

Andersen, T.G., Bollerslev, T., Christoffersen, P.F., and Diebold, F.X. (2006),  
"Volatility and Correlation Forecasting,"  
in G. Elliott, C.W.J. Granger, and Allan Timmermann (eds.),  
Handbook of Economic Forecasting. Amsterdam: North-Holland, 778-878.

Chapter 15

## VOLATILITY AND CORRELATION FORECASTING\*

TORBEN G. ANDERSEN

*Kellogg School of Management, Northwestern University, Evanston, IL 60208, and NBER*  
*e-mail: t-andersen@kellogg.northwestern.edu*

TIM BOLLERSLEV

*Department of Economics, Duke University, Durham, NC 27708, and NBER*  
*e-mail: boller@econ.duke.edu*

PETER F. CHRISTOFFERSEN

*Faculty of Management, McGill University, Montreal, Quebec, H3A 1G5, and CIRANO*  
*e-mail: peter.christoffersen@mcgill.ca*

FRANCIS X. DIEBOLD

*Department of Economics, University of Pennsylvania, Philadelphia, PA 19104, and NBER*  
*e-mail: fdiebold@sas.upenn.edu*

### Contents

Abstract	779
Keywords	779
1. Introduction	780
1.1. Basic notation and notions of volatility	781
1.2. Final introductory remarks	786
2. Uses of volatility forecasts	786
2.1. Generic forecasting applications	787
2.1.1. Point forecasting	787
2.1.2. Interval forecasting	788
2.1.3. Probability forecasting including sign forecasting	788
2.1.4. Density forecasting	789
2.2. Financial applications	789
2.2.1. Risk management: Value-at-Risk (VaR) and Expected Shortfall (ES)	790
2.2.2. Covariance risk: Time-varying betas and conditional Sharpe ratios	792

\* The work of Andersen, Bollerslev and Diebold was supported by grants from the U.S. National Science Foundation, and the work of Christoffersen was supported by FQRSC, SSHRC and IFM2. The authors gratefully acknowledge extremely helpful comments from two anonymous referees. The usual disclaimer applies.

*Handbook of Economic Forecasting, Volume 1*

*Edited by Graham Elliott, Clive W.J. Granger and Allan Timmermann*

© 2006 Elsevier B.V. All rights reserved

DOI: 10.1016/S1574-0706(05)01015-3

1	2.2.3. Asset allocation with time-varying covariances	793	1
2	2.2.4. Option valuation with dynamic volatility	794	2
3	2.3. Volatility forecasting in fields outside finance	796	3
4	2.4. Further reading	797	4
5	<b>3. GARCH volatility</b>	798	5
6	3.1. Rolling regressions and RiskMetrics	798	6
7	3.2. GARCH(1, 1)	800	7
8	3.3. Asymmetries and “leverage” effects	803	8
9	3.4. Long memory and component structures	805	9
10	3.5. Parameter estimation	807	10
11	3.6. Fat tails and multi-period forecast distributions	809	11
12	3.7. Further reading	812	12
13	<b>4. Stochastic volatility</b>	814	13
14	4.1. Model specification	815	14
15	4.1.1. The mixture-of-distributions hypothesis	815	15
16	4.1.2. Continuous-time stochastic volatility models	818	16
17	4.1.3. Estimation and forecasting issues in SV models	820	17
18	4.2. Efficient method of simulated moments procedures for inference and forecasting	823	18
19	4.3. Markov Chain Monte Carlo (MCMC) procedures for inference and forecasting	826	19
20	4.4. Further reading	828	20
21	<b>5. Realized volatility</b>	830	21
22	5.1. The notion of realized volatility	830	22
23	5.2. Realized volatility modeling	834	23
24	5.3. Realized volatility forecasting	835	24
25	5.4. Further reading	837	25
26	<b>6. Multivariate volatility</b>	839	26
27	6.1. Exponential smoothing and RiskMetrics	840	27
28	6.2. Multivariate GARCH models	841	28
29	6.3. Multivariate GARCH estimation	843	29
30	6.4. Dynamic conditional correlations	845	30
31	6.5. Multivariate stochastic volatility and factor models	847	31
32	6.6. Realized covariances and correlations	849	32
33	6.7. Further reading	851	33
34	<b>7. Evaluating volatility forecasts</b>	853	34
35	7.1. Point forecast evaluation from general loss functions	854	35
36	7.2. Volatility forecast evaluation	855	36
37	7.3. Interval forecast and Value-at-Risk evaluation	859	37
38	7.4. Probability forecast evaluation and market timing tests	860	38
39	7.5. Density forecast evaluation	861	39
40	7.6. Further reading	863	40
41	<b>8. Concluding remarks</b>	864	41
42	Uncited references	865	42
43	References	865	43

**Abstract**

Volatility has been one of the most active and successful areas of research in time series econometrics and economic forecasting in recent decades. This chapter provides a selective survey of the most important theoretical developments and empirical insights to emerge from this burgeoning literature, with a distinct focus on forecasting applications. Volatility is inherently latent, and Section 1 begins with a brief intuitive account of various key volatility concepts. Section 2 then discusses a series of different economic situations in which volatility plays a crucial role, ranging from the use of volatility forecasts in portfolio allocation to density forecasting in risk management. Sections 3–5 present a variety of alternative procedures for univariate volatility modeling and forecasting based on the GARCH, stochastic volatility and realized volatility paradigms, respectively. Section 6 extends the discussion to the multivariate problem of forecasting conditional covariances and correlations, and Section 7 discusses volatility forecast evaluation methods in both univariate and multivariate cases. Section 8 concludes briefly.

**Keywords**

volatility modeling, covariance forecasting, GARCH, stochastic volatility, realized volatility

*JEL classification:* C1, C5, G1

## 1. Introduction

In everyday language, volatility refers to the fluctuations observed in some phenomenon over time. Within economics, it is used slightly more formally to describe, without a specific implied metric, the variability of the random (unforeseen) component of a time series. More precisely, or narrowly, in financial economics, volatility is often defined as the (instantaneous) standard deviation (or “sigma”) of the random Wiener-driven component in a continuous-time diffusion model. Expressions such as the “implied volatility” from option prices rely on this terminology. In this chapter, we use the term volatility in the looser descriptive sense, characteristic of economics and econometrics, rather than the precise notion often implied in finance. Nonetheless, much of our discussion will be motivated by the need for forecasting the volatility of financial asset return series.

Return volatility is, of course, central to financial economics. Indeed, as noted by Campbell, Lo and MacKinlay (1997):

“... what distinguishes financial economics is the central role that uncertainty plays in both financial theory and its empirical implementation... Indeed in the absence of uncertainty, the problems of financial economics reduce to exercises in basic microeconomics” (p. 3).

This departure of finance from standard microeconomics is even more striking once one recognizes that volatility is inherently unobserved, or latent, and evolves stochastically through time. Not only is there nontrivial uncertainty to deal with in financial markets, but the level of uncertainty is latent. This imbues financial decision making with complications rarely contemplated within models of optimizing behavior in other areas of economics.

Depending on the data availability as well as the intended use of the model estimates and associated forecasts, volatility models are cast either in discrete time or continuous time. It is clear, however, that the trading and pricing of securities in many of today’s liquid financial asset markets is evolving in a near continuous fashion throughout the trading day. As such, it is natural to think of the price and return series of financial assets as arising through discrete observations from an underlying continuous-time process. It is, however, in many situations useful – and indeed standard practice – to formulate the underlying model directly in discrete time, and we shall consider both approaches in the chapter. Formally, there is also no necessary contradiction between the two strategies, as it is always, in principle, possible to deduce the distributional implications for a price series observed only discretely from an underlying continuous-time model. At the same time, formulation and estimation of empirically realistic continuous-time models often presents formidable challenges. Thus, even though many of the popular discrete-time models in current use are not formally consistent with an underlying continuous-time price processes, they are typically much easier to deal with from an inferential perspective, and as such, discrete-time models and forecasting procedures remain the method of choice in most practical applications.

## 1.1. Basic notation and notions of volatility

We first introduce some notation that will allow us to formalize the discussion of the different models and volatility concepts considered throughout the chapter. As noted above, although it is often natural to think about the process being forecasted as evolving in continuous time, many of the key developments in volatility forecasting have been explicitly formulated in terms of models for discretely sampled observations. In the univariate case, with observations available at equally spaced discrete points in time, we shall refer to such a process as

$$y_t \equiv y(t), \quad t = 1, 2, \dots, \quad (1.1)$$

where  $y(t)$  in turn may be thought of as the underlying continuously evolving process. We shall assume throughout that the *conditional* second moments of the  $y_t$  process exist, and refer to the corresponding conditional mean and variance as

$$\mu_{t|t-1} = E[y_t | \mathcal{F}_{t-1}], \quad (1.2)$$

and

$$\sigma_{t|t-1}^2 = \text{Var}[y_t | \mathcal{F}_{t-1}] = E[(y_t - \mu_{t|t-1})^2 | \mathcal{F}_{t-1}], \quad (1.3)$$

respectively, where the information set,  $\mathcal{F}_{t-1}$ , is assumed to reflect all relevant information through time  $t - 1$ . Just as the conditional mean may differ from the unconditional mean by effectively incorporating the most recent information into the one-step-ahead forecasts,  $\mu_{t|t-1} \neq E(y_t)$ , so will the conditional variance in many applications in macroeconomics and finance,  $\sigma_{t|t-1}^2 \neq \text{Var}(y_t)$ . This difference between conditional and unconditional moments is, of course, what underlies the success of time series based forecasting procedures more generally. For notational simplicity we will focus our discussion on the univariate case, but many of the same ideas readily carry over to the multivariate case. In the case of vector processes, discussed in detail in Section 6, we shall use the notation  $Y_t$ , with the corresponding vector of conditional means denoted by  $M_{t|t-1}$ , and the conditional covariance matrix denote by  $\Omega_{t|t-1}$ .

As previously noted, most of the important developments and applications in volatility modeling and forecasting have come within financial economics. Thus, to help fix ideas, we focus on the case of return volatility modeling and forecasting in the remainder of this section. To facilitate subsequent discussions, it will sometimes prove convenient to refer to the corresponding “price” and “return” processes by the letters  $p$  and  $r$ , respectively. Specifically, let  $p(t)$  denote the logarithmic price of an asset. The return over the discrete interval  $[t - h, t]$ ,  $h > 0$ , is then given by

$$r(t, h) = p(t) - p(t - h). \quad (1.4)$$

When measuring the returns over one time unit,  $h = 1$ , indicating, say, daily returns, we will generally drop the second indicator, so that

$$r(t) \equiv r(t, 1) = p(t) - p(t - 1). \quad (1.5)$$

Also, for discrete-time models and procedures, we shall follow the convention set out above, indicating the timing of the returns by subscripts in lieu of parentheses,

$$r_t = p_t - p_{t-1}. \quad (1.6)$$

Similarly, we shall refer to the multivariate case involving vectors of returns by the upper case letter,  $R_t$ .

Consider the discretely sampled return process,  $r_t$ . This one-period return is readily decomposed into an expected conditional mean return and an innovation, where the latter may be expressed as a standardized white noise process scaled by the time-varying conditional volatility. Specifically, using the notation in Equations (1.2) and (1.3),

$$r_t = \mu_{t|t-1} + \varepsilon_t = \mu_{t|t-1} + \sigma_{t|t-1}z_t, \quad (1.7)$$

where  $z_t$  denotes a mean zero, variance one, serially uncorrelated disturbance (white noise) process. This is the decomposition and volatility concept underlying the popular, and empirically highly successful, ARCH and GARCH type models discussed in Section 3. One reason that this approach is very convenient and tractable is that – conditional on the null hypothesis that all relevant information is observed and the model correctly specified – the volatility is known, or predetermined, as of time  $t - 1$ .

The assumption that all relevant information is observed and used in the formation of conditional expectations in accordance with the true model is obviously strong, but has powerful and very convenient implications. In contrast, if some relevant information is not directly observable, then it is only possible to exploit a genuine subset of the full information set, say  $\mathfrak{F}_{t-1} \subset \mathcal{F}_{t-1}$ . Under this scenario, the “true” conditional variance will be unobservable, even under correct model specification, and the volatility process becomes genuinely latent,

$$E[(r_t - E[r_t | \mathfrak{F}_{t-1}])^2 | \mathfrak{F}_{t-1}] \neq \sigma_{t|t-1}^2 \equiv E[\varepsilon_t^2 | \mathcal{F}_{t-1}].$$

Treating the volatility process as latent effectively transforms the volatility estimation problem into a filtering problem in which the “true” volatility cannot be determined exactly, but only extracted with some degree of error. This general line of reasoning is relevant for our discussion of stochastic volatility models in Section 4, and for the relationship between continuous and discrete-time modeling and forecasting procedures.

