributional implications of Jacod (1994) and Jacod and Protter (1998), the results of Barndorff-Nielsen and Shephard (2002a, 2004a) provide the following feasible mixed-Gaussian asymptotic (for $n \rightarrow \infty$) approximation to the distribution of the measurement errors.

Proposition 6 *Asymptotic Mixed Normality of Realized Volatility*

The realized volatility errors for the continuous sample path diffusion in Eq. (4.14) is distributed as

$$[\nu^2(t, h; n) - \nu^2(t, h)] \cdot [2/3 \cdot \nu^{[4]}(t, h; n)]^{-1/2} \rightarrow N(0, 1), \quad (4.15)$$

for $n \rightarrow \infty$, where

$$v^{[4]}(t, h; n) \equiv \sum_{i=1, \dots, n} r(t - h + i \cdot (h/n), h/n)^4. \quad (4.16)$$

Importantly, as shown in Barndorff-Nielsen and Shephard (2006b), this proposition remain valid in the presence of leverage effects, or correlations between the $\sigma(t)$ volatility process and the Brownian motion, $W(t)$, dictating the price innovations.³⁶ Thus, this results considerably strengthens the aforementioned convergence of realized volatility to notional volatility (in probability) by providing the asymptotic distribution of the corresponding errors. Formally, the variance of the realized volatility errors is given by $2/3 \cdot \int_{t-h}^t \sigma^4(\tau) d\tau$, which is consistently estimated by $2/3 \cdot v^{[4]}(t, h; n)$ as defined in Eq. (4.16). Hence, the magnitude of the errors depends upon the level of the (latent) volatility. This result represents a fundamental extension of the corresponding expression for the variance of the (uncentered) sample variance for the continuous-time random walk model discussed in Section 4.1 above, $2 \cdot \sigma^4 \cdot h^2 \cdot n^{-1} + O(n^{-1})$. Of course, the general expression for the variance of the realized volatility errors, or $v^{[4]}(t, h; n)$, is straightforward to calculate in practice.

The more powerful distributional result in Proposition 6, compared to the weak convergence in Proposition 4, comes at the cost of the stronger assumption on the underlying log-price process. Although the additional conditions, most notably the absence of jumps in the price path, likely are violated empirically at the highest sampling frequencies, the asymptotic distribution should nonetheless serve as a useful theoretical benchmark for assessing the properties of the realized volatility measures and further assist in guiding empirical procedures.

In this regard, Barndorff-Nielsen and Shephard (2005) find that an improved finite-sample (finite n) approximation may be obtained by the log-linearization,

$$[\log(v^2(t, h; n)) - \log(v^2(t, h)) + 1/2 \cdot s(t, h; n)^2] \cdot s(t, h; n)^{-1} \sim N(0, 1), \quad (4.17)$$

where

$$s(t, h; n)^2 \equiv \max \left\{ 2/3 v^{[4]}(t, h; n) \cdot v^2(t, h; n)^{-2}, 2/n \right\}. \quad (4.18)$$

The upper bound of $2/n$ in Eq. (4.18) arises from imposing the theoretical lower bound for $n \rightarrow \infty$ on the first ratio. This approximation seems to work well, even for moderately sized n (say $n \geq 10$), in a (stylized) simulation setting. The improvement is related to the logarithm delivering a variance-stabilizing transformation. In that sense, the improved finite-sample distribution obtained by Eqs. (4.17) and (4.18) is directly in line with

³⁶This is also corroborated by the related finite-sample (finite n) simulation evidence reported in Andersen et al. (2005).

the evidence for the parametric discrete-time SV models discussed in Section 3.1.1, for which the innovation process for the formulations involving the logarithmic volatility typically exhibits much reduced (conditional) heteroskedasticity.³⁷

The above results speak to the precision in extracting information about the (realized) notional volatility from the realized volatility measures. Realizations of the notional volatility are, of course, of direct interest as indicators of return variability. However, they also provide an indication of the character of the underlying return distribution itself. In particular, it follows under appropriate conditions that the returns, $r(t, h)$, conditional on the notional volatility (and the mean return) over the $[t - h, t]$ return interval will be Gaussian.

Proposition 7 *Normal Mixture Distribution*

The discrete-time returns $r(t, h)$ over $[t - h, t]$, $0 < h \leq t \leq T$, for the continuous sample path diffusion in Eq. (4.14) is distributed as a normal mixture,

$$r(t, h) | \sigma \{ \mu(t, h), v^2(t, h) \} \sim N(\mu(t, h), v^2(t, h)), \quad (4.19)$$

provided that the Brownian Motion, $W(t)$, is independent of $\mu(p(t), \sigma(t))$ and $\sigma(t)$.

Of course, the (ex-ante) mean return and the notional volatility is not directly observable. However, integrating out $\sigma \{ \mu(t, h), v^2(t, h) \}$, the proposition implies that the return distribution conditional on time $t - h$ information should be governed by a normal mixture distribution.³⁸ This is directly in line with the implications of the MDH pioneered by Clark (1973) which, as discussed in Section 3.2.2, has motivated the formulation of some of the most widely used empirical discrete-time SV models.

The consistency of the realized volatility for the notional volatility in Proposition 4 along with the approximate log-normality of the realized volatility distribution and the normal mixture distribution in Proposition 7 suggest a simple alternative empirical return-volatility modeling strategy. Assume that the demeaned returns standardized by the realized volatilities, $[r(t, h) - \mu(t, h)] \cdot v^2(t, h; n)^{-1}$, are (approximately) Gaussian, coupled with a simple reduced form (approximately) Gaussian time series model for the logarithmic realized volatilities, $\log[v^2(t, h; n)]$. Effectively, this modeling strategy relies exclusively on forecasts for the distribution of the future notional volatilities through the observed realized volatilities, and as such is in principle straightforward to implement in practice.³⁹ Moreover, leverage effects, or asymmetries, in the notional volatility are

³⁷This is also consistent with the empirical evidence in Andersen et al. (2001a,b) suggesting that the unconditional distribution of realized volatility is approximately log-normal.

³⁸As discussed further below, the presence of jumps in the price process will generally render the corresponding distribution of the (standardized) returns nonnormal. This may be exploited in the formulation of tests for (the importance of) jumps, as in, e.g., Drost et al. (1998), Ait-Sahalia (2002, 2004), Andersen et al. (2007b), and Andersen et al. (2009).

³⁹Of course, the realized volatility invariably differs from the true notional volatility for finite n . However, the measurement errors are (approximately) serially uncorrelated, and therefore, effectively averaged out in any reduced form time-series model for $v^2(t, h; n)$.

easily incorporated by allowing the time series model for $\log[v^2(t, h; n)]$ to depend (nontrivially) on the level of the (past) returns. This empirical modeling framework has been pursued successfully by Andersen et al. (2003a), who report impressive forecast performance from the estimation of simple standard time series models for the realized volatilities.⁴⁰ Related empirical work by Fleming et al. (2003) also suggests that important improvement can be obtained by using this realized volatility modeling approach in lieu of more standard parametric volatility modeling procedures in practical portfolio allocation decisions.⁴¹

As mentioned repeatedly, the realized volatility approach of holding $h > 0$ fixed is motivated by the fact that it is undesirable, and due to the presence of market microstructure frictions indeed practically infeasible, to sample returns infinitely often ($n \rightarrow \infty$) over infinitesimally short time intervals ($h \rightarrow 0$). To more directly illustrate these issues, suppose that the observed logarithmic price process, say $p^o(t)$, is equal to the true (latent) semimartingale price process that would obtain in the absence of any frictions, $p(t)$, plus a “noise” term, $u(t)$, coming from the use of discrete price grids, bid-ask spreads, and other pertinent market microstructure frictions; see, e.g., Hasbrouck (1996) and Stoll (2000). In this situation, the continuously compounded observed return over the $[t - (i/n)h, t - ((i - 1)/n)h]$ time interval, $i = 1, 2, \dots, n$, is then given by,

$$\begin{aligned} r^o(t - ((i - 1)/n) \cdot h, h/n) &= p^o(t - ((i - 1)/n) \cdot h) - p^o(t - (i/n) \cdot h) \\ &= r(t - ((i - 1)/n) \cdot h, h/n) \\ &\quad + u(t - ((i - 1)/n) \cdot h) - u(t - (i/n) \cdot h). \end{aligned} \tag{4.20}$$

Hence, the realized volatility constructed from the summation of these n -squared returns within $[t - h, t]$ will typically not provide a consistent estimate of the increment to the quadratic variation of the true latent return process, or the notional volatility $v^2(t, h) \equiv [r, r]_t - [r, r]_{t-h}$. Indeed, assuming that the variance of the $u(t - ((i - 1)/n)h)$ process does not depend upon the value of n and is $O(1)$, the realized volatility estimator constructed from the observed returns, $r^o(t - ((i - 1)/n) \cdot h, h/n)$, will generally diverge for $n \rightarrow \infty$. This directly motivates choosing n sufficiently large so as to render the asymptotic results discussed above reliable, yet not too large so as not to overwhelm the estimate by the variation stemming from the noise component. The realized volatility signature plots of Andersen et al. (2000a), in which the sample means of $v^2(t, h; n)$, $t = 1, 2, \dots, T$,

⁴⁰These empirical results have been further corroborated by the corresponding theoretical implications for specific continuous-time SV models derived in Andersen et al. (2004).