For now, however, we proceed under the convenient assumption that we are dealing with correctly specified models and the associated full information sets, so that the conditional first and second moments are directly observable and well specified. In this situation, the one-period-ahead volatility defined in (1.3) provides an unbiased estimate of the subsequent squared return innovation. Consequently, model specification and forecast evaluation tests can be constructed by comparing the realization of the squared return innovations to the corresponding one-step-ahead forecasts,

$$\varepsilon_t^2 = \sigma_{t|t-1}^2 z_t^2 = \sigma_{t|t-1}^2 + \sigma_{t|t-1}^2 (z_t^2 - 1). \quad (1.8)$$

The second term on the right-hand side has mean zero, confirming the unbiasedness of the conditional variance. However, there is typically a large amount of noise in the one-period squared return innovations relative to the underlying volatility, as manifest by a large idiosyncratic error component governed by the variance of  $z_t^2$ . In fact, for daily or weekly return data, this variance term is an order of magnitude larger than the period-per-period variation in the volatility process. Hence, even if the conditional variance can be seen as the proper forecasts of the corresponding “realized volatility”, as given by the squared return innovation, the latter provides a poor ex-post indicator of the actual volatility over the period, and would consequently not provide a very reliable way of judging the quality of the forecasts. We return to this point below.

Before doing so, however, it is useful to think of the returns as arising from an underlying continuous-time process. In particular, suppose that this underlying model involves a continuous sample path for the (logarithmic) price process. The return process may then, under general assumptions, be written in standard stochastic differential equation (sde) form as

$$dp(t) = \mu(t) dt + \sigma(t) dW(t), \quad t \geq 0, \quad (1.9)$$

where  $\mu(t)$  denotes the drift,  $\sigma(t)$  refers to the point-in-time or spot volatility, and  $W(t)$  denotes a standard Brownian motion. We will be more specific regarding the additional properties of these processes later on in the chapter. Intuitively, over (infinitesimal) small time intervals,  $\Delta$ ,

$$r(t, \Delta) \equiv p(t) - p(t - \Delta) \simeq \mu(t - \Delta) \cdot \Delta + \sigma(t - \Delta) \Delta W(t),$$

where  $\Delta W(t) \equiv W(t) - W(t - \Delta) \sim N(0, \Delta)$ . Of course, for  $\Delta = 1$ , and constant drift,  $\mu(\tau) \equiv \mu_{t|t-1}$ , and volatility,  $\sigma(\tau) \equiv \sigma_{t|t-1}$ , for  $t - 1 < \tau \leq t$ , this reduces to the discrete-time return decomposition in (1.7) with the additional assumption that  $z_t$  is i.i.d.  $N(0, 1)$ . Importantly, however, the drift,  $\mu(t)$ , and instantaneous volatility,  $\sigma(t)$ , for the continuous-time model in (1.9) need not be constant over the  $[t - 1, t]$  time interval, resulting in the general expression for the one-period return,

$$r(t) = p(t) - p(t - 1) = \int_{t-1}^t \mu(s) ds + \int_{t-1}^t \sigma(s) dW(s). \quad (1.10)$$

The semblance between this representation and the previous one-period return for the discrete-time model in (1.7) is clear. The conditional mean and variance processes in the discrete formulation are replaced by the corresponding integrated (averaged) realizations of the (potentially stochastically time-varying) mean and variance process over the following period, with the return innovations driven by the continuously evolving standard Brownian motion. For full generality, the above continuous-time model can be extended with a jump process allowing for discontinuities in the price path, as discussed further in Section 4.

Intuitively, the volatility for the continuous-time process in (1.9) over  $[t - 1, t]$  is intimately related to the evolution of the diffusive coefficient,  $\sigma(t)$ , which is also known as the spot volatility. In fact, given the i.i.d. nature of the return innovations governed by

the Brownian motion process, the return variation should be related to the cumulative (integrated) spot variance. It is, indeed, possible to formalize this intuition: the conditional return variation is linked closely and – under certain conditions in an ex-post sense – equal to the so-called integrated variance (volatility),

$$IV(t) \equiv \int_{t-1}^t \sigma^2(s) ds. \quad (1.11)$$

We provide more in-depth discussion and justification for this integrated volatility measure and its relationship to the conditional return distribution in Section 4. It is, however, straightforward to motivate the association through the approximate discrete return process,  $r(t, \Delta)$ , introduced above. If the variation in the drift is an order of magnitude less than the variation in the volatility over the  $[t - 1, t]$  time interval – which holds empirically over daily or weekly horizons and is consistent with a no-arbitrage condition – it follows, for small (infinitesimal) time intervals,  $\Delta$ ,

$$\text{Var}(r(t) \mid \mathcal{F}_{t-1}) \simeq E \left[ \sum_{j=1}^{1/\Delta} \sigma^2(t - j/\Delta) \cdot \Delta \mid \mathcal{F}_{t-1} \right] \simeq E[IV(t) \mid \mathcal{F}_{t-1}].$$

Hence, the integrated variance measure corresponds closely to the conditional variance,  $\sigma_{t|t-1}^2$ , for discretely sampled returns. It represents the realized volatility over the same one-period-ahead forecast horizon, and it simply reflects the cumulative impact of the spot volatility process over the return horizon. In other words, integrated variances are ex-post realizations that are directly comparable to ex-ante volatility forecasts. Moreover, in contrast to the one-period-ahead squared return innovations, which, as discussed in the context of (1.8), are plagued by large idiosyncratic errors, the integrated volatility measure is not distorted by error. As such, it serves as an ideal theoretical ex-post benchmark for assessing the quality of ex-ante volatility forecasts.

To more clearly illustrate these differences between the various volatility concepts, Figure 1 graphs the simulations from a continuous-time stochastic volatility process. The simulated model is designed to induce temporal dependencies consistent with the popular, and empirically successful, discrete-time GARCH(1, 1) model discussed in Section 3.<sup>1</sup> The top left panel displays sample path realization of the spot volatility or variance,  $\sigma^2(t)$ , over the 2500 “days” in the simulated sample. The top panel on the right shows the corresponding “daily” integrated volatility or variance,  $IV(t)$ . The two bottom panels show the “optimal” one-step-ahead discrete-time GARCH(1, 1) forecasts,  $\sigma_{t|t-1}^2$ , along with the “daily” squared returns,  $r_t^2$ . A number of features in these displays are of interest. First, it is clear that even though the “daily” squared returns generally track

<sup>1</sup> The simulated continuous-time GARCH diffusion shown in Figure 1 is formally defined by  $dp(t) = \sigma(t) dW_1(t)$  and  $d\sigma^2(t) = 0.035[0.636 - \sigma^2(t)] dt + 0.144\sigma^2(t) dW_2(t)$ , where  $W_1(t)$  and  $W_2(t)$  denote two independent Brownian motions. The same model has previously been analyzed in Andersen and Bollerslev (1998a), Andersen, Bollerslev and Meddahi (2004, 2005), among others.



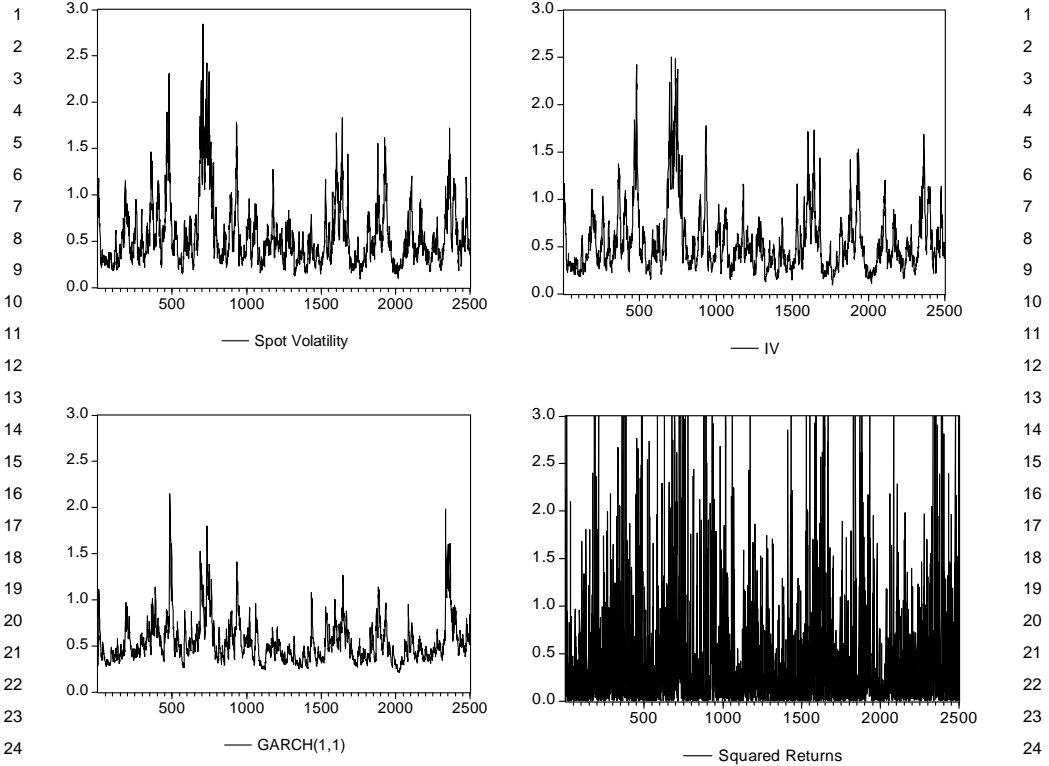


Figure 1. Different volatility concepts. We show the “daily” spot volatility,  $\sigma^2(t)$ , the integrated volatility,  $IV(t)$ , the discrete-time GARCH based volatility forecasts,  $\sigma_{t|t-1}^2$ , and the corresponding squared returns,  $r_T^2$ , from a simulated continuous-time GARCH diffusion.

the overall level of the volatility in the first two panels, as an unbiased measure should, it is an extremely noisy proxy. Hence, a naive assessment of the quality of the GARCH based forecasts in the third panel based on a comparison with the ex-post squared returns in panel four invariable will suggest very poor forecast quality, despite the fact that by construction the GARCH based procedure is the “optimal” discrete-time forecast. We provide a much more detailed discussion of this issue in Section 7 below. Second, the integrated volatility provides a mildly smoothed version of the spot volatility process. Since the simulated series has a very persistent volatility component the differences are minor, but still readily identifiable. Third, the “optimal” discrete-time GARCH forecasts largely appear as smoothed versions of the spot and integrated volatility series. This is natural as forecasts, by construction, should be less variable than the corresponding ex-post realizations. Fourth, it is also transparent, however, that the GARCH based forecasts fail to perfectly capture the nature of the integrated volatility series. The largest spike in volatility (around the 700–750 “day” marks) is systematically under-

1 estimated by the GARCH forecasts while the last spike (around the 2300–2350 “day” 1  
2 marks) is exaggerated relative to the actual realizations. This reflects the fact that the 2  
3 volatility is not constant over the “day”, and as such the (realized) integrated volatility 3  
4 is not equal to the (optimal) forecast from the discrete-time GARCH model which only 4  
5 utilizes the past “daily” return observations. Instead, there is a genuine random component 5  
6 to the volatility process as it evolves stochastically over the “trading day”. As 6  
7 a result, the “daily” return observations do not convey all relevant information and the 7  
8 GARCH model simply cannot produce fully efficient forecasts compared to what is the- 8  
9oretically possible given higher frequency “intraday” data. At the same time, in practice 9  
10 it is not feasible to produce exact real-time measures of the integrated, let alone the spot, 10  
11 volatility, as the processes are latent and we only have a limited and discretely sampled 11  
12 set of return observations available, even for the most liquid asset markets. As such, an 12  
13 important theme addressed in more detail in Sections 4 and 5 below involves the con- 13  
14struction of practical measures of ex-post realized volatility that mimic the properties of 14  
15the integrated volatility series. 15