⁴¹Many other empirical studies highlighting the potential benefits of the realized volatility framework in volatility forecasting, asset and option pricing, risk management, and other practical financial decision making have emerged over the past few years; see, e.g., Andersen et al. (2005), Areal and Taylor (2002), Bandi et al. (2006), Bollerslev and Zhang (2003), Corsi (2003), Deo et al. (2006), Engle and Gallo (2006), Koopman et al. (2005), Maheu and McCurdy (2002), Martens (2002), and Thomakos and Wang (2003). For a survey of some of these methods, see also Andersen et al. (2006b).

for a large value of T , are plotted against different values of n , provides a simple informal tool for gauging this trade-off and identifying the highest possible sampling frequency at which the impact of the noise appears negligible. More advanced techniques for directly determining the optimal, in a mean-square error sense, value of n has also been developed by Aït-Sahalia et al. (2005), and Bandi and Russell (2006a, 2008). For many actively traded assets, this often implies a value of n equivalent to about five minutes. Although the resulting measurement errors in the realized volatilities invariably depend upon the true underlying model and the exact form of the frictions, it is nonetheless evident that in most empirically realistic situations, the errors are often nontrivial.⁴²

This in turn has inspired the development of several modified realized volatility measures designed to circumvent the impact of the microstructure frictions in estimating the notional volatility. In particular, suppose that the $u(t)$ noise process is i.i.d. It follows then readily from Eq. (4.20) that the discretely observed returns will inherit an MA(1) error structure. Motivated by this, early work along these lines relied on different MA (and AR) filters to mitigate the impact of the noise component; e.g., Andersen et al. (2001a), Areal and Taylor (2002), Bollen and Inder (2002), Corsi et al. (2001), among many others. Similarly, Zhou (1996) first proposed a kernel-based estimator, adding twice the first-order autocovariance to the realized variance as a way to account for the spurious first-order serial correlation induced by the i.i.d. noise component. More sophisticated kernel-based estimators, allowing for a wider variety of dependent noise processes, have been developed by Barndorff-Nielsen et al. (2008) and Hansen and Lunde (2006).⁴³ Alternatively, Zhang et al. (2005) suggested the use of subsampling schemes to correct for the bias induced by the noise component. Intuitively, under fairly general assumptions about the noise process, the bias in the realized volatility will grow at rate n . Thus, by properly combining realized volatilities for different sampling frequencies, it becomes possible to annihilate this first-order bias through a Jackknife-type estimator. As shown in Barndorff-Nielsen et al. (2008), this two-scale estimator may be expressed as a kernel-type estimator. Refined multiscale estimators have also been developed by Aït-Sahalia et al. (2006). Comprehensive surveys of this rapidly growing literature and the many different methods proposed therein can be found in Aït-Sahalia (2007), Bandi and Russell (2006b), Barndorff-Nielsen and Shephard (2006b), and McAleer and Medeiros (2006).

⁴²Model-specific calculations and simulations by Andersen and Bollerslev (1998a), Andersen et al. (1999), Andersen et al. (2004, 2005), Andreou and Ghysels (2002), Bai et al. (2004), Barndorff-Nielsen and Shephard (2005), Barucci and Renò (2002), and Zumbach et al. (2002), among others, illustrate the effects of finite n (and h) in a variety of different settings.

⁴³These estimators have a parallel in the so-called Heteroskedasticity and Autocorrelation Consistent (HAC) estimators used for estimating long-run covariance matrices (e.g., Newey and West, 1987, and Andrews, 1991). Importantly, however, the realized kernel-based estimators are not scaled by the sample size, which make their asymptotic properties very different. This also mirrors (in many ways) earlier developments related to the estimation of Capital Asset Pricing Model (CAPM) beta's in the presence of asynchronous trading effects by Scholes and Williams (1977), and the adjustment to the sample variance in French et al. (1987) obtained by including the cross-product between successive returns.

Meanwhile, the desire to guard against the potentially distorting impact of high-frequency real-world frictions has also inspired the use of alternative robust variation measures. One such measure, dating back to the work of Garman and Klass (1980) and Parkinson (1980), is the range; i.e., the difference between the maximum and the minimum price over some nontrivial $[t - h, t]$ time interval. As argued in Alizadeh et al. (2002) and Brandt and Diebold (2006), range-based volatility measures that involve only two as opposed to a large number of intrainterval price observations are less susceptible to both bid-ask bounce and asynchronous trading effects.⁴⁴ However, this desirable robustness feature must be weighed against the fact that formal statistical analysis of range-based estimators generally require specific distributional assumptions conveniently avoided by the realized volatility measures, as well as other power-based variation measures.

There is a long history in statistics of relying on absolute returns rather than squared returns as more robust (to outliers) measure of the ex-post variation; e.g., Davidian and Carroll (1987).⁴⁵ These results have a direct analog for the continuous sample path diffusion in Eq. (4.14). In particular, returning to the general frictionless arbitrage-free setting, the following definitions of *notional* and *realized power variation*, adapted from Barndorff-Nielsen and Shephard (2003), directly parallel the notional and realized volatility concepts discussed earlier.

Definition 7 *Power Variation Measures*

The notional sth order power variation and the realized sth order power variation, $s > 0$, for the diffusion in Eq. (4.14) over $[t - h, t]$, $0 < h \leq t \leq T$, are defined, respectively, as

$$v^{[s]}(t, h) \equiv \int_{t-h}^t \sigma^s(\tau) d\tau \quad (4.21)$$

and

$$v^{[s]}(t, h; n) \equiv \mu_s^{-1} (h/n)^{1-s/2} \sum_{i=1, \dots, n} |r(t - h + i \cdot (h/n), h/n)|^s, \quad (4.22)$$

where $\mu_s = E(|Z|^s)$, and Z denotes a standard normal distribution.

It is apparent that, for $s = 2$, the definitions correspond directly to the previously discussed notional and realized volatility concepts; i.e., $v^{[2]}(t, h) \equiv v^2(t, h)$ and $v^{[2]}(t, h; n) \equiv v^2(t, h; n)$, respectively. However, other values of s may allow for more robust measurements. In particular, extending the distributional results for the (standard)

⁴⁴More formal statistical properties of realized range-based estimator and comparisons with other realized volatility estimators are discussed in Christensen and Podolskij (2007), Martens and van Dijk (2007), and Dobrev (2007).

⁴⁵Also, as noted in Section 4.1 above, the optimal ARCH filters for discrete-time SV models may entail a distributed lag of past absolute returns as opposed to the squared returns (Nelson and Foster, 1994).

realized volatility in Proposition 5 to the generalized power variation measures defined above, the following proposition follows directly from Barndorff-Nielsen and Shephard (2003).

Proposition 8 *Asymptotic Mixed Normality of Realized Power Variation*

The realized s th order power variation errors, $s \geq 1/2$, for the continuous sample path diffusion in Eq. (4.11) is distributed as

$$\mu_s \cdot \omega_s^{-1/2} (h/n)^{s/2-1} \cdot [v^{[s]}(t, h; n) - v^{[s]}(t, h)] v^{[2s]}(t, h; n)^{-1/2} \rightarrow N(0, 1), \quad (4.23)$$

for $n \rightarrow \infty$, where $\mu_s = E(|Z|^s)$, $\omega_s = \text{Var}(|Z|^s)$, and Z denotes a standard normal distribution.