### 16 *1.2. Final introductory remarks* 16

17  
18 This section has introduced some of the basic notation used in our subsequent discus- 18  
19sion of the various volatility forecasting procedures and evaluation methods. Our initial 19  
20account also emphasizes a few conceptual features and practical considerations. First, 20  
21volatility forecasts and measurements are generally restricted to (nontrivial) discrete- 21  
22time intervals, even if the underlying process may be thought of as evolving in con- 22  
23tinuous time. Second, differences between ARCH and stochastic volatility models may 23  
24be seen as direct consequences of assumptions about the observable information set. 24  
25Third, it is important to recognize the distinction between ex-ante forecasts and ex-post 25  
26realizations. Only under simplifying – and unrealistic – assumptions are the two iden- 26  
27tical. Fourth, standard ex-post measurements of realized volatility are often hampered 27  
28by large idiosyncratic components. The ideal measure is instead, in cases of general in- 28  
29terest, given by the so-called integrated volatility. The relationships among the various 29  
30concepts are clearly illustrated by the simulations in Figure 1. 30

31 The rest of the chapter unfolds as follows. Section 2 provides an initial motivating 31  
32discussion of several practical uses of volatility forecasts. Sections 3–5 present a variety 32  
33of alternative procedures for univariate volatility forecasting based on the GARCH, 33  
34stochastic volatility and realized volatility paradigms, respectively. Section 6 extends 34  
35the discussion to the multivariate problem of forecasting conditional covariances and 35  
36correlations, and Section 7 discusses practical volatility forecast evaluation techniques. 36  
37Section 8 concludes briefly. 37

## 38 **2. Uses of volatility forecasts** 38

39  
40 This section surveys how volatility forecasts are used in practical applications along 40  
41with applications in the academic literature. While the emphasis is on financial appli- 41  
42cations, 42  
43 43

1 cations the discussion is kept at a general level. Thus, we do not yet assume a specific 1  
 2 volatility forecasting model. The issues involved in specifying and estimating particular 2  
 3 volatility forecasting models will be discussed in subsequent sections. 3

4 We will first discuss a number of general statistical forecasting applications where 4  
 5 volatility dynamics are important. Then we will go into some detail on various applica- 5  
 6 tions in finance. Lastly we will briefly mention some applications in macroeconomics 6  
 7 and in other disciplines. 7

### 9 2.1. Generic forecasting applications 9

10 For concreteness, assume that the future realization of the variable of interest can be 10  
 11 written as a decomposition similar to the one already developed in Equation (1.7), 11  
 12

$$13 \quad y_{t+1} = \mu_{t+1|t} + \sigma_{t+1|t} z_{t+1}, \quad z_{t+1} \sim \text{i.i.d. } F, \quad (2.1) \quad 14$$

15 where  $\{y_{t+1}\}$  denotes a discrete-time real-valued univariate stochastic process, and  $F$  15  
 16 refers to the distribution of the zero-mean, unit-variance innovation,  $z_{t+1}$ . This represen- 16  
 17 tation is not entirely general as there could be higher-order conditional dependence in 17  
 18 the innovations. Such higher-moment dynamics would complicate some of the results, 18  
 19 but the qualitative insights would remain the same. Thus, to facilitate the presentation 19  
 20 we continue our discussion of the different forecast usages under slightly less than full 20  
 21 generality. 21

#### 22 2.1.1. Point forecasting 22

23 We begin by defining the forecast loss function which maps the ex-ante forecasts  $\hat{y}_{t+1|t}$  23  
 24 and the ex-post realization  $y_{t+1}$  into a loss value  $L(y_{t+1}, \hat{y}_{t+1|t})$ , which by assumption 24  
 25 increases with the discrepancy between the realization and the forecast. The exact form 25  
 26 of the loss function depends, of course, directly on the use of the forecast. However, in 26  
 27 many situations the loss function may reasonably be written in the form of an additive 27  
 28 error,  $e_{t+1} \equiv y_{t+1} - \hat{y}_{t+1|t}$ , as the argument, so that  $L(y_{t+1}, \hat{y}_{t+1|t}) = L(e_{t+1})$ . We will 28  
 29 refer to this as the forecast error loss function. 29  
 30

31 In particular, under the symmetric quadratic forecast error loss function, which is 31  
 32 implicitly used in many practical applications, the optimal point forecast is simply the 32  
 33 conditional mean of the process, regardless of the shape of the conditional distribution. 33  
 34 That is, 34  
 35

$$36 \quad \hat{y}_{t+1|t} \equiv \underset{\hat{y}}{\text{Arg min}} E[(y_{t+1} - \hat{y})^2 \mid \mathcal{F}_t] = \mu_{t+1|t}. \quad 36$$

37 Volatility forecasting is therefore irrelevant for calculating the optimal point forecast, 37  
 38 unless the conditional mean depends directly on the conditional volatility. However, 38  
 39 this exception is often the rule in finance, where the expected return generally involves 39  
 40 some function of the volatility of market wide risk factors. Of course, as discussed 40  
 41 further below, even if the conditional mean does not explicitly depend on the conditional 41  
 42 42  
 43

1 volatility, volatility dynamics are still relevant for assessing the uncertainty of the point  
2 forecasts.

3 In general, when allowing for asymmetric loss functions, the volatility forecast will  
4 be a key part of the optimal forecast. Consider for example the asymmetric linear loss  
5 function,

$$6 \quad L(e_{t+1}) = a|e_{t+1}|\mathbf{I}(e_{t+1} > 0) + b|e_{t+1}|\mathbf{I}(e_{t+1} \leq 0), \quad (2.2) \quad 7$$

8 where  $a, b > 0$ , and  $\mathbf{I}(\cdot)$  denotes the indicator function equal to zero or one depending  
9 on the validity of its argument. In this case positive and negative forecast errors have  
10 different weights ( $a$  and  $b$ , respectively) and thus different losses. Now the optimal  
11 forecast can be shown to be

$$12 \quad \hat{y}_{t+1|t} = \mu_{t+1|t} + \sigma_{t+1|t} F^{-1}(a/(a+b)), \quad (2.3) \quad 13$$

14 which obviously depends on the relative size of  $a$  and  $b$ . Importantly, the volatility  
15 plays a key role even in the absence of conditional mean dynamics. Only if  $F^{-1}(a/$   
16  $(a+b)) = 0$  does the optimal forecast equal the conditional mean.

17 This example is part of a general set of results in Granger (1969) who shows that  
18 if the conditional distribution is symmetric (so that  $F^{-1}(1/2) = 0$ ) and if the forecast  
19 error loss function is also symmetric (so that  $a/(a+b) = 1/2$ ) but not necessarily  
20 quadratic, then the conditional mean is the optimal point forecast.

### 21 2.1.2. Interval forecasting

22 Constructing accurate interval forecasts around the conditional mean forecast for in-  
23 flation was a leading application in Engle's (1982) seminal ARCH paper. An interval  
24 forecast consists of an upper and lower limit. One version of the interval forecast puts  
25  $p/2$  probability mass below and above the lower and upper limit, respectively. The in-  
26 terval forecast can then be written as

$$27 \quad \hat{y}_{t+1|t} = \{\mu_{t+1|t} + \sigma_{t+1|t} F^{-1}(p/2), \mu_{t+1|t} + \sigma_{t+1|t} F^{-1}(1-p/2)\}. \quad (2.4) \quad 28$$

29 Notice that the volatility forecast plays a key role again. Note also the direct link be-  
30 tween the interval forecast and the optimal point forecast for the asymmetric linear loss  
31 function in (2.3).

### 32 2.1.3. Probability forecasting including sign forecasting

33 A forecaster may care about the variable of interest falling above or below a certain  
34 threshold value. As an example, consider a portfolio manager who might be interested  
35 in forecasting whether the return on a stock index will be larger than the known risk-  
36 free bond return. Another example might be a rating agency forecasting if the value of  
37 a firm's assets will end up above or below the value of its liabilities and thus trigger  
38 bankruptcy. Yet another example would be a central bank forecasting the probability of  
39  
40  
41  
42  
43

inflation – or perhaps an exchange rate – falling outside its target band. In general terms, if the concern is about a variable  $y_{t+1}$  ending up above some fixed (known) threshold,  $c$ , the loss function may be expressed as

$$L(y_{t+1}, \hat{y}_{t+1|t}) = (\mathbf{I}(y_{t+1} > c) - \hat{y}_{t+1|t})^2. \quad (2.5)$$

Minimizing the expected loss by setting the first derivative equal to zero then readily yields

$$\begin{aligned} \hat{y}_{t+1|t} &= E[\mathbf{I}(y_{t+1} > c) \mid \mathcal{F}_t] = P(y_{t+1} > c \mid \mathcal{F}_t) \\ &= 1 - F((c - \mu_{t+1|t})/\sigma_{t+1|t}). \end{aligned} \quad (2.6)$$

Thus, volatility dynamics are immediately important for these types of probability forecasts, even if the conditional mean is constant and not equal to  $c$ ; i.e.,  $c - \mu_{t+1|t} \neq 0$ .

The important special case where  $c = 0$  is sometimes referred to as sign forecasting. In this situation,

$$\hat{y}_{t+1|t} = 1 - F(-\mu_{t+1|t}/\sigma_{t+1|t}). \quad (2.7)$$

Hence, the volatility dynamics will affect the forecast as long as the conditional mean is not zero, or the conditional mean is not directly proportional to the standard deviation.

#### 2.1.4. Density forecasting

In many applications the entire conditional density of the variable in question is of interest. That is, the forecast takes the form of a probability distribution function

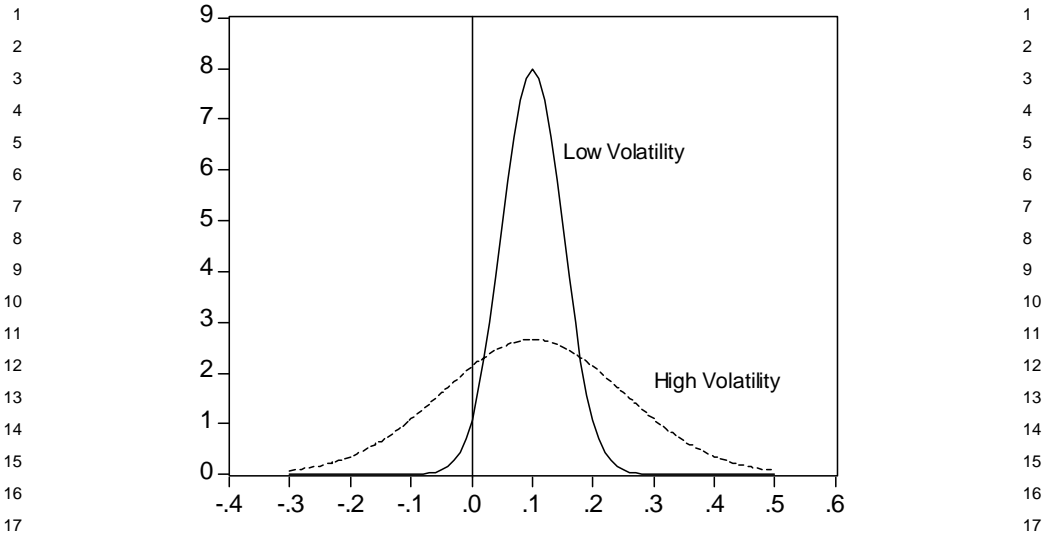
$$\hat{y}_{t+1|t} = f_{t+1|t}(y) \equiv f(y_{t+1} = y \mid \mu_{t+1|t}, \sigma_{t+1|t}) = f(y_{t+1} = y \mid \mathcal{F}_t). \quad (2.8)$$

Of course, the probability density function may itself be time-varying, for example, due to time-varying conditional skewness or kurtosis, but as noted earlier for simplicity we rule out these higher order effects here.

Figure 2 shows two stylized density forecasts corresponding to a high and low volatility day, respectively. Notice that the mean outcome is identical (and positive) on the two days. However, on the high volatility day the occurrence of a large negative (or large positive) outcome is more likely. Notice also that the probability of a positive outcome (of any size) is smaller on the high volatility day than on the low volatility day. Thus, as discussed in the preceding sections, provided that the level of the volatility is forecastable, the figure indicates some degree of sign predictability, despite the constant mean.