The special case corresponding to $s = 1$ is naturally termed *absolute variation*. The realized absolute variation is, of course, simply constructed by the (scaled) summation of the n absolute returns, $|r(t - h + i \cdot (h/n), h/n)|$, $i = 1, 2, \dots, n$, within the $[t - h, t]$ time interval. From Proposition 8, the asymptotic (for $n \rightarrow \infty$) distribution of the corresponding measurement error for the notional absolute variation thus satisfies

$$(\pi/2 - 1)^{-1/2} \cdot (h/n)^{-1/2} \cdot [v^{[1]}(t, h; n) - v^{[1]}(t, h)] v^2(t, h; n)^{-1/2} \rightarrow N(0, 1).$$

This provides a formal theoretical basis for gauging the empirical results in Andersen and Bollerslev (1998b) among others based on $v^{[1]}(t, h; n)$. Similarly, these distributional results may be helpful in better understanding the so-called Taylor Effect (e.g., Granger and Ding, 1995), according to which the autocorrelations of power transforms of the absolute returns are maximized (empirically) for values of s close to unity.

Meanwhile, the most desirable feature of the power variation measures arguable relates to their robustness to jumps for appropriate choice of s . In particular, consider the jump-diffusion model discussed earlier in Section 3.1.2 expressed in short-hand sde form,

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) + \kappa(t)dq(t), \quad 0 \leq t \leq T, \quad (4.24)$$

where $q(t)$ denote a Poisson point process, with $dq(t) = 1$ indicating a jump at time t , and $dq(t) = 0$ otherwise, and the random jump size is determined by the $\kappa(t)$ process (which is only defined for $dq(t) = 1$). From the discussion in Section 3.1.2 and Eq. (3.8) along with Proposition 5, the realized volatility is then consistent for the integrated volatility plus the squared jumps,

$$\text{plim}_{n \rightarrow \infty} v^2(t, h; n) = v^2(t, h) = \int_{t-h}^t \sigma^2(\tau) d\tau + \sum_{t-h \leq \tau \leq t} \kappa^2(\tau) \cdot dq(\tau), \quad 0 < h \leq t \leq T. \quad (4.25)$$

One may further show (e.g., Aït-Sahalia, 2004; Barndorff-Nielsen and Shephard, 2003) that even in the presence of jumps, but for $s < 2$, the s th-order realized power variation is unaffected by jumps and remains consistent for the notional s th-order power variation, as defined above,

$$\text{plim}_{n \rightarrow \infty} v^{[s]}(t, h; n) = v^{[s]}(t, h) \equiv \int_{t-h}^t \sigma^s(\tau) d\tau, \quad 0 < h \leq t \leq T. \quad (4.26)$$

Hence, by summing high-frequency absolute returns raised to powers less than two, it is possible to mitigate the impact of the discontinuous jump component in the volatility measurement. Related realized power variation measures have also recently been explored empirically in a series of paper by Ghysels et al. (2004, 2006) in the form of so-called MIDAS, or mixed-data-sample, regressions.

More general so-called multipower variation measures have also recently been analyzed in the literature, with the following definition adapted from Barndorff-Nielsen and Shephard (2006a).

Definition 8 *Multipower Variation Measures*

The realized multipower variation of order $\{s_1, s_2, \dots, s_j\}$ over $[t - h, t]$, for $0 < h \leq t \leq T$, is defined by

$$v^{[s_1, s_2, \dots, s_j]}(t, h; n) \equiv \mu_{s_1}^{-1} \cdot \dots \cdot \mu_{s_j}^{-1} \cdot (h/n)^{1-sm/2} \sum_{i=-j, \dots, n} |r(t - h + i \cdot (h/n), h/n)|^{s_1} \cdot |r(t - h + (i - 1) \cdot (h/n), h/n)|^{s_2} \cdot \dots \cdot |r(t - h + (i - j + 1) \cdot (h/n), h/n)|^{s_j}, \quad (4.27)$$

where $s_1 \geq 0, \dots, s_j \geq 0$, $sm \equiv s_1 + \dots + s_j$, $\mu_s = E(|Z|^s)$, and Z denotes a standard normal distribution.

This definition obviously includes the standard realized volatility,

$$v^{[2, 0, \dots, 0]}(t, h; n) = v^2(t, h; n),$$

and the s th-order power variation measure,

$$v^{[s, 0, \dots, 0]}(t, h; n) = v^{[s]}(t, h; n),$$

as special cases. However, in contrast to the realized volatility and power variation measures, which are based on the summation of power transforms of the absolute returns, the more general multipower variation measures are constructed by the summation of the product of *sequential* transformed absolute returns. Just like the power variation measures may be rendered robust to jumps by considering $s < 2$, the realized multipower variation measures may similarly be insulated from the impact of jumps by appropriately

choosing the different orders of the power transforms. The following proposition, due Barndorff-Nielsen and Shephard (2006a) and Barndorff-Nielsen et al. (2006), formally justifies this idea.

Proposition 9 *Consistency of Realized Multipower Variation*

The realized multipower variation for the jump-diffusion in Eq. (4.24) provides a consistent estimate for the corresponding integrated variation,

$$\text{plim}_{n \rightarrow \infty} v^{[s_1, s_2, \dots, s_j]}(t, h; n) = \int_{t-h}^t \sigma^{sm}(\tau) d\tau, \quad 0 < h \leq t \leq T, \quad (4.28)$$

where $sm \equiv s_1 + \dots + s_j$, and $s_1 < 2, \dots, s_j < 2$

Thus, the sum of the powers in the multipower variation measure directly dictates the specific limiting integrated variation measure. In particular, the so-called bipower variation measure, constructed by the summation of adjacent absolute returns, consistently (for $n \rightarrow \infty$) estimates conventional integrated volatility,

$$\text{plim}_{n \rightarrow \infty} v^{[1,1]}(t, h; n) = \int_{t-h}^t \sigma^2(\tau) d\tau. \quad (4.29)$$

Combining this result with the consistency of the realized volatility in Eq. (4.25), the difference between the two measures affords a relatively simple-to-implement consistent estimate of the squared jumps that occurred over the $[t - h, t]$ time interval,

$$\text{plim}_{n \rightarrow \infty} [v^2(t, h; n) - v^{[1,1]}(t, h; n)] = \sum_{t-h \leq \tau \leq t} \kappa^2(\tau) \cdot dq(\tau). \quad (4.30)$$

Moreover, Barndorff-Nielsen and Shephard (2006a) have shown that for the continuous sample path diffusion in Eq. (4.14), or $q(t) \equiv 0$ in Eq. (4.24),

$$(h/n)^{-1/2} \cdot [v^2(t, h; n) - v^{[1,1]}(t, h; n)] \cdot [(\mu_1^{-4} + 2\mu_1^{-2} - 5) \cdot v^{[1,1,1,1]}(t, h; n)]^{-1/2} \rightarrow N(0, 1), \quad (4.31)$$

for $n \rightarrow \infty$. This, therefore, allows for the construction of high-frequency-based nonparametric tests for the existence jumps. Further, theoretical refinements and actual empirical applications involving this test have recently been pursued by Andersen et al. (2007a) and Huang and Tauchen (2005), among others.

Research in the realized volatility area has evolved rapidly over the past few years, and it is still too early to draw firm conclusions or consensus opinion about the preferred procedures. However, the theoretical and empirical results reported to date have been very promising. Recent research into multivariate extensions of the different variation

measures and propositions discussed above should also help in further establishing a firm theoretical foundation for corresponding new realized covariation measures, CAPM betas, and factor loadings. The formulation of feasible co-jump measures and test statistics present another theoretically challenging set of problems. Of course, as discussed repeatedly, the development of reliable empirical procedures for dealing with the inherent market microstructure frictions at the highest possible sampling frequencies, both univariate and multivariate, and across different assets and market mechanisms, remains of the utmost importance from a practical perspective.

5. DIRECTIONS FOR FUTURE RESEARCH

In the last 10 years, there has been a movement toward the use of newly available high-frequency asset return data, and away from restrictive and hard-to-estimate parametric models toward flexible and computationally simple nonparametric approaches. Those trends will continue. Two related directions for future research are apparent: (i) continued development of methods for exploiting the volatility information in high-frequency data, and (ii) volatility modeling and forecasting in the high-dimensional multivariate environments of practical financial economic relevance. The realized volatility concept tackles both: it incorporates the highly useful information in high-frequency data while dispensing with the need to actually model the high-frequency data, and it requires only the most trivial of computations, thereby bringing within reach the elusive goal of accurate and high-dimensional volatility measurement, modeling, and forecasting. We look forward to realization of that goal in the foreseeable future.

ACKNOWLEDGMENTS

We acknowledge an anonymous referee, Lars Hansen, Nour Meddahi, and Neil Shephard for providing us with exceptionally helpful comments on the initial draft of the chapter. The work of Andersen and Bollerslev is supported by a grant from the NSF to NBER and CREATES funded by the Danish Research Foundation, and Diebold's work is funded by a grant from the NSF to the University of Pennsylvania.

REFERENCES

- Anh, V.V., C.C. Heyde and N.N. Leonenko (2002), "Dynamic Models of Long-Memory Processes Driven by Lévy Noise," *Journal of Applied Probability*, 39, 730–747.
- Aït-Sahalia, Y. (2002), "Telling from Discrete Data Whether the Underlying Continuous-Time Model is a Diffusion," *Journal of Finance*, 57, 2075–2121.
- Aït-Sahalia, Y. (2004), "Disentangling Volatility from Jumps," *Journal of Financial Economics*, 74, 487–528.

- Aït-Sahalia (2007), "Estimating Continuous-Time Models with Discretely Sampled Data," in R. Blundell, P. Torsten and W.K. Newey (eds.), *Advances in Economics and Econometrics, Ninth World Congress*, 261–327. Cambridge, UK: Cambridge University Press.
- Aït-Sahalia, Y., P.A. Mykland and L. Zhang (2005), "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise," *Review of Financial Studies*, 18, 351–416.
- Aït-Sahalia, Y., P.A. Mykland and L. Zhang (2006), "Ultra High Frequency Volatility Estimation with Dependent Microstructure Noise," Manuscript, Department of Economics, Princeton University.
- Aït-Sahalia, Y., L.P. Hansen and J. Scheinkman (2010), "Operator Methods for Continuous-Time Markov Processes," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Akgiray, V. and G. Booth (1986), "Stock Price Processes with Discontinuous Time Paths: An Empirical Investigation," *The Financial Review*, 21, 163–184.
- Alizadeh, S., M.W. Brandt and F.X. Diebold (2002), "Range-Based Estimation of Stochastic Volatility Models," *Journal of Finance*, 57, 1047–1092.
- Andersen, T.G. (1992), "Volatility," Manuscript, Kellogg Graduate School of Management, Northwestern University.
- Andersen, T.G. (1994), "Stochastic Autoregressive Volatility: A Framework for Volatility Modeling," *Mathematical Finance*, 4, 75–102.
- Andersen, T.G. (1996), "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility," *Journal of Finance*, 51, 169–204.
- Andersen, T.G., L. Benzoni and J. Lund (2002), "An Empirical Investigation of Continuous Time Equity Return Models," *Journal of Finance*, 57, 1239–1284.
- Andersen, T.G. and T. Bollerslev (1997), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," *Journal of Finance*, 52, 975–1005.
- Andersen, T.G. and T. Bollerslev (1998a), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885–905.
- Andersen, T.G. and T. Bollerslev (1998b), "Deutschemark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies," *Journal of Finance*, 53, 219–265.
- Andersen, T.G. and T. Bollerslev (1998c), "ARCH and GARCH Models," in S. Kotz, C.B. Read and D.L. Banks (eds.), *Encyclopedia of Statistical Sciences, Vol. II* (6–16). New York: John Wiley and Sons.
- Andersen, T.G., T. Bollerslev, P. Christoffersen and F.X. Diebold (2006a), "Volatility and Correlation Forecasting," in G. Elliott, C.W.J. Granger and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Chapter 15, 777–878. Amsterdam: North-Holland.
- Andersen, T.G., T. Bollerslev, P. Christoffersen and F.X. Diebold (2006b), "Practical Volatility and Correlation Modeling for Financial Market Risk Management," pp. 512–548 (with discussion); M. Carey and R. Stulz (eds.), *Risks of Financial Institutions*. Chicago: University of Chicago Press and National Bureau of Economic Research.
- Andersen, T.G., T. Bollerslev and F.X. Diebold (2007a), "Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility," *Review of Economics and Statistics*, 89, 701–720.

- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2000a), "Great Realizations," *Risk*, 13, 105–108.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2000b), "Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian," *Multinational Finance Journal*, 4, 159–179.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens (2001a), "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43–76.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001b), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42–55.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2003a), "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and C. Vega (2003b), "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review*, 93, 38–62.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and G. Wu (2005), "A Framework for Exploring the Macroeconomic Determinants of Systematic Risk," *American Economic Review*, 95, 398–404.
- Andersen, T.G., T. Bollerslev and D.P. Dobrev (2007b), "No-Arbitrage Semi-Martingale Restrictions for Continuous-Time Volatility Models subject to Leverage Effects, Jumps and i.i.d. Noise: Theory and Testable Distributional Implications," *Journal of Econometrics*, 138, 125–180.
- Andersen, T.G., T. Bollerslev, P.H. Frederiksen and M.Ø. Nielsen (2009), "Continuous-Time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns," *Journal of Applied Econometrics*, in press.
- Andersen, T.G., T. Bollerslev and S. Lange (1999), "Forecasting Financial Market Volatility: Sample Frequency vis-a-vis Forecast Horizon," *Journal of Empirical Finance*, 6, 457–477.
- Andersen, T.G., T. Bollerslev and N. Meddahi (2004), "Analytic Evaluation of Volatility Forecasts," *International Economic Review*, 45, 1079–1110.
- Andersen, T.G., T. Bollerslev and N. Meddahi (2005), "Correcting the Errors: Volatility Forecast Evaluation Using High-Frequency Data and Realized Volatilities," *Econometrica*, 73, 279–296.
- Andersen, T.G., H.J. Chung and B.E. Sørensen (1999), "Efficient Method of Moments Estimation of a Stochastic Volatility Model: A Monte Carlo Study," *Journal of Econometrics*, 91, 61–87.
- Andersen, T.G., and J. Lund (1997), "Estimating Continuous-Time Stochastic Volatility Models of the Short Term Interest Rate," *Journal of Econometrics*, 77, 343–377.
- Anderson, B.D.O. and J.B. Moore (1979), *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.
- Andreou, E. and E. Ghysels (2002), "Rolling Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results," *Journal of Business and Economic Statistics*, 20, 363–376.
- Andrews, D.W.K. (1991), "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817–858.
- Ané, T. and H. Geman (2000), "Order Flow, Transaction Clock, and Normality of Asset Returns," *Journal of Finance*, 55, 2259–2284.
- Areal, N.M.P.C. and S.J. Taylor (2002), "The Realized Volatility of FTSE-100 Futures Prices," *Journal of Futures Markets*, 22, 627–648.
- Back, K. (1991), "Asset Prices for General Processes," *Journal of Mathematical Economics*, 20, 371–395.
- Bai, X., J.R. Russell and G.C. Tiao (2004), "Effects on Non-Normality and Dependence on the Precision of Variance Estimates using High-Frequency Data," Manuscript, Graduate School of Business, University of Chicago.