#### 2.2. Financial applications

The trade-off between risk and expected return, where risk is associated with some notion of price volatility, constitute one of the key concepts in modern finance. As such, measuring and forecasting volatility is arguably among the most important pursuits in empirical asset pricing finance and risk management.



18 Figure 2. Density forecasts on high volatility and low volatility days. The figure shows two hypothetical  
 19 return distributions for a low volatility (solid line) and high volatility (dashed line) day. The areas to the left  
 20 of the vertical line represent the probability of a negative return.  
 21

22 *2.2.1. Risk management: Value-at-Risk (VaR) and Expected Shortfall (ES)*  
 23

24 Consider a portfolio of returns formed from a vector of  $N$  risky assets,  $R_{t+1}$ , with  
 25 corresponding vector of portfolio weights,  $W_t$ . The portfolio return is defined as  
 26

$$27 \quad r_{w,t+1} = \sum_{i=1}^N w_{i,t} r_{i,t+1} \equiv W_t' R_{t+1}, \quad (2.9)$$

29 where the  $w$  subscript refers to the fact that the portfolio distribution depends on the  
 30 actual portfolio weights.  
 31

32 Financial risk managers often report the riskiness of the portfolio using the concept  
 33 of Value-at-Risk (VaR) which is simply the quantile of the conditional portfolio distrib-  
 34 ution. If we model the portfolio returns directly as a univariate process,  
 35

$$36 \quad r_{w,t+1} = \mu_{w,t+1|t} + \sigma_{w,t+1|t} z_{w,t+1}, \quad z_{w,t+1} \sim \text{i.i.d. } F_w, \quad (2.10)$$

37 then the VaR is simply  
 38

$$39 \quad \text{VaR}_{t+1|t}^p = \mu_{w,t+1|t} + \sigma_{w,t+1|t} F_w^{-1}(p). \quad (2.11)$$

40 This, of course, corresponds directly to the lower part of the interval forecast previously  
 41 defined in Equation (2.4).  
 42

42 Figure 3 shows a typical simulated daily portfolio return time series with dynamic  
 43 volatility (solid line). The short-dashed line, which tracks the lower range of the return,  
 43




Figure 3. Simulated portfolio returns with dynamic volatility and historical simulation VaRs. The solid line shows a time series of typical simulated daily portfolio returns. The short-dashed line depicts the true one-day-ahead, 1% VaR. The long-dashed line gives the 1% VaR based on the so-called Historical Simulation (HS) technique and a 500-day moving window.

depicts the true 1-day, 1% VaR corresponding to the simulated portfolio return. Notice that the true VaR varies considerably over time and increases in magnitude during bursts in the portfolio volatility. The relatively sluggish long-dashed line calculates the VaR using the so-called Historical Simulation (HS) technique. This is a very popular approach in practice. Rather than explicitly modeling the volatility dynamics, the HS technique calculates the VaR as an empirical quantile based on a moving window of the most recent 250 or 500 days. The HS VaR in Figure 3 is calculated using a 500-day window. Notice how this HS VaR reacts very sluggishly to changes in the volatility, and generally is too large (in absolute value) when the volatility is low, and more importantly too small (in absolute value) when the volatility is high. Historical simulation thus underestimates the risk when the risk is high. This is clearly not a prudent risk management practice. As such, these systematic errors in the HS VaR clearly highlight the value of explicitly modeling volatility dynamics in financial risk management.

The VaR depicted in Figure 3 is a very popular risk-reporting measure in practice, but it obviously only depicts a very specific aspect of the risk; that is with probability  $p$  the loss will be at least the VaR. Unfortunately, the VaR measure says nothing about the expected magnitude of the loss on the days the VaR is breached.

Alternatively, the Expected Shortfall (ES) risk measure was designed to provide additional information about the tail of the distribution. It is defined as the expected loss

on the days when losses are larger than the VaR. Specifically,

$$ES_{t+1|t}^P \equiv E[r_{w,t+1} \mid r_{w,t+1} < \text{VaR}_{t+1|t}^P] = \mu_{w,t+1|t} + \sigma_{w,t+1|t} \text{EF}_w^P. \quad (2.12)$$

Again, it is possible to show that if  $z_{w,t}$  is i.i.d., the multiplicative factor,  $\text{EF}_w^P$ , is constant and depends only on the shape of the distribution,  $F_w$ . Thus, the volatility dynamics plays a similar role in the ES risk measure as in the VaR in Equation (2.11).

The analysis above assumed a univariate portfolio return process specified as a function of the portfolio weights at any given time. Such an approach is useful for risk measurement but is not helpful, for example, for calculating optimal portfolio weights. If active risk management is warranted, say maximizing expected returns subject to a VaR constraint, then a multivariate model is needed. If we assume that each return is modeled separately then the vector of returns can be written as

$$R_{t+1} = M_{t+1|t} + \Omega_{t+1|t}^{1/2} Z_{t+1}, \quad Z_{t+1} \sim \text{i.i.d. } F, \quad (2.13)$$

where  $M_{t+1|t}$  and  $\Omega_{t+1|t}$  denote the vector of conditional mean returns and the covariance matrix for the returns, respectively, and all of the elements in the vector random process,  $Z_t$ , are independent with mean zero and variance one. Consequently, the mean and the variance of the portfolio returns,  $W_t' R_{t+1}$ , may be expressed as

$$\mu_{w,t+1|t} = W_t' M_{t+1|t}, \quad \sigma_{w,t+1|t}^2 = W_t' \Omega_{t+1|t} W_t. \quad (2.14)$$

In the case of the normal distribution,  $Z_{t+1} \sim N(0, I)$ , linear combinations of multivariate normal variables are themselves normally distributed, so that  $r_{w,t+1} \equiv W_t' R_{t+1} \sim N(\mu_{w,t+1|t}, \sigma_{w,t+1|t}^2)$ , but this aggregation property does not hold in general for other multivariate distributions. Hence, except in special cases, such as the multivariate normal, the VaR and ES measures are not known in closed form, and will have to be calculated using Monte Carlo simulation.

### 2.2.2. Covariance risk: Time-varying betas and conditional Sharpe ratios

The above discussion has focused on measuring the risk of a portfolio from purely statistical considerations. We now turn to a discussion of the more fundamental economic issue of the expected return on an asset given its risk profile. Assuming the absence of arbitrage opportunities a fundamental theorem in finance then proves the existence of a stochastic discount factor, say  $\text{SDF}_{t+1}$ , which can be used to price any asset, say  $I$ , via the conditional expectation

$$E[\text{SDF}_{t+1}(1 + r_{i,t+1}) \mid \mathcal{F}_t] = 1. \quad (2.15)$$

In particular, the return on the risk free asset, which pays one dollar for sure the next period, must satisfy  $1 + r_{f,t} = E[\text{SDF}_{t+1} \mid \mathcal{F}_t]^{-1}$ . It follows also directly from (2.15) that the expected excess return on any risky asset must be proportional to its covariance with the stochastic discount factor,

$$E[r_{i,t+1} - r_{f,t} \mid \mathcal{F}_t] = -(1 + r_{f,t}) \text{Cov}(\text{SDF}_{t+1}, r_{i,t+1} \mid \mathcal{F}_t). \quad (2.16)$$



Now, assuming that the stochastic discount factor is linearly related to the market return,

$$\text{SDF}_{t+1} = a_t - b_t(1 + r_{M,t+1}), \quad (2.17)$$

it follows from  $E[\text{SDF}_{t+1}(1 + r_{M,t+1}) | \mathcal{F}_t] = 1$  and  $1 + r_{f,t} = E[\text{SDF}_{t+1} | \mathcal{F}_t]^{-1}$  that

$$a_t = (1 + r_{f,t})^{-1} + b_t \mu_{M,t+1|t}, \quad (2.18)$$

$$b_t = (1 + r_{f,t})^{-1} (\mu_{M,t+1|t} - r_{f,t}) / \sigma_{M,t+1|t}^2,$$

where  $\mu_{M,t+1|t} \equiv E[1 + r_{M,t+1} | \mathcal{F}_t]$  and  $\sigma_{M,t+1|t}^2 \equiv \text{Var}[r_{M,t+1} | \mathcal{F}_t]$ . Notice that the dynamics in the moments of the market return (along with any dynamics in the risk-free rate) render the coefficients in the SDF time varying. Also, in parallel to the classic one-period CAPM model of Markowitz (1952) and Sharpe (1964), the *conditional* expected excess returns must satisfy the relation,

$$E[r_{i,t+1} - r_{f,t} | \mathcal{F}_t] = \beta_{i,t} (\mu_{M,t+1|t} - r_{f,t}), \quad (2.19)$$

where the *conditional* “beta” is defined by  $\beta_{i,t} \equiv \text{Cov}(r_{M,t+1}, r_{i,t+1} | \mathcal{F}_t) / \sigma_{M,t+1|t}^2$ . Moreover, the expected risk adjusted return, also known as the *conditional* Sharpe ratio, equals

$$\begin{aligned} \text{SR}_t &\equiv E[r_{i,t+1} - r_{f,t} | \mathcal{F}_t] / \text{Var}(r_{i,t+1} | \mathcal{F}_t)^{1/2} \\ &= \text{Corr}(r_{M,t+1}, r_{i,t+1} | \mathcal{F}_t) / \sigma_{M,t+1|t}. \end{aligned} \quad (2.20)$$

The simple asset pricing framework above illustrates how the expected return (raw and risk adjusted) on various assets will be driven by the mean and volatility dynamics of the overall market return as well as the dynamics of the covariance between the market and the individual assets. Covariance forecasting is thus at least as important as volatility forecasting in the context of financial asset pricing, and we discuss each in subsequent sections.

### 2.2.3. Asset allocation with time-varying covariances

The above CAPM model imposes a very restrictive structure on the covariance matrix of asset returns. In this section we instead assume a generic dynamic covariance matrix and study the optimization problem of an investor who constructs a portfolio of  $N$  risky assets by minimizing the portfolio variance subject to achieving a certain target portfolio return,  $\mu_p$ .

Formally, the investor chooses a vector of portfolio weights,  $W_t$ , by solving the quadratic programming problem

$$\min W_t' \Omega_{t+1|t} W_t \quad \text{s.t.} \quad W_t' M_{t+1|t} = \mu_p. \quad (2.21)$$

From the corresponding first order conditions, the resulting portfolio weights for the risky assets satisfy

$$W_t^* = \frac{\Omega_{t+1|t}^{-1} M_{t+1|t}}{M'_{t+1|t} \Omega_{t+1|t}^{-1} M_{t+1|t}} \mu_p, \quad (2.22)$$

with the optimal portfolio weight for the risk-free asset given by

$$w_{f,t}^* = 1 - \sum_{i=1}^N w_{i,t}^*. \quad (2.23)$$

Moreover, from (2.21) the portfolio Sharpe ratio equals

$$SR_t = \mu_p / \sqrt{W_t^{*'} \Omega_{t+1|t} W_t^*}. \quad (2.24)$$

Just as in the CAPM pricing model discussed above, both volatility and covariance dynamics are clearly important for asset allocation. Notice also that even if we rule out exploitable conditional mean dynamics, the optimal portfolio weights would still be time-varying from the second moment dynamics alone.

#### 2.2.4. Option valuation with dynamic volatility

The above tools are useful for the analysis of primitive securities with linear payoffs such as stocks, bonds, foreign exchange and futures contracts. Consider now instead a European call option which gives the owner the right but not the obligation to buy the underlying asset (say a stock or currency) on a future date,  $T$ , at a strike price,  $K$ . The option to exercise creates a nonlinear payoff which in turn requires a special set of tools for pricing and risk management.

In the Black–Scholes–Merton (BSM) option pricing model the returns are assumed to be normally distributed with constant volatility,  $\sigma$ , along with the possibility of (cost-less) continuous trading and a constant risk free rate,  $r_f$ . In this situation, the call price of an option equals

$$c_t = \text{BSM}(s_t, \sigma^2, K, r_f, T) = s_t \Phi(d) - K \exp(-r_f T) \Phi(d - \sigma \sqrt{T}), \quad (2.25)$$

where  $s_t$  denotes the current price of the asset,  $d = (\ln(s_t/K) + T(r_f + \sigma^2/2))/(\sigma \sqrt{T})$ , and  $\Phi(\cdot)$  refers to the cumulative normal distribution function.