- Baillie, R.T., T. Bollerslev and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 74, 3–30.
- Bakshi, G., C. Cao and Z. Chen (1997), "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, 2003–2049.
- Ball, C.A. and W.N. Torous (1985), "On Jumps in Common Stock Prices and Their Impact on Call Option Pricing," *Journal of Finance*, 40, 155–173.
- Bandi, F. and P.C.B. Phillips (2010), "Nonstationary Continuous-Time Processes," in Y. Ait-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Bandi, F. and J.R. Russell (2006a), "Separating Microstructure Noise from Volatility," *Journal of Financial Economics*, 79, 655–692.
- Bandi, F. and J.R. Russell (2006b), "Volatility," in J.R. Birge and V. Linetsky (eds.), *Handbook of Financial Engineering*. Amsterdam: North-Holland, inpress.
- Bandi, F. and J.R. Russell (2008), "Microstructure Noise, Realized Volatility, and Optimal Sampling," *Review of Economic Studies*, 75, 339–369.
- Bandi, F., J.R. Russell and Y. Zhou (2006), "Using High-Frequency Data in Dynamic Portfolio Choice," *Econometric Reviews*, inpress.
- Banon, G. (1978), "Nonparametric Identification for Diffusion Processes," *SIAM Journal of Optimization and Control*, 16, 380–395.
- Barndorff-Nielsen, O.E. and N. Shephard (2001), "Non-Gaussian Ornstein-Uhlenbeck-Based Models and Some of Their Uses in Financial Economics," *Journal of the Royal Statistical Society, Series B*, 63, 167–241.
- Barndorff-Nielsen, O.E. and N. Shephard (2002a), "Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society, Series B*, 64, 253–280.
- Barndorff-Nielsen, O.E. and N. Shephard (2002b), "Estimating Quadratic Variation Using Realised Variance," *Journal of Applied Econometrics*, 17, 457–477.
- Barndorff-Nielsen, O.E. and N. Shephard (2003), "Realized Power Variation and Stochastic Volatility Models," *Bernoulli*, 9, 243–265.
- Barndorff-Nielsen, O.E. and N. Shephard (2004a), "Econometric Analysis of Realised Covariation: High-Frequency Covariance, Regression and Correlation in Financial Econometrics," *Econometrica*, 72, 885–925.
- Barndorff-Nielsen, O.E. and N. Shephard (2004b), "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics*, 2, 1–37.
- Barndorff-Nielsen, O.E. and N. Shephard (2005), "How Accurate is the Asymptotic Approximation to the Distribution of Realized Volatility?" in D.W.F. Andrews, J.L. Powell, P.A. Ruud and J.H. Stock (eds.), *Identification and Inference for Econometric Models. A Festschrift in Honour of T.J. Rothenberg*. Cambridge, UK: Cambridge University Press.
- Barndorff-Nielsen, O.E. and N. Shephard (2006a), "Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation," *Journal of Financial Econometrics*, 4, 1–30.
- Barndorff-Nielsen, O.E. and N. Shephard (2006b), "Variation, Jumps and High-Frequency Data in Financial Econometrics," in R. Blundell, P. Torsten and W.K. Newey (eds.), *Advances in Economics and Econometrics, Ninth World Congress*. Cambridge, UK: Cambridge University Press.
- Barndorff-Nielsen, O.E., S.E. Graversen, J. Jacod and N. Shephard (2006), "Limit Theorems for Realised Bipower Variation in Econometrics," *Econometric Theory*, 22, 677–719.

- Barndorff-Nielsen, O.E., P.R. Hansen, A. Lunde and N. Shephard (2008), "Designing Realised Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise," *Econometrica*, 76, 1481–1536.
- Barucci, E. and R. Renò (2002), "On Measuring Volatility and the GARCH Forecasting Performance," *Journal of International Financial Markets, Institutions & Money*, 12, 183–200.
- Bates, D. (1996a), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options," *Review of Financial Studies*, 9, 69–107.
- Bates (1996b), "Testing Option Pricing Models," in G.S. Maddala (ed.), *Handbook of Statistics Vol. 14, Statistical Methods in Finance*, 567–611. Amsterdam: North-Holland.
- Bekaert, G. and G. Wu (2000), "Asymmetric Volatility and Risk in Equity Markets," *Review of Financial Studies*, 13, 1–42.
- Bibby, B.M., M. Jacobsen and M. Sørensen (2010), "Estimating Functions for Discretely Sampled Diffusion-Type Models," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Black, F. (1976), "Studies of Stock Market Volatility Changes," *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.
- Bollen, B. and B. Inder (2002), "Estimating Daily Volatility in Financial Markets Utilizing Intraday Data," *Journal of Empirical Finance*, 9, 551–562.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.
- Bollerslev, T. (1990), "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Approach," *Review of Economic and Statistics*, 72, 498–505.
- Bollerslev, T., R.Y. Chou and K.F. Kroner (1992), "ARCH Modeling in Finance: A Selective Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5–59.
- Bollerslev, T., R.F. Engle and D.B. Nelson (1994), "ARCH Models," in R.F. Engle and D. McFadden (eds.), *Handbook of Econometrics, Volume IV*, 2959–3038. Amsterdam: North-Holland.
- Bollerslev, T., R.F. Engle and J.M. Wooldridge (1988), "A Capital Asset Pricing Model with Time Varying Covariances," *Journal of Political Economy*, 96, 116–131.
- Bollerslev, T., M. Gibson and H. Zhou (2005), "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," Manuscript, Department of Economics, Duke University.
- Bollerslev, T. and H.O. Mikkelsen (1996), "Modeling and Pricing Long Memory in Stock Market Volatility," *Journal of Econometrics*, 73, 151–184.
- Bollerslev, T. and J.M. Wooldridge (1992), "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances," *Econometric Reviews*, 11, 143–172.
- Bollerslev, T. and B.Y.B. Zhang (2003), "Measuring and Modeling Systematic Risk in Factor Pricing Models Using High-Frequency Data," *Journal of Empirical Finance*, 10, 533–558.
- Bollerslev, T. and H. Zhou (2002), "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics*, 109, 33–65.
- Bollerslev, T. and H. Zhou (2006), "Volatility Puzzles: A Simple Framework for Gauging Return-Volatility Regressions," *Journal of Econometrics*, 131, 123–150.
- Bougerol, P. and N. Picard (1992), "Stationarity of GARCH Processes and Some Non-Negative Time Series," *Journal of Econometrics*, 52, 115–128.

- Brandt, M.W. and F.X. Diebold (2006), "A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations," *Journal of Business*, 79, 61–73.
- Breidt, F.J., N. Crato and P. de Lima (1998), "On the Detection and Estimation of Long Memory in Stochastic Volatility," *Journal of Econometrics*, 73, 325–334.
- Britten-Jones, M. and A. Neuberger (2000), "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839–866.
- Brockwell, P. and T. Marquardt (2005), "Lévy-Driven and Fractional Integrated ARMA Processes with Continuous Time Parameters," *Statistica Sinica*, 15, 477–494.
- Campbell, J.Y. and L. Hentschel (1992), "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns," *Journal of Financial Economics*, 31, 281–318.
- Canina, L. and S. Figlewski (1993), "The Information Content of Implied Volatility," *Review of Financial Studies*, 6, 659–681.
- Carr, P. and D. Madan (1998), "Towards a Theory of Volatility Trading," in R. Jarrow (ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*, 417–428. London: Risk Books.
- Carr, P. and L. Wu (2006), "A Tale of Two Indices," *Journal of Derivatives*, Spring, 13–29.
- Carr, P. and L. Wu (2009), "Variance Risk Premia," *Review of Financial Studies*, 22, 1311–1342.
- Chan, K.C., G.A. Karolyi, F.A. Longstaff and A. Sanders (1992), "An Empirical Comparison of Alternative Models of the Short-Term Interest Rate," *Journal of Finance*, 47, 1209–1227.
- Chen, X., L.P. Hansen and M. Carrasco (2003), "Nonlinearity and Temporal Dependence," Manuscript, Department of Economics, University of Chicago.
- Chernov, M., A.R. Gallant, E. Ghysels and G.E. Tauchen (2003), "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, 116, 225–257.
- Chernov, M. and E. Ghysels (2000), "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation," *Journal of Financial Economics*, 57, 407–458.
- Chib, S., F. Nardari and N. Shephard (2006), "Analysis of High Dimensional Multivariate Stochastic Volatility Models," *Journal of Econometrics*, 134, 341–371.
- Christensen, K. and M. Podolskij (2007), "Realized Range-Based Estimation of Integrated Variance," *Journal of Econometrics*, 141, 323–349.
- Christensen, B.J. and N.R. Prabhala (1998), "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 37, 125–150.
- Chung, K.L. and R.J. Williams (1983), *Introduction to Stochastic Integration*. Boston: Birkhäuser, Inc.
- Clark, P.K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41, 135–155.
- Comte, F. and E. Renault (1996), "Long Memory Continuous Time Models," *Journal of Econometrics*, 73, 101–149.
- Comte, F. and E. Renault (1998), "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291–323.
- Corsi, F. (2003), "A Simple Long-Memory Model of Realized Volatility," Manuscript, University of Southern Switzerland.
- Corsi, F., G. Zumbach, U. Müller and M.M. Dacorogna (2001), "Consistent High-Precision Volatility from High-Frequency Data," *Economic Notes*, 30, 183–204.
- Cox, J.C., J.E. Ingersoll, Jr. and S.A. Ross (1985), "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–407.