Meanwhile, the constant volatility assumption in BSM causes systematic pricing errors when comparing the theoretical prices with actual market prices. One manifestation of this is the famous volatility-smiles which indicate systematic underpricing by the BSM model for in- or out-of-the-money options. The direction of these deviations, however, are readily explained by the presence of stochastic volatility, which creates fatter tails than the normal distribution, in turn increasing the value of in- and out-of-the-money options relative to the constant-volatility BSM model.

In response to this, Hull and White (1987) explicitly allow for an independent stochastic volatility factor in the process for the underlying asset return. Assuming that this additional volatility risk factor is not priced in equilibrium, the Hull–White call option price simply equals the expected BSM price, where the expectation is taken over the future integrated volatility. More specifically, defining the integrated volatility as the integral of the spot volatility during the remaining life of the option,

$$IV(T, t) = \int_t^T \sigma^2(u) du,$$

where  $IV(T, t) = IV(T) + IV(T-1) + \dots + IV(t+1)$  generalizes the integrated variance concept from Equation (1.11) to a multi-period horizon in straightforward fashion. The Hull–White option valuation formula may then be succinctly written as

$$C_t = E[\text{BSM}(IV(T, t)) \mid \mathcal{F}_t]. \quad (2.26)$$

In discrete time, the integrated volatility may be approximated by the sum of the corresponding one-period conditional variances,

$$IV(T, t) \approx \sum_{\tau=t}^{T-1} \sigma_{\tau+1|\tau}^2.$$

Several so-called realized volatility measures have also recently been proposed in the literature for (ex-post) approximating the integrated volatility. We will return to a much more detailed discussion of these measures in Sections 4 and 5 below.

Another related complication that arises in the pricing of equity options, in particular, stems from the apparent negative correlation between the returns and the volatility. This so-called leverage effect, as discussed further below, induces negative skewness in the return distribution and causes systematic asymmetric pricing errors in the BSM model.

Assuming a mean-reverting stochastic volatility process, Heston (1993) first developed an option pricing formula where the innovations to the returns and the volatility are correlated, and where the volatility risk is priced by the market. In contrast to the BSM setting, where an option can be hedged using a dynamically rebalanced stock and bond portfolio alone, in the Heston model an additional position must be taken in another option in order to hedge the volatility risk.

Relying on Heston's formulation, Fouque, Papanicolaou and Sircar (2000) show that the price may conveniently be expressed as

$$C_t = E[\text{BSM}(\xi_{t,T} s_t, (1 - \rho^2) IV(T, t)) \mid \mathcal{F}_t], \quad (2.27)$$

where  $\rho$  refers to the (instantaneous) correlation between the returns and the volatility, and  $\xi_{t,T}$  denotes a stochastic scaling factor determined by the volatility risk premium, with the property that  $E[\xi_{t,T} \mid \mathcal{F}_t] = 1$ . Importantly, however, the integrated volatility remains the leading term as in the Hull–White valuation formula.

### 2.3. Volatility forecasting in fields outside finance

Although volatility modeling and forecasting has proved to be extremely useful in finance, the motivation behind Engle's (1982) original ARCH model was to provide a tool for measuring the dynamics of inflation uncertainty. Tools for modeling volatility dynamics have been applied in many other areas of economics and indeed in other areas of the social sciences, the natural sciences and even medicine. In the following we list a few recent papers in various fields showcasing the breath of current applications of volatility modeling and forecasting. It is by no means an exhaustive list but these papers can be consulted for further references.

Related to Engle's original work, the modeling of inflation uncertainty and its relationship with labor market variables has recently been studied by Rich and Tracy (2004). They corroborate earlier findings of an inverse relationship between desired labor contract durations and the level of inflation uncertainty. Analyzing the inflation and output forecasts from the Survey of Professional Forecasters, Giordani and Soderlind (2003) find that while each forecaster on average tends to underestimate uncertainty, the disagreement between forecasters provides a reasonable proxy for inflation and output uncertainty. The measurement of uncertainty also plays a crucial role in many micro-economic models. Meghir and Pistaferri (2004), for instance, estimate the conditional variance of income shocks at the microlevel and find strong evidence of temporal variance dynamics.

Lastrapes (1989) first analyzed the relationship between exchange rate volatility and U.S. monetary policy. In a more recent study, Ruge-Murcia (2004) developed a model of a central bank with asymmetric preferences for unemployment above versus below the natural rate. The model implies an inflation bias proportional to the conditional variance of unemployment. Empirically, the conditional variance of unemployment is found to be positively related to the rate of inflation. In another central banking application, Tse and Yip (2003) use volatility models to study the effect on changes in the Hong Kong currency board on interbank market rates.

Volatility modeling and forecasting methods have also found several interesting uses in agricultural economics. Ramirez and Fadiga (2003), for instance, find evidence of asymmetric volatility patterns in U.S. soybean, sorghum and wheat prices. Building on the earlier volatility spill-over models used in analyzing international financial market linkages in the papers by Engle, Ito and Lin (1990) and King, Sentana and Wadhvani (1994), Buguk, Hudson and Hanson (2003) have recently used similar methods in documenting strong price volatility spillovers in the supply-chain of fish production. The volatility in feeding material prices (e.g., soybeans) affects the volatility of fish feed prices which in turn affect fish farm price volatility and finally wholesale price volatility. Also, Barrett (1999) uses a GARCH model to study the effect of real exchange rate depreciations on stochastic producer prices in low-income agriculture.

The recent deregulation in the utilities sector has also prompted many new applications of volatility modeling of gas and power prices. Shawky, Marathe and Barrett (2003) use dynamic volatility models to determine the minimum variance hedge ratios

1 for electricity futures. [Linn and Zhu \(2004\)](#) study the effect of natural gas storage report 1  
2 announcements on intraday volatility patterns in gas prices. They also find evidence of 2  
3 strong intraday patterns in natural gas price volatility. [Battle and Barquin \(2004\)](#) use a 3  
4 multivariate GARCH model to simulate gas and oil price paths, which in turn are shown 4  
5 to be useful for risk management in the wholesale electricity market. 5

6 In a related context, [Taylor and Buizza \(2003\)](#) use weather forecast uncertainty to 6  
7 model electricity demand uncertainty. The variability of wind measurements is found 7  
8 to be forecastable using GARCH models in [Dripps and Dunsmuir \(?\)](#), while tempera- 8  
9 ture forecasting with seasonal volatility dynamics is explored in [Campbell and Diebold](#) 9  
10 [\(2005\)](#). [Marinova and McAleer \(2003\)](#) model volatility dynamics in ecological patents. 10

11 In political science, [Maestas and Preuhs \(2000\)](#) suggest modeling political volatility 11  
12 broadly defined as periods of rapid and extreme change in political processes, while 12  
13 [Gronke and Brehm \(2002\)](#) use ARCH models to assess the dynamics of volatility in 13  
14 presidential approval ratings. 14

15 Volatility forecasting has recently found applications even in medicine. [Ewing, Piette](#) 15  
16 [and Payne \(2003\)](#) forecast time varying volatility in medical net discount rates which 16  
17 are in turn used to determine the present value of future medical costs. Also, [Johnson,](#) 17  
18 [Elashoff and Harkema \(2003\)](#) use a heteroskedastic time series process to model neuro- 18  
19 muscular activation patterns in patients with spinal cord injuries, while [Martin-Guerrero](#) 19  
20 [et al. \(2003\)](#) use a dynamic volatility model to help determine the optimal EPO dosage 20  
21 for patients with secondary anemia. 21

#### 23 2.4. Further reading 23

24  
25 Point forecasting under general loss functions when allowing for dynamic volatility 25  
26 has been analyzed by [Christoffersen and Diebold \(1996, 1997\)](#). [Patton and Timmer-](#) 26  
27 [mann \(2004\)](#) have recently shown that under general specifications of the loss function, 27  
28 the optimal forecast error will have a conditional expected value that is a function of 28  
29 the conditional variance of the underlying process. Methods for incorporating time- 29  
30 varying volatility into interval forecasts are suggested in [Granger, White and Kamstra](#) 30  
31 [\(1989\)](#). Financial applications of probability forecasting techniques are considered in 31  
32 [Christoffersen and Diebold \(2003\)](#). 32

33 Financial risk management using dynamic volatility models is surveyed in [Christof-](#) 33  
34 [fersen \(2003\)](#) and [Jorion \(2000\)](#). [Berkowitz and O'Brien \(2002\)](#), [Pritsker \(2001\)](#), and 34  
35 [Barone-Adesi, Giannopoulos and Vosper \(1999\)](#) explicitly document the value added 35  
36 from incorporating volatility dynamics into daily financial risk management systems. 36

37 Volatility forecasting at horizons beyond a few weeks is found to be difficult by [West](#) 37  
38 [and Cho \(1995\)](#) and [Christoffersen and Diebold \(2000\)](#). However, [Brandt and Jones](#) 38  
39 [\(2002\)](#) show that using intraday information improves the longer horizon forecasts con- 39  
40 siderably. [Guidolin and Timmermann \(2005a\)](#) uncover VaR dynamics at horizons of up 40  
41 to two years. [Campbell \(1987, 2003\)](#), [Shanken \(1990\)](#), [Ait-Sahalia and Brandt \(2001\)](#), 41  
42 [Harvey \(2001\)](#), [Lettau and Ludvigsson \(2003\)](#) and [Marquering and Verbeek \(2004\)](#) find 42  
43 that interest rate spreads and financial ratios help predict volatility at longer horizons. 43

A general framework for conditional asset pricing allowing for time-varying betas is laid out in Cochrane (2001). Market timing arising from time-varying Sharpe ratios is analyzed in Whitelaw (1997), while volatility timing has been explicitly explored by Johannes, Polson and Stroud (2004). The relationship between time-varying volatility and return has been studied in Engle, Lilien and Robbins (1987), French, Schwert and Stambaugh (1987), Bollerslev, Engle and Wooldridge (1988), Bollerslev, Chou and Kroner (1992), Glosten, Jagannathan and Runkle (1993), among many others.

The value of modeling volatility dynamics for asset allocation in a single-period setting have been highlighted in the series of papers by Fleming, Kirby and Oestdiek (2001, 2003), with multi-period extensions considered by Wang (2004). The general topic of asset allocation under predictable returns is surveyed in Brandt (2004). Brandt (1999) and Ait-Sahalia and Brandt (2001) suggest portfolio allocation methods which do not require the specification of conditional moment dynamics.

The literature on option valuation allowing for volatility dynamics is very large and active. In addition to some of the key theoretical contributions mentioned above, noteworthy empirical studies based on continuous-time methods include Bakshi, Cao and Chen (1997), Bates (1996), Chernov and Ghysels (2000), Eraker (2004), Melino and Turnbull (1990), and Pan (2002). Recent discrete-time applications, building on the theoretical work of Duan (1995) and Heston (1993), can be found in Christoffersen and Jacobs (2004a, 2004b) and Heston and Nandi (2000).

### 3. GARCH volatility

The current interest in volatility modeling and forecasting was spurred by Engle's (1982) path breaking ARCH paper, which set out the basic idea of modeling and forecasting volatility as a time-varying function of current information. The GARCH class of models, of which the GARCH(1, 1) remains the workhorse, were subsequently introduced by Bollerslev (1986), and also discussed independently by Taylor (1986). These models, including their use in volatility forecasting, have been extensively surveyed elsewhere and we will not attempt yet another exhaustive survey here. Instead we will try to highlight some of the key features of the models which help explain their dominant role in practical empirical applications. We will concentrate on univariate formulations in this section, with the extension to multivariate GARCH-based covariance and correlation forecasting discussed in Section 6.

#### 3.1. Rolling regressions and RiskMetrics

Rolling sample windows arguably provides the simplest way of incorporating actual data into the estimation of time-varying volatilities, or variances. In particular, consider the rolling sample variance based on the  $p$  most recent observations as of time  $t$ ,

$$\hat{\sigma}_t^2 = p^{-1} \sum_{i=0}^{p-1} (y_{t-i} - \hat{\mu})^2 \equiv p^{-1} \sum_{i=0}^{p-1} \hat{\varepsilon}_{t-i}^2. \quad (3.1)$$

1 Interpreting  $\hat{\sigma}_t^2$  as an estimate of the current variance of  $y_t$ , the value of  $p$  directly 1  
 2 determines the variance-bias tradeoff of the estimator, with larger values of  $p$  reducing 2  
 3 the variance but increasing the bias. For instance, in the empirical finance literature, it is 3  
 4 quite common to rely on rolling samples of five-years of monthly data, corresponding 4  
 5 to a value of  $p = 60$ , in estimating time varying-variances, covariances, and CAPM 5  
 6 betas. 6

7 Instead of weighting each of the most recent  $p$  observations the same, the bias of the 7  
 8 estimator may be reduced by assigning more weights to the most recent observations. 8  
 9 An exponentially weighted moving average filter is commonly applied in doing so, 9

$$10 \hat{\sigma}_t^2 = \gamma(y_t - \hat{\mu})^2 + (1 - \gamma)\hat{\sigma}_{t-1}^2 \equiv \gamma \sum_{i=1}^{\infty} (1 - \gamma)^{i-1} \hat{\varepsilon}_{t-i}^2. \quad (3.2) \quad 10$$

11 In practice, the sum will, of course, have to be truncated at  $I = t - 1$ . This is typically 11  
 12 done by equating the pre-sample values to zero, and adjusting the finite sum by the 12  
 13 corresponding multiplication factor  $1/[1 - (1 - \gamma)^t]$ . Of course, for large values of 13  
 14  $t$  and  $(1 - \gamma) < 1$ , the effect of this truncation is inconsequential. This approach is 14  
 15 exemplified by RiskMetrics [J.P. Morgan (1996)], which rely on a value of  $\gamma = 0.06$  15  
 16 and  $\mu \equiv 0$  in their construction of daily (monthly) volatility measures for wide range 16  
 17 of different financial rates of returns. 17  
 18 18  
 19 19  
 20 20

21 Although it is possible to write down explicit models for  $y_t$  which would justify the 21  
 22 rolling window approach and the exponential weighted moving average as the optimal 22  
 23 estimators for the time-varying variances in the models, the expressions in (3.1) and 23  
 24 (3.2) are more appropriately interpreted as data-driven filters. In this regard, the the- 24  
 25oretical properties of both filters as methods for extracting consistent estimates of the 25  
 26 current volatility as the sampling frequencies of the underlying observations increases 26  
 27 over fixed-length time intervals – or what is commonly referred to as continuous record, 27  
 28 or fill-in, asymptotics – has been extensively studied in a series of influential papers 28  
 29 by Dan Nelson [these papers are collected in the edited volume of readings by Rossi 29  
 30 (1996)]. 30

31 It is difficult to contemplate optimal volatility forecasting without the notion of a 31  
 32 model, or data generating process. Of course, density or VaR forecasting, as discussed 32  
 33 in Section 2, is even more problematic. Nonetheless, the filters described above are 33  
 34 often used in place of more formal model building procedures in the construction of 34  
 35  $h$ -period-ahead volatility forecasts by simply equating the future volatility of interest 35  
 36 with the current filtered estimate, 36

$$37 \text{Var}(y_{t+h} | \mathcal{F}_t) \equiv \sigma_{t+h|t}^2 \approx \hat{\sigma}_t^2. \quad (3.3) \quad 37$$

38 In the context of forecasting the variance of multi-period returns, assuming that the 38  
 39 corresponding one-period returns are serially uncorrelated so that the forecast equals 39  
 40 the sum of the successive one-period variance forecasts, it follows then directly that 40  
 41 41

$$42 \text{Var}(y_{t+k} + y_{t+k-1} + \dots + y_{t+1} | \mathcal{F}_t) \equiv \sigma_{t:t+k|t}^2 \approx k\hat{\sigma}_t^2. \quad (3.4) \quad 42$$

43 43

Hence, the multi-period return volatility scales with the forecast horizon,  $k$ . Although this approach is used quite frequently by finance practitioners it has, as discussed further below, a number of counterfactual implications. In contrast, the GARCH(1, 1) model, to which we now turn, provides empirically realistic mean-reverting volatility forecasts within a coherent and internally consistent, yet simple, modeling framework.

### 3.2. GARCH(1, 1)

In order to define the GARCH class of models, consider the decomposition of  $y_t$  into the one-step-ahead conditional mean,  $\mu_{t|t-1} \equiv E(y_t | \mathcal{F}_{t-1})$ , and variance,  $\sigma_{t|t-1}^2 \equiv \text{Var}(y_t | \mathcal{F}_{t-1})$ , in parallel to the expression in Equation (1.7) above,

$$y_t = \mu_{t|t-1} + \sigma_{t|t-1} z_t, \quad z_t \sim \text{i.i.d.}, \quad E(z_t) = 0, \quad \text{Var}(z_t) = 1. \quad (3.5)$$

The GARCH(1, 1) model for the conditional variance is then defined by the recursive relationship,

$$\sigma_{t|t-1}^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1|t-2}^2, \quad (3.6)$$

where  $\varepsilon_t \equiv \sigma_{t|t-1} z_t$ , and the parameters are restricted to be nonnegative,  $\omega > 0$ ,  $\alpha \geq 0$ ,  $\beta \geq 0$ , in order to ensure that the conditional variance remains positive for all realizations of the  $z_t$  process. The model is readily extended to higher order GARCH( $p$ ,  $q$ ) models by including additional lagged squared innovations and/or conditional variances on the right-hand side of the equation.

By recursive substitution, the GARCH(1, 1) model may alternatively be expressed as an ARCH( $\infty$ ) model,

$$\sigma_{t|t-1}^2 = \omega(1 - \beta)^{-1} + \alpha \sum_{i=1}^{\infty} \beta^{i-1} \varepsilon_{t-i}^2. \quad (3.7)$$

This obviously reduces to the exponentially weighted moving average filter in (3.2) for  $\omega = 0$ ,  $\alpha = \gamma$ , and  $\beta = 1 - \gamma$ . The corresponding GARCH model in which  $\alpha + \beta = 1$  is also sometimes referred to as an Integrated GARCH, or IGARCH(1, 1) model. Importantly, however, what sets the GARCH(1, 1) model, and more generally the ARCH class of models, apart from the filters discussed above is the notion of a data generating process embedded in the distributional assumptions for  $z_t$ . This means that the construction of optimal variance forecasts is a well-posed question within the context of the model.

In particular, it follows directly from the formulation of the model that the optimal, in a mean-square error sense, one-step ahead variance forecasts equals  $\sigma_{t+1|t}^2$ . Corresponding expressions for the longer run forecasts,  $\sigma_{t+h|t}^2$  for  $h > 1$ , are also easily constructed by recursive procedures. To facilitate the presentation, assume that the conditional mean is constant and equal to zero, or  $\mu_{t|t-1} = 0$ , and that  $\alpha + \beta < 1$  so that the unconditional variance of the process exists,

$$\sigma^2 = \omega(1 - \alpha - \beta)^{-1}. \quad (3.8)$$






Figure 4. GARCH volatility term structure. The first panel shows the unconditional distribution of  $\sigma_{t+1|t}^2$ . The second panel shows the term-structure-of-variance,  $k^{-1}\sigma_{t:t+k|t}^2$ , for  $\sigma_{t+1|t}^2$  equal to the mean, together with the fifth and the ninety-fifth percentiles in the unconditional distribution.

The  $h$ -step ahead forecast is then readily expressed as

$$\sigma_{t+h|t}^2 = \sigma^2 + (\alpha + \beta)^{h-1}(\sigma_{t+1|t}^2 - \sigma^2), \tag{3.9}$$

showing that the forecasts revert to the long-run unconditional variance at an exponential rate dictated by the value of  $\alpha + \beta$ .

Moreover, with serially uncorrelated returns, so that the conditional variance of the sum is equal to the sum of the conditional variances, the optimal forecast for the variance of the  $k$ -period return may be expressed as

$$\sigma_{t:t+k|t}^2 = k\sigma^2 + (\sigma_{t+1|t}^2 - \sigma^2)(1 - (\alpha + \beta)^k)(1 - \alpha - \beta)^{-1}. \tag{3.10}$$

Thus, the longer the forecast horizon (the higher the value of  $k$ ), the less variable will be the forecast per unit time-interval. That is, the term-structure-of-variance, or  $k^{-1}\sigma_{t:t+k|t}^2$ , flattens with  $k$ .

To illustrate, consider Figure 4. The left-hand panel plots the unconditional distribution of  $\sigma_{t+1|t}^2$  for the same GARCH(1, 1) model depicted in Figure 1. The mean of the distribution equals  $\sigma^2 = 0.020(1 - 0.085 - 0.881)^{-1} \approx 0.588$ , but there is obviously considerable variation around that value, with a much longer tail to the right. The panel on the right gives the corresponding term-structure for  $k = 1, 2, \dots, 250$ , and  $\sigma_{t+1|t}^2$  equal to the mean, five, and ninety-five percentiles in the unconditional distribution. The slope of the volatility-term-structure clearly flattens with the horizon. The figure also illustrates that the convergence to the long-run unconditional variance occurs much slower for a given percentage deviation of  $\sigma_{t+1|t}^2$  above the median than for the same percentage deviation below the median.

To further illustrate the dynamics of the volatility-term structure, Figure 5 graphs  $k^{-1}\sigma_{t:t+k|t}^2$  for  $k = 1, 5, 22$  and  $66$ , corresponding to daily, weekly, monthly and

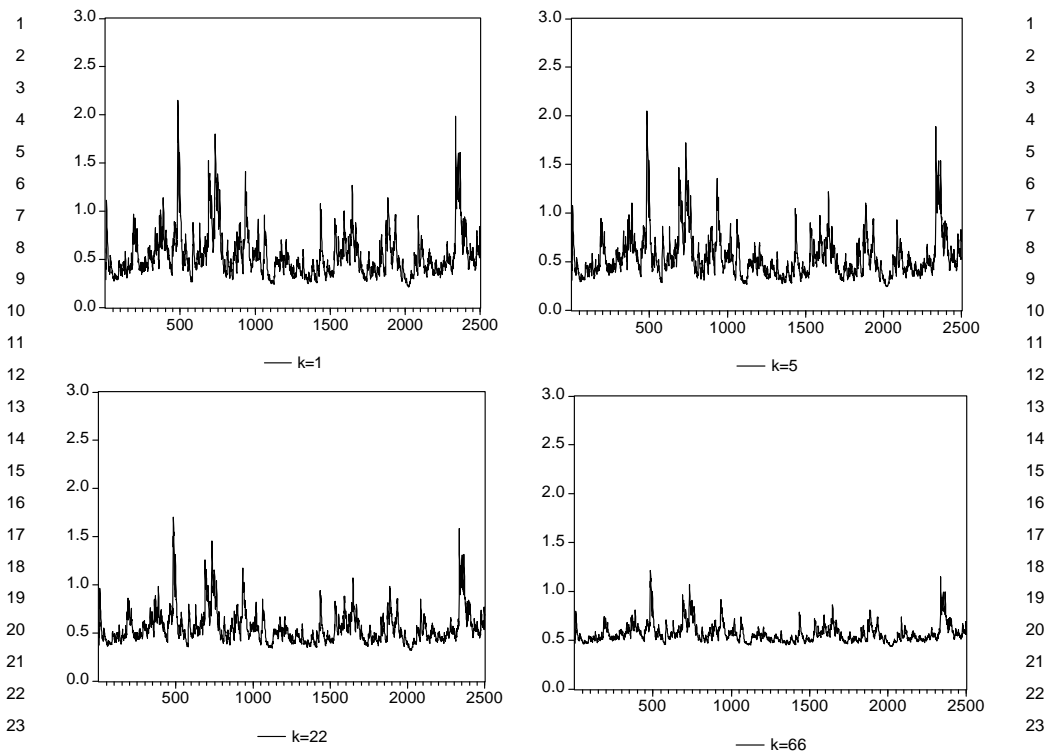


Figure 5. GARCH volatility forecasts and horizons. The four panels show the standardized “daily” GARCH(1, 1) volatility forecasts,  $k^{-1}\sigma_{t:t+k|t}^2$ , for horizons  $k = 1, 5, 22, 66$ .

quarterly forecast horizons, for the same  $t = 1, 2, \dots, 2500$  GARCH(1, 1) simulation sample depicted in Figure 1. Comparing the four different panels, the volatility-of-the-volatility clearly diminishes with the forecast horizon.