- Cox, J.C. and S.A. Ross (1976), "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 7, 229–263.
- Dacorogna, M.M., R. Gencay, U. Müller, R.B. Olsen and O.V. Pictet (2001). *An Introduction to High-Frequency Finance*. San Diego, CA: Academic Press.
- Dai, Q. and K.J. Singleton (2000), "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943–1978.
- Danielsson, J. (1994), "Stochastic Volatility in Asset Prices: Estimation with Simulated Maximum Likelihood," *Journal of Econometrics*, 61, 375–400.
- Davidian, M. and R.J. Carroll (1987), "Variance Function Estimation," *Journal of the American Statistical Association*, 82, 1079–1091.
- Day, T.E. and C.M. Lewis (1992), "Stock Market Volatility and the Information Content of Stock Index Options," *Journal of Econometrics*, 52, 267–287.
- Demeterfi, K., E. Derman, M. Kamal and J. Zou (1999), "A Guide to Volatility and Variance Swaps," *Journal of Derivatives*, 6, 9–32.
- Deo, R., C. Hurvich and Y. Lu (2006), "Forecasting Realized Volatility Using a Long-Memory Stochastic Volatility Model: Estimation, Prediction and Seasonal Adjustment," *Journal of Econometrics*, 131, 29–58.
- Diebold, F.X. and A. Inoue (2001), "Long Memory and Structural Change," *Journal of Econometrics*, 105, 131–159.
- Diebold, F.X. and J.A. Lopez (1995), "Modeling Volatility Dynamics," in K.V. Hoover (ed.), *Macroeconometrics: Developments, Tensions and Prospects*, 427–472. Boston: Kluwer Academic Press.
- Diebold, F.X. and M. Nerlove (1989), "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1–22.
- Ding, Z., C.W.J. Granger and R.F. Engle (1993), "A Long Memory Property of Stock Market Returns and a New Model," *Journal of Empirical Finance*, 1, 83–106.
- Dohnal, G. (1987), "On Estimating the Diffusion Coefficient," *Journal of Applied Probability*, 24, 105–114.
- Dobrev, D.P. (2007), "Capturing Volatility from Large Price Moves: Generalized Range Theory and Applications," Manuscript, Department of Finance, Kellogg School, Northwestern University.
- Drost, F.C. and T.E. Nijman (1993), "Temporal Aggregation of GARCH Processes," *Econometrica*, 61, 909–927.
- Drost, F.C., T.E. Nijman and B.J.M. Werker (1998), "Estimation and Testing in Models Containing Both Jumps and Conditional Heteroskedasticity," *Journal of Business and Economic Statistics*, 16, 237–243.
- Drost, F.C. and B.J.M. Werker (1996), "Closing the GARCH Gap: Continuous Time GARCH Modeling," *Journal of Econometrics*, 74, 31–58.
- Duan, J.C. (1995), "The GARCH Option Pricing Model," *Mathematical Finance*, 5, 13–32.
- Duan, J.C. (1997), "Augmented GARCH(p,q) Process and Its Diffusion Limit," *Journal of Econometrics*, 79, 97–127.
- Duffie, D. and K.J. Singleton (1993), "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–952.
- Duffie, D., J. Pan and K. Singleton (2000), "Transform Analysis and Asset Pricing for Affine Jump Diffusions," *Econometrica*, 68, 1343–1376.

- Dumas, B., J. Fleming and R.E. Whaley (1998), "Implied Volatility Functions: Empirical Tests," *Journal of Finance*, 53, 2059–2106.
- Dybvig, P.H. (1993), "Exploration of Interest Rate Data," Manuscript, Department of Finance, Olin School of Business, Washington University.
- Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 50, 987–1008.
- Engle, R.F. (1995). *ARCH: Selected Readings*. Oxford, UK: Oxford University Press.
- Engle, R.F. (2002), "Dynamic Conditional Correlations: A Simple Class of Multivariate GARCH Models," *Journal of Business and Economic Statistics*, 20, 339–350.
- Engle, R.F. (2004), "Risk and Volatility: Econometric Models and Financial Practice," *American Economic Review*, 94, 405–420.
- Engle, R.F. and T. Bollerslev (1986), "Modeling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1–50.
- Engle, R.F. and G.M. Gallo (2006), "A Multiple Indicators Model for Volatility Using Intra-Daily Data," *Journal of Econometrics*, 131, 3–27.
- Engle, R.F. and K. Kroner (1995), "Multivariate Simultaneous Generalized ARCH," *Econometric Theory*, 11, 122–150.
- Engle, R.F. and G.G.J. Lee (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*, 475–497. Oxford, UK: Oxford University Press.
- Engle, R.F., D.M. Lilien and R.P. Robins (1987), "Estimating Time-Varying Risk Premia in the Term Structure: The ARCH-M Model," *Econometrica*, 58, 525–542.
- Engle, R.F. and A.J. Patton (2001), "What Good is a Volatility Model?" *Quantitative Finance*, 1, 237–245.
- Epps, T.W. and M.L. Epps (1976), "The Stochastic Dependence of Security Prices Changes and Transaction Volumes: Implications for the Mixture-of-Distributions Hypothesis," *Econometrica*, 44, 305–321.
- Eraker, B. (2004), "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," *Journal of Finance*, 59, 1367–1403.
- Eraker, B., M. Johannes and N. Polson (2003), "The Impact of Jumps in Equity Index Volatility and Returns," *Journal of Finance*, 58, 1269–1300.
- Fleming, J. (1998), "The Quality of Market Volatility Forecasts Implied by S&P100 Index Options Prices," *Journal of Empirical Finance*, 5, 317–345.
- Fleming, J., C. Kirby and B. Ostdiek (2003), "The Economic Value of Volatility Timing Using Realized Volatility," *Journal of Financial Economics*, 67, 473–509.
- Fleming, J., B. Ostdiek and R.E. Whaley (1995), "Predicting Stock Market Volatility: A New Measure," *Journal of Futures Market*, 15, 265–302.
- Florens-Zmirou, D. (1993), "On Estimating the Diffusion Coefficient from Discrete Observations," *Journal of Applied Probability*, 30, 790–804.
- Fornari, F. and A. Mele (2001), "Recovering the Probability Density Function of Asset Prices Using GARCH as Diffusion Approximations," *Journal of Empirical Finance*, 8, 83–109.
- Foster, D. and D.B. Nelson (1996), "Continuous Record Asymptotics for Rolling Sample Estimators," *Econometrica*, 64, 139–174.