It is also informative to compare and contrast the optimal GARCH(1, 1) volatility forecasts to the common empirical practice of horizon volatility scaling by  $k$ . In this regard, it follows directly from the expressions in (3.9) and (3.10) that

$$E(k\sigma_{t+1|t}^2) = k\sigma^2 = E(\sigma_{t:t+k|t}^2),$$

so that the level of the scaled volatility forecasts will be right on average. However, comparing the variance of the scaled  $k$ -period forecasts to the variance of the optimal forecast,

$$\begin{aligned} \text{Var}(k\sigma_{t+1|t}^2) &= k^2 \text{Var}(\sigma_{t+1|t}^2) > (1 - (\alpha + \beta)^k)^2 (1 - \alpha - \beta)^{-2} \text{Var}(\sigma_{t:t+k|t}^2) \\ &= \text{Var}(\sigma_{t:t+k|t}^2), \end{aligned}$$

1 it is obvious that by not accounting for the mean-reversion in the volatility, the scaled 1  
 2 forecasts exaggerate the volatility-of-the-volatility relative to the true predictable varia- 2  
 3 tion. On tranquil days the scaled forecasts underestimate the true risk, while the risk is 3  
 4 inflated on volatile days. Obviously not a very prudent risk management procedure. 4

5 This tendency for the horizon scaled forecasts to exhibit excessive variability is also 5  
 6 directly evident from the term structure plots in Figure 5. Consider the optimal  $k$ -period 6  
 7 ahead variance forecasts defined by  $k$  times the  $k^{-1}\sigma_{t:t+k|t}^2$  series depicted in the last 7  
 8 three panels. Contrasting these correct multi-step forecasts with their scaled counter- 8  
 9 parts defined by  $k$  times the  $\sigma_{t+1|t}^2$  series in the first panel, it is obvious, that although 9  
 10 both forecasts will be centered around the right unconditional value of  $k\sigma^2$ , the horizon 10  
 11 scaled forecasts will result in too large “day-to-day” fluctuations. This is especially true 11  
 12 for the longer run “monthly” ( $k = 22$ ) and “quarterly” ( $k = 66$ ) forecasts in the last two 12  
 13 panels. 13

### 14 3.3. Asymmetries and “leverage” effects 15

16 The basic GARCH model discussed in the previous section assumes that positive and 16  
 17 negative shocks of the same absolute magnitude will have the identical influence on the 17  
 18 future conditional variances. In contrast, the volatility of aggregate equity index return, 18  
 19 in particular, has been shown to respond asymmetrically to past negative and positive 19  
 20 return shocks, with negative returns resulting in larger future volatilities. This asym- 20  
 21 metry is generally referred to as a “leverage” effect, although it is now widely agreed 21  
 22 that financial leverage alone cannot explain the magnitude of the effect, let alone the 22  
 23 less pronounced asymmetry observed for individual equity returns. Alternatively, the 23  
 24 asymmetry has also been attributed to a “volatility feedback” effect, whereby height- 24  
 25 ened volatility requires an increase in the future expected returns to compensate for 25  
 26 the increased risk, in turn necessitating a drop in the current price to go along with 26  
 27 the initial increase in the volatility. Regardless of the underlying economic explanation 27  
 28 for the phenomenon, the three most commonly used GARCH formulations for describ- 28  
 29 ing this type of asymmetry are the GJR or Threshold GARCH (TGARCH) models of 29  
 30 Glosten, Jagannathan and Runkle (1993) and Zakoian (1994), the Asymmetric GARCH 30  
 31 (AGARCH) model of Engle and Ng (1993), and the Exponential GARCH (EGARCH) 31  
 32 model of Nelson (1991). 32  
 33

34 The conditional variance in the GJR(1, 1), or TGARCH(1, 1), model simply aug- 34  
 35 ments the standard GARCH(1, 1) formulation with an additional ARCH term condi- 35  
 36 tional on the sign of the past innovation, 36

$$37 \sigma_{t|t-1}^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 I(\varepsilon_{t-1} < 0) + \beta \sigma_{t-1|t-2}^2, \quad (3.11) \quad 37$$

38 where  $I(\cdot)$  denotes the indicator function. It is immediately obvious that for  $\gamma > 0$ , past 38  
 39 negative return shocks will have a larger impact on the future conditional variances. 39  
 40 Mechanically, the calculation of multi-period variance forecast works exactly as for the 40  
 41 standard symmetric GARCH model. In particular, assuming that 41  
 42

$$42 P(z_t \equiv \sigma_{t|t-1}^{-1} \varepsilon_t < 0) = 0.5, \quad 42$$

$$43 \quad 43$$

1 it follows readily that

$$2 \quad \sigma_{t+h|t}^2 = \sigma^2 + (\alpha + 0.5\gamma + \beta)^{h-1} (\sigma_{t+1|t}^2 - \sigma^2), \quad (3.12) \quad 3$$

4 where the long-run, or unconditional variance, now equals

$$5 \quad \sigma^2 = \omega(1 - \alpha - 0.5\gamma - \beta)^{-1}. \quad (3.13) \quad 6$$

7 Although the forecasting formula looks almost identical to the one for the GARCH(1, 1) model in Equation (3.9), the inclusion of the asymmetric term may materially affect the forecasts by importantly altering the value of the current conditional variance,  $\sigma_{t+1|t}^2$ .

11 The news impact curve, defined by the functional relationship between  $\sigma_{t|t-1}^2$  and  $\varepsilon_{t-1}$  holding all other variables constant, provides a simple way of characterizing the influence of the most recent shock on next periods conditional variance. In the standard GARCH model this curve is obviously quadratic around  $\varepsilon_{t-1} = 0$ , while the GJR model with  $\gamma > 0$  has steeper slopes for negative values of  $\varepsilon_{t-1}$ . In contrast, the Asymmetric GARCH, or AGARCH(1, 1), model,

$$12 \quad \sigma_{t|t-1}^2 = \omega + \alpha(\varepsilon_{t-1} - \gamma)^2 + \beta\sigma_{t-1|t-2}^2, \quad (3.14) \quad 13$$

14 shifts the center of the news impact curve from zero to  $\gamma$ , affording an alternative way of capturing asymmetric effects. The GJR and AGARCH model may also be combined to achieve even more flexible parametric formulations.

18 Instead of directly parameterizing the conditional variance, the EGARCH model is formulated in terms of the logarithm of the conditional variance, as in the EGARCH(1, 1) model,

$$20 \quad \log(\sigma_{t|t-1}^2) = \omega + \alpha(|z_{t-1}| - E(|z_{t-1}|)) + \gamma z_{t-1} + \beta \log(\sigma_{t-1|t-2}^2), \quad (3.15) \quad 21$$

22 where as previously defined,  $z_t \equiv \sigma_{t|t-1}^{-1} \varepsilon_t$ . As for the GARCH model, the EGARCH model is readily extended to higher order models by including additional lags on the right-hand side. The parameterization in terms of logarithms has the obvious advantage of avoiding nonnegativity constraints on the parameters, as the variance implied by the exponentiated logarithmic variance from the model is guaranteed to be positive. As in the GJR and AGARCH models above, values of  $\gamma > 0$  in the EGARCH model directly captures the asymmetric response, or “leverage” effect. Meanwhile, because of the nondifferentiability with respect to  $z_{t-1}$  at zero, the EGARCH model is often somewhat more difficult to estimate and analyze numerically. From a forecasting perspective, the recursions defined by the EGARCH equation (3.15) readily deliver the optimal – in a mean-square error sense – forecast for the future logarithmic conditional variances,  $E(\log(\sigma_{t+h}^2) | \mathcal{F}_t)$ . However, in most applications the interest centers on point forecasts for  $\sigma_{t+h}^2$ , as opposed to  $\log(\sigma_{t+h}^2)$ . Unfortunately, the transformation of the  $E(\log(\sigma_{t+h}^2) | \mathcal{F}_t)$  forecasts to  $E(\sigma_{t+h}^2 | \mathcal{F}_t)$  generally depends on the entire  $h$ -step ahead forecast distribution,  $f(y_{t+h} | \mathcal{F}_t)$ . As discussed further in Section 3.6 below, this distribution is generally not available in closed-form, but it may

be approximated by Monte Carlo simulations from the convolution of the corresponding  $h$  one-step-ahead predictive distributions implied by the  $z_t$  innovation process using numerical techniques. In contrast, the expression for  $\sigma_{t+h|t}^2$  in Equation (3.12) for the GJR or TGARCH model is straightforward to implement, and only depends upon the assumption that  $P(z_t < 0) = 0.5$ .

### 3.4. Long memory and component structures

The GARCH, TGARCH, AGARCH, and EGARCH models discussed in the previous sections all imply that shocks to the volatility decay at an exponential rate. To illustrate, consider the GARCH(1, 1) model. It follows readily from Equation (3.9) that the impulse effect of a time- $t$  shock on the forecast of the variance  $h$  period into the future is given by  $\partial\sigma_{t+h|t}^2/\partial\varepsilon_t^2 = \alpha(\alpha + \beta)^{h-1}$ , or more generally

$$\partial\sigma_{t+h|t}^2/\partial\varepsilon_t^2 = \kappa\delta^h, \tag{3.16}$$

where  $0 < \delta < 1$ . This exponential decay typically works well when forecasting over short horizons. However, numerous studies, including Ding, Granger and Engle (1993) and Andersen and Bollerslev (1997), have argued that the autocorrelations of squared and absolute returns decay at a much slower hyperbolic rate over longer lags. In the context of volatility forecasting using GARCH models parameterized in terms of  $\varepsilon_t^2$ , this suggests that better long term forecasts may be obtained by formulating the conditional variance in such a way that the impulse effect behaves as

$$\partial\sigma_{t+h|t}^2/\partial\varepsilon_t^2 \approx \kappa h^\delta, \tag{3.17}$$

for large values of  $h$ , where again  $0 < \delta < 1$ . Several competing long-memory, or fractionally integrated, GARCH type models have been suggested in the literature to achieve this goal.

In the Fractionally Integrated FIGARCH(1,  $d$ , 1) model proposed by Baillie, Bollerslev and Mikkelsen (1996) the conditional variance is defined by

$$\sigma_{t|t-1}^2 = \omega + \beta\sigma_{t-1|t-2}^2 + [1 - \beta L - (1 - \alpha L - \beta L)(1 - L)^d]\varepsilon_t^2. \tag{3.18}$$

For  $d = 0$  the model reduces to the standard GARCH(1, 1) model, but for values of  $0 < d < 1$  shocks to the point volatility forecasts from the model will decay at a slow hyperbolic rate. The actual forecasts are most easily constructed by recursive substitution in

$$\sigma_{t+h|t+h-1}^2 = \omega(1 - \beta)^{-1} + \lambda(L)\sigma_{t+h-1|t+h-2}^2, \tag{3.19}$$

with  $\sigma_{t+h|t+h-1}^2 \equiv \varepsilon_t^2$  for  $h < 0$ , and the coefficients in  $\lambda(L) \equiv 1 - (1 - \beta L)^{-1}(1 - \alpha L - \beta L)(1 - L)^d$  calculated from the recursions,

$$\lambda_1 = \alpha + d,$$

$$\lambda_j = \beta\lambda_{j-1} + [(j - 1 - d)j^{-1} - (\alpha + \beta)]\delta_{j-1}, \quad j = 2, 3, \dots,$$

where  $\delta_j \equiv \delta_{j-1}(j-1-d)j^{-1}$  refer to the coefficients in the MacLaurin series expansion of the fractional differencing operator,  $(1-L)^d$ . Higher order FIGARCH models, or volatility forecast filters, may be defined in an analogous fashion. Asymmetries are also easily introduced into the recursions by allowing for separate influences of past positive and negative innovations as in the GJR or TGARCH model. Fractional Integrated EGARCH, or FIEGARCH, models may be similarly defined by parameterizing the logarithmic conditional variance as a fractionally integrated distributed lag of past values.

An alternative, and often simpler, approach for capturing longer-run dependencies involves the use of component type structures. Granger (1980) first showed that the superposition of an infinite number of stationary AR(1) processes may result in a true long-memory process. In fact, there is a long history in statistics and time series econometrics for approximating long memory by the sum of a few individually short-memory components. This same idea has successfully been used in the context of volatility modeling by Engle and Lee (1999) among others.