- French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3–29.
- Fridman, M. and L. Harris (1998), "A Maximum Likelihood Approach for Non-Gaussian Stochastic Volatility Models," *Journal of Business and Economic Statistics*, 16, 284–291.
- Gallant, A.R., C.T. Hsu and G.E. Tauchen (1999), "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance," *Review of Economics and Statistics*, 81, 617–631.
- Gallant, A.R. and G.E. Tauchen (1997), "Estimation of Continuous-Time Models for Stock Returns and Interest Rates," *Macroeconomic Dynamics*, 1, 135–168.
- Gallant, A.R. and G.E. Tauchen (1998), "Reprojecting Partially Observed Systems with Applications to Interest Rate Diffusions," *Journal of the American Statistical Association*, 93, 10–24.
- Gallant, A.R. and G.E. Tauchen (2000), "SNP, A Program for Nonparametric Time Series Analysis: A Users Guide," available at <ftp.econ.duke.edu>, folder pub/arg/snp.
- Gallant, A.R. and G.E. Tauchen (2010), "Simulated Score Methods and Indirect Inference for Continuous-Time Models," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Garcia, R., M.A. Lewis and E. Renault (2001), "Estimation of Objective and Risk-Neutral Distributions Based on Moments of Integrated Volatility," Manuscript, Department of Economics, University of Montreal and CIRANO.
- Garcia, R., E. Ghysels and E. Renault (2010), "The Econometrics of Option Pricing," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Garman, M.B. and M.J. Klass (1980), "On the Estimation of Security Price Volatilities from Historical Data," *Journal of Business*, 53, 67–78.
- Genon-Catalot, V., C. Laredo and D. Picard (1992), "Non-parametric Estimation of the Diffusion Coefficient by Wavelets Methods," *Scandinavian Journal of Statistics*, 19, 317–335.
- Geweke, J. (1989), "Exact Predictive Densities for Linear Models with ARCH Disturbances," *Journal of Econometrics*, 40, 63–86.
- Ghysels, E., A. Harvey and E. Renault (1996), "Stochastic Volatility," in G.S. Maddala (ed.), *Handbook of Statistics Vol. 14, Statistical Methods in Finance*, 119–191. Amsterdam: North-Holland.
- Ghysels, E., P. Santa-Clara and R. Valkanov (2004), "The MIDAS Touch: Mixed Data Sampling Regressions," Manuscript, Department of Economics, University of North Carolina.
- Ghysels, E., P. Santa-Clara and R. Valkanov (2006), "Predicting Volatility: How to Get the Most out of Returns Data Sampled at Different Frequencies," *Journal of Econometrics*, 131, 59–95.
- Giraitis, L., P. Kokoszka and R. Leipus (2000), "Stationary ARCH Models: Dependence Structure and Central Limit Theorem," *Econometric Theory*, 16, 3–22.
- Giraitis, L., R. Leipus, P.M. Robinson and D. Surgailis (2004), "LARCH, Leverage and Long Memory," *Journal of Financial Econometrics*, 2, 177–210.
- Giraitis, L., R. Leipus and D. Surgailis (2005), "Recent Advances in ARCH Modeling," in A. Kirman and G. Teysiere (eds.), *Long-Memory in Economics*, 3–38. Berlin: Springer Verlag.
- Glosten, L.R., R. Jagannathan and D. Runkle (1993), "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779–1801.
- Granger, C.W.J. and Z. Ding (1995), "Some Properties of Absolute Returns, an Alternative Measure of Risk," *Annals d'Economie et de Statistique*, 40, 67–91.

- Hansen, L.P. (1982), "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.
- Hansen, P.R. and A. Lunde (2006), "Realized Variance and Market Microstructure Noise," *Journal of Business and Economic Statistics*, 24, 127–161.
- Harrison, J.M. and D. Kreps (1978), "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.
- Harvey, A.C. (1998), "Long Memory in Stochastic Volatility," in J. Knight and J. Satchell (eds.), *Forecasting Volatility in Financial Markets*, 307–320. London, UK: Butterworth-Heinemann.
- Harvey, A.C., E. Ruiz and N. Shephard (1994), "Multivariate Stochastic Variance Models," *Review of Economic Studies*, 61, 247–264.
- Hasbrouck, J. (1996), "Modeling Market Microstructure Time Series," in G.S. Maddala and C.R. Rao (eds.), *Handbook of Statistics, Vol. 14*, 647–692. Amsterdam: North-Holland.
- Heston, S.L. (1993), "A Closed Form Solution for Options with Stochastic Volatility, with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–343.
- Heston, S.L. and S. Nandi (2000), "A Closed-Form GARCH Option Valuation Model," *Review of Financial Studies*, 13, 585–625.
- Hoffmann-Jørgensen, J. (1994), *Probability with a View Toward Statistics*, Volume 1. Chapman and Hall Probability Series, New York: Chapman & Hall.
- Huang, X. and G. Tauchen (2005), "The Relative Contribution of Jumps to Total Price Variation," *Journal of Financial Econometrics*, 3, 456–499.
- Hull, J. and A. White (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 281–300.
- Hsieh, D.A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets," *Journal of Finance*, 46, 1839–1877.
- Morgan, J.P. (1997), *RiskMetrics, Technical Documents*, (Fourth Edition). New York.
- Jacod, J. (1994), "Limit of Random Measures Associated with the Increments of a Brownian Semimartingale," Manuscript, Laboratoire de Probabilités, Université Pierre et Marie Curie, Paris.
- Jacod, J. (2010), "Inference for Stochastic Processes," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.
- Jacod, J. and P. Protter (1998), "Asymptotic Error Distributions for the Euler Method for Stochastic Differential Equations," *Annals of Probability*, 26, 267–307.
- Jacquier, E., N.G. Polson and P.E. Rossi (1994), "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics*, 12, 371–389.
- Jacquier, E., N.G. Polson and P.E. Rossi (2004), "Bayesian Analysis of Stochastic Volatility Models with Fat Tails and Leverage Effect," *Journal of Econometrics*, 122, 185–212.
- Jarrow, R. and E. Rosenfeld (1984), "Jump Risk and the Intertemporal Asset Pricing Model," *Journal of Business*, 57, 337–351.
- Jiang, G. and Y. Tian (2005), "Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, 1305–1342.
- Johannes, M.S. (2004), "The Statistical and Economic Role of Jumps in Continuous-Time Interest Rate Models," *Journal of Finance*, 59, 227–260.
- Johannes, M. and N. Polson (2010), "MCMC Methods for Continuous-Time Financial Econometrics," in Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.

- Karatzas, I. and S.E. Shreve (1991), *Brownian Motion and Stochastic Calculus* (Second Edition). New York: Springer-Verlag.
- Kawakatsu, H. (2006), "Matrix Exponential GARCH," *Journal of Econometrics*, 134, 95–128.
- Kim, S., N. Shephard and S. Chib (1998), "Stochastic Volatility: Likelihood Inference and Comparisons with ARCH Models," *Review of Economic Studies*, 65, 361–393.
- King, M., E. Sentana and S. Wadhvani (1994), "Volatility and Links Between National Stock Markets," *Econometrica*, 62, 901–933.
- Kitagawa, G. (1987), "Non-Gaussian State Space Modeling of Nonstationary Time Series," *Journal of the American Statistical Association*, 82, 1032–1063.
- Koopman, S.J., B. Jungbacker and E. Hol (2005), "Forecasting Daily Variability of the S&P500 Stock Index Using Historical, Realised and Implied Volatility Measures," *Journal of Empirical Finance*, 12, 445–475.
- Latané, H. and R.J. Rendleman (1976), "Standard Deviations of Stock Price Ratios Implied in Option Prices," *Journal of Finance*, 31, 369–381.
- Ledoit, O., P. Santa-Clara and M. Wolf (2003), "Flexible Multivariate GARCH Modeling with an Application to International Stock Markets," *Review of Economics and Statistics*, 85, 735–747.
- Ling, S. and M. McAleer (2002), "Necessary and Sufficient Moment Conditions for the GARCH(r,s) and Asymmetric Power GARCH(r,s) Models," *Econometric Theory*, 18, 722–729.
- Liu, M. (2000), "Modeling Long Memory in Stock Market Volatility," *Journal of Econometrics*, 99, 139–171.
- Maheu, J.M. and T.H. McCurdy (2002), "Nonlinear Features of Realized FX Volatility," *Review of Economics and Statistics*, 84, 668–681.
- Maheswaran, S. and C.A. Sims (1993), "Empirical Implications of Arbitrage-Free Asset Markets," in P.C.B. Phillips (ed.), *Models, Methods and Applications of Econometrics*, 301–316. Oxford: Basil Blackwell.
- Martens, M. (2002), "Measuring and Forecasting S&P500 Index-Futures Volatility Using High-Frequency Data," *Journal of Futures Markets*, 22, 497–518.
- Martens, M. and D. van Dijk (2007), "Measuring Volatility with the Realized Range," *Journal of Econometrics*, 138, 112–133.
- McAleer, M. and M.C. Medeiros (2006), "Realized Volatility: A Review," Manuscript, Department of Economics, Pontifical Catholic University of Rio de Janeiro, Brazil.
- Meddahi, N. (2002), "A Theoretical Comparison Between Integrated and Realized Volatilities," *Journal of Applied Econometrics*, 17, 479–508.
- Meddahi, N. and E. Renault (1996), "Aggregation and Marginalization of GARCH and Stochastic Volatility Models," Manuscript, Department of Economics, University of Montreal and CIRANO.
- Meddahi, N. and E. Renault (2004), "Temporal Aggregation of Volatility Models," *Journal of Econometrics*, 119, 355–379.
- Mele, A. and F. Fornari (2000), *Stochastic Volatility in Financial Markets: Crossing the Bridge to Continuous Time*. Boston: Kluwer Academic Publishers.
- Merton, R.C. (1976), "Option Pricing when Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics*, 3, 125–144.
- Merton, R.C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.

- Merton, R.C. (1982), "On the Mathematics and Economics Assumptions of Continuous-Time Models," in W.F. Sharpe and C.M. Cootner (eds.), *Financial Economics: Essays in Honor of Paul Cootner*. Englewood Cliffs, NJ: Prentice-Hall Inc.
- Nelson, D.B. (1988), "The Time Series Behavior of Stock Market Volatility and Returns," Unpublished Ph.D. Dissertation, Department of Economics, MIT.
- Nelson, D.B. (1990a), "ARCH Models as Diffusion Approximations," *Journal of Econometrics*, 45, 7–39.
- Nelson, D.B. (1990b), "Stationarity and Persistence in the GARCH(1,1) Model," *Econometric Theory*, 6, 318–334.
- Nelson, D.B. (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347–370.
- Nelson, D.B. (1992), "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," *Journal of Econometrics*, 52, 61–90.
- Nelson, D.B. (1996a), "Asymptotic Filtering Theory for Multivariate ARCH Models," *Journal of Econometrics*, 71, 1–47.
- Nelson, D.B. (1996b), "Asymptotically Optimal Smoothing with ARCH Models," *Econometrica*, 64, 561–573.
- Nelson, D.B. and C.Q. Cao (1992), "Inequality Constraints in the Univariate GARCH Model," *Journal of Business and Economic Statistics*, 10, 229–235.
- Nelson, D.B. and D. Foster (1994), "Asymptotic Filtering Theory for Univariate ARCH Models," *Econometrica*, 62, 561–573.
- Nelson, D.B. and D. Foster (1995), "Filtering and Forecasting with Misspecified ARCH Models II: Making the Right Forecast with the Wrong Model," *Journal of Econometrics*, 67, 303–335.
- Nelson, D.B. and B. Schwartz (1992), "Filtering with ARCH Models: A Monte Carlo Investigation," *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 1–6.
- Newey, W.K. and K.D. West (1987), "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.
- Nicolato, E. and E. Venardos (2003), "Option Pricing in Stochastic Volatility Models of the Ornstein-Uhlenbeck type," *Mathematical Finance*, 13, 445–466.
- Nijman, T.H. and E. Sentana (1996), "Marginalization and Contemporaneous Aggregation of Multivariate GARCH Processes," *Journal of Econometrics*, 71, 71–86.
- Pagan, A. and G.W. Schwert (1990), "Alternative Models for Conditional Stock Volatility," *Journal of Econometrics*, 45, 267–290.
- Palandri, A. (2006), "Sequential Conditional Correlations: Inference and Evaluation," Manuscript, Department of Economics, Duke University.
- Pan, J. (2002), "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics*, 63, 3–50.
- Parkinson, M. (1980), "The Extreme Value Method for Estimating the Variance of the Rate of Return," *Journal of Business*, 53, 61–65.
- Pelletier, D. (2006), "Regime Switching for Dynamic Correlations," *Journal of Econometrics*, 131, 445–473.
- Piazzesi, M. (2010), "Affine Term Structure Models," in Y. Ait-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.

- Pitt, M.K. and N. Shephard (1999), "Filtering via Simulation: Auxiliary Particle Filter," *Journal of the American Statistical Association*, 94, 590–599.
- Poterba, J. and L. Summers (1986), "The Persistence of Volatility and Stock Market Fluctuations," *American Economic Review*, 76, 1124–1141.
- Press, J.S. (1967), "A Compound Events Model for Security Prices," *Journal of Business*, 40, 317–335.
- Protter, P. (1992), *Stochastic Integration and Differential Equations: A New Approach* (Second Edition). New York: Springer-Verlag.
- Renault, E. (1997), "Econometric Models of Option Pricing Errors," in D.M. Kreps and K.F. Wallis (eds.) *Advances in Economics and Econometrics: Theory and Applications*, 223–278. Cambridge, UK: Cambridge University Press.
- Robinson, P.M. (1991), "Testing for Strong Serial Correlation and Dynamic Conditional Heteroskedasticity in Multiple Regression," *Journal of Econometrics*, 47, 67–84.
- Robinson, P.M. (2001), "The Memory of Stochastic Volatility Models," *Journal of Econometrics*, 101, 195–218.
- Rossi, P.E. (1996), *Modelling Stock Market Volatility: Bridging the Gap to Continuous Time*. San Diego, CA: Academic Press.
- Sandmann, G. and S.J. Koopman (1998), "Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood," *Journal of Econometrics*, 87, 271–301.
- Scholes, M. and J. Williams (1977), "Estimating Beta from Non-Synchronous Data," *Journal of Financial Economics*, 5, 309–327.
- Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?," *Journal of Finance*, 44, 1115–1153.
- Schwert, G.W. (1990), "Stock Volatility and the Crash of '87," *Review of Financial Studies*, 3, 77–102.
- Shephard, N. (1996), "Statistical Aspects of ARCH and Stochastic Volatility Models," in D.R. Cox, D.V. Hinkley and O.E. Barndorff-Nielsen (eds.), *Time Series Models in Econometrics, Finance and Other Fields*, 1–67. London: Chapman & Hall.
- Stoll, H.R. (2000), "Frictions," *Journal of Finance*, 55, 1479–1514.
- Tauchen, G.E. and M. Pitts (1983), "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, 51, 485–505.
- Taylor, S.J. (1982), "Financial Returns Modeled by the Product of two Stochastic Processes – a Study of the Daily Sugar Prices 1961–75," in O.D. Anderson (ed.), *Time Series Analysis: Theory and Practice*, 203–226. Amsterdam: North-Holland.
- Taylor, S.J. (1986), *Modeling Financial Time Series*. Chichester, UK: John Wiley and Sons.
- Taylor, S.J. (1994), "Modeling Stochastic Volatility: A Review and Comparative Study," *Mathematical Finance*, 4, 183–204.
- Taylor, S.J. and X. Xu (1997), "The Incremental Volatility Information in One Million Foreign Exchange Quotations," *Journal of Empirical Finance*, 4, 317–340.
- Thomakos, D.D. and T. Wang (2003), "Realized Volatility in the Futures Markets," *Journal of Empirical Finance*, 10, 321–353.
- Tse, Y.K. and A.K.C. Tsui (2002), "A Multivariate GARCH Model with Time-Varying Correlations," *Journal of Business and Economic Statistics*, 20, 351–362.
- Whaley, R.E. (1993), "Derivatives on Market Volatility: Hedging Tools Long Overdue," *Journal of Derivatives*, 1, 71–84.

- Wiggins, J.B. (1992), "Estimating the Volatility of S&P500 Futures Prices Using the Extreme Value Method," *Journal of Futures Market*, 12, 265–273.
- Wong, E. (1964), "The Construction of a Class of Stationary Markoff Processes," in R. Bellman (ed.), *Sixteenth Symposium in Applied Mathematics – Stochastic Processes in Mathematical Physics and Engineering*, 264–276. Providence, RI: American Mathematical Society.
- Zhang, L., P.A. Mykland and Y. Ait-Sahalia (2005), "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-frequency Data," *Journal of the American Statistical Association*, 100, 1394–1411.
- Zhou, B. (1996), "High-Frequency Data and Volatility in Foreign Exchange Rates," *Journal of Business and Economic Statistics*, 14, 45–52.
- Zumbach, G. (2004), "Volatility Processes and Volatility Forecasts with Long Memory," *Quantitative Finance*, 4, 70–86.
- Zumbach, G., F. Corsi and A. Trapletti (2002), "Efficient Estimation of Volatility using High Frequency Data," Manuscript, Olsen & Associates, Zürich, Switzerland.