In order to motivate the Component GARCH model of Engle and Lee (1999), rewrite the standard GARCH(1, 1) model in (3.6) as

$$(\sigma_{t|t-1}^2 - \sigma^2) = \alpha(\varepsilon_{t-1}^2 - \sigma^2) + \beta(\sigma_{t-1|t-2}^2 - \sigma^2), \quad (3.20)$$

where it is assumed that  $\alpha + \beta < 1$ , so that the model is covariance stationary and the long term forecasts converge to the long-run, or unconditional, variance  $\sigma^2 = \omega(1 - \alpha - \beta)^{-1}$ . The component model then extends the basic GARCH model by explicitly allowing the long-term level to be time-varying,

$$(\sigma_{t|t-1}^2 - \zeta_t^2) = \alpha(\varepsilon_{t-1}^2 - \zeta_t^2) + \beta(\sigma_{t-1|t-2}^2 - \zeta_t^2), \quad (3.21)$$

with  $\zeta_t^2$  parameterized by the separate equation,

$$\zeta_t^2 = \omega + \rho\zeta_{t-1}^2 + \varphi(\varepsilon_{t-1}^2 - \sigma_{t-1|t-2}^2). \quad (3.22)$$

Hence, the transitory dynamics is governed by  $\alpha + \beta$ , while the long-run dependencies are described by  $\rho > 0$ . It is possible to show that for the model to be covariance stationary, and the unconditional variance to exist, the parameters must satisfy  $(\alpha + \beta)(1 - \rho) + \rho < 1$ . Also, substituting the latter equation into the first, the model may be expressed as the restricted GARCH(2, 2) model,

$$\begin{aligned} \sigma_{t|t-1}^2 = & \omega(1 - \alpha - \beta) + (\alpha + \varphi)\varepsilon_{t-1}^2 - (\varphi(\alpha + \beta) + \rho\alpha)\varepsilon_{t-2}^2 \\ & + (\rho + \beta + \varphi)\sigma_{t-1|t-2}^2 + (\varphi(\alpha + \beta) - \rho\beta)\sigma_{t-2|t-3}^2. \end{aligned}$$

As for the GARCH(1, 1) model, volatility shocks therefore eventually dissipate at the exponential rate in Equation (3.15). However, for intermediate forecast horizons and values of  $\rho$  close to unity, the volatility forecasts from the component GARCH model will display approximate long memory.

To illustrate, consider Figure 6 which graphs the volatility impulse response function,  $\partial\sigma_{t+h|t}^2/\partial\varepsilon_t^2$ ,  $h = 1, 2, \dots, 250$ , for the RiskMetrics forecasts, the standard




Figure 6. Volatility impulse response coefficients. The left panel graphs the volatility impulse response function,  $\partial\sigma_{t+h|t}^2/\partial\varepsilon_t^2$ ,  $h = 1, 2, \dots, 250$ , for the RiskMetrics forecasts, the standard GARCH(1, 1) model in (3.6), the FIGARCH(1,  $d$ , 1) model in (3.18), and the component GARCH model in (3.21) and (3.22). The right panel plots the corresponding logarithmic values.

GARCH(1, 1) model in (3.6), the FIGARCH(1,  $d$ , 1) model in (3.18), and the component GARCH model defined by (3.21) and (3.22). The parameters for the different GARCH models are calibrated to match the volatilities depicted in Figure 1. To facilitate comparisons and exaggerate the differences across models, the right-hand panel depicts the logarithm of the same impulse response coefficients. The RiskMetrics forecasts, corresponding to an IGARCH(1, 1) model with  $\alpha = 0.06$ ,  $\beta = 1 - \alpha = 0.94$  and  $\omega = 0$ , obviously results in infinitely persistent volatility shocks. In contrast, the impulse response coefficients associated with the GARCH(1, 1) forecasts die out at the exponential rate  $(0.085 + 0.881)^h$ , as manifest by the log-linear relationship in the right-hand panel. Although the component GARCH model also implies an exponential decay and therefore a log-linear relationship, it fairly closely matches the hyperbolic decay rate for the long-memory FIGARCH model for the first 125 steps. However, the two models clearly behave differently for forecasts further into the future. Whether these differences and potential gains in forecast accuracy over longer horizons are worth the extra complications associated with the implementation of a fractional integrated model obviously depends on the specific uses of the forecasts.

### 3.5. Parameter estimation

The values of the parameters in the GARCH models are, of course, not known in practice and will have to be estimated. By far the most commonly employed approach for doing so is Maximum Likelihood Estimation (MLE) under the additional assumption that the standardized innovations in Equation (3.5),  $z_t \equiv \sigma_{t|t-1}^{-1}(y_t - \mu_{t|t-1})$ , are i.i.d.

1 normally distributed, or equivalently that the conditional density for  $y_t$  takes the form, 1

$$2 \quad f(y_t | \mathcal{F}_{t-1}) = (2\pi)^{-1/2} \sigma_{t|t-1}^{-1} \exp(-1/2 \sigma_{t|t-1}^{-2} (y_t - \mu_{t|t-1})^2). \quad (3.23) \quad 2$$

3  
4 In particular, let  $\theta$  denote the vector of unknown parameters entering the conditional 4  
5 mean and variance functions to be estimated. By standard recursive conditioning argu- 5  
6 ments, the log-likelihood function for the  $y_T, y_{T-1}, \dots, y_1$  sample is then simply given 6  
7 by the sum of the corresponding  $T$  logarithmic conditional densities, 7

$$8 \quad \log L(\theta; y_T, \dots, y_1) \quad 8$$

$$9 \quad = -\frac{T}{2} \log(2\pi) - \frac{1}{2} \sum_{t=1}^T [\log \sigma_{t|t-1}^2(\theta) - \sigma_{t|t-1}^{-2}(\theta) (y_t - \mu_{t|t-1}(\theta))^2]. \quad (3.24) \quad 9$$

10  
11  
12  
13 The likelihood function obviously depends upon the parameters in a highly nonlinear 13  
14 fashion, and numerical optimization techniques are required in order to find the value of 14  
15  $\theta$  which maximizes the function, say  $\hat{\theta}_T$ . Also, to start up the recursions for calculating 15  
16  $\sigma_{t|t-1}^2(\theta)$ , pre-sample values of the conditional variances and squared innovations are 16  
17 also generally required. If the model is stationary, these initial values may be fixed at 17  
18 their unconditional sample counterparts, without affecting the asymptotic distribution 18  
19 of the resulting estimates. Fortunately, there now exist numerous software packages for 19  
20 estimating all of the different GARCH formulations discussed above based upon this 20  
21 likelihood approach. 21

22 Importantly, provided that the model is correctly specified and satisfies a necessary 22  
23 set of technical regularity conditions, the estimates obtained by maximizing the func- 23  
24 tion in (3.24) inherit the usual optimality properties associated with MLE, allowing for 24  
25 standard parameter inference based on an estimate of the corresponding information 25  
26 matrix. This same asymptotic distribution may also be used in incorporating the para- 26  
27 meter estimation error uncertainty in the distribution of the volatility forecasts from the 27  
28 underlying model. However, this effect is typically ignored in practice, instead relying 28  
29 on a simple plugin approach using  $\hat{\theta}_T$  in place of the true unknown parameters in the 29  
30 forecasting formulas. Of course, in many financial applications the size of the sample 30  
31 used in the parameter estimation phase is often very large compared to the horizon of the 31  
32 forecasts, so that the additional influence of the parameter estimation error is likely to be 32  
33 relatively minor compared to the inherent uncertainty in the forecasts from the model. 33  
34 Bayesian inference procedures can, of course, also be used in directly incorporating the 34  
35 parameter estimation error uncertainty in the model forecasts. 35

36 More importantly from a practical perspective, the log-likelihood function in Equa- 36  
37 tion (3.24) employed in almost all software packages is based on the assumption that 37  
38  $z_t$  is i.i.d. normally distributed. Although this assumption coupled with time-varying 38  
39 volatility implies that the unconditional distribution of  $y_t$  has fatter tails than the normal, 39  
40 this is typically not sufficient to account for all of the mass in the tails in the distributions 40  
41 of daily or weekly returns. Hence, the likelihood function is formally misspecified. 41

42 However, if the conditional mean and variance are correctly specified, the corre- 42  
43 sponding Quasi-Maximum Likelihood Estimates (QMLE) obtained under this auxiliary 43



assumption of conditional normality will generally be consistent for the true value of  $\theta$ . Moreover, asymptotically valid robust standard errors may be calculated from the so-called “sandwich-form” of the covariance matrix estimator, defined by the outer product of the gradients post- and pre-multiplied by the inverse of the usual information matrix estimator. Since the expressions for the future conditional variances for most of the GARCH models discussed above do not depend upon the actual distribution of  $z_t$ , as long as  $E(z_t | \mathcal{F}_{t-1}) = 0$  and  $E(z_t^2 | \mathcal{F}_{t-1}) = 1$ , this means that asymptotically valid point volatility forecasts may be constructed from the conditionally normal QMLE for  $\theta$  without fully specifying the distribution of  $z_t$ .

Still, the efficiency of the parameter estimates, and therefore the accuracy of the resulting point volatility forecasts obtained by simply substituting  $\hat{\theta}_T$  in place of the unknown parameters in the forecasting formulas, may be improved by employing the correct conditional distribution of  $z_t$ . A standardized Student  $t$  distribution with degrees of freedom  $\nu > 2$  often provides a good approximation to this distribution. Specifically,

$$f(y_t | \mathcal{F}_{t-1}) = \Gamma\left(\frac{\nu+1}{2}\right) \Gamma\left(\frac{\nu}{2}\right)^{-1} ((\nu-2)\sigma_{t|t-1}^2)^{-1/2} \\ \times (1 + (\nu-2)^{-1}\sigma_{t|t-1}^{-2}(y_t - \mu_{t|t-1})^2)^{-(\nu+1)/2} \quad (3.25)$$

with the log-likelihood function given by the sum of the corresponding  $T$  logarithmic densities, and the degrees of freedom parameter  $\nu$  estimated jointly with the other parameters of the model entering the conditional mean and variance functions. Note, that for  $\nu \rightarrow \infty$  the distribution converges to the conditional normal density in (3.23). Of course, more flexible distributions allowing for both fat tails and asymmetries could be, and have been, employed as well. Additionally, semi-nonparametric procedures in which the parameters in  $\mu_{t|t-1}(\theta)$  and  $\sigma_{t|t-1}^2(\theta)$  are estimated sequentially on the basis of nonparametric kernel type estimates for the distribution of  $\hat{z}_t$  have also been developed to enhance the efficiency of the parameter estimates relative to the conditionally normal QMLEs. From a forecasting perspective, however, the main advantage of these more complicated conditionally nonnormal estimation procedures lies not so much in the enhanced efficiency of the plugin point volatility forecasts,  $\sigma_{T+h|T}^2(\hat{\theta}_T)$ , but rather in their ability to better approximate the tails in the corresponding predictive distributions,  $f(y_{T+h} | \mathcal{F}_T; \hat{\theta}_T)$ . We next turn to a discussion of this type of density forecasting.

### 3.6. Fat tails and multi-period forecast distributions

The ARCH class of models directly specifies the one-step-ahead conditional mean and variance,  $\mu_{t|t-1}$  and  $\sigma_{t|t-1}^2$ , as functions of the time  $t-1$  information set,  $\mathcal{F}_{t-1}$ . As such, the one-period-ahead predictive density for  $y_t$  is directly determined by the distribution of  $z_t$ . In particular, assuming that  $z_t$  is i.i.d. standard normal,

$$f_z(z_t) = (2\pi)^{-1/2} \exp(-1/2z_t^2),$$

the conditional density of  $y_t$  is then given by the expression in Equation (3.23) above, where the  $\sigma_{t|t-1}^{-1}$  term is associated with the Jacobian of the transformation from  $z_t$  to  $y_t$ . Thus, in this situation, the one-period-ahead VaR at level  $p$  is readily calculated by  $\text{VaR}_{t+1|t}^p = \mu_{t+1|t} + \sigma_{t+1|t} F_z^{-1}(p)$ , where  $F_z^{-1}(p)$  equals the  $p$ th quantile in the standard normal distribution.

Meanwhile, as noted above the distributions of the standardized GARCH innovations often have fatter tails than the normal distribution. To accommodate this feature alternative conditional error distributions, such as the Student  $t$  distribution in Equation (3.25) discussed above, may be used in place of the normal density in Equation (3.23) in the construction of empirically more realistic predictive densities. In the context of quantile predictions, or VARs, this translates into multiplication factors,  $F_z^{-1}(p)$ , in excess of those for the normal distribution for small values of  $p$ . Of course, the exact value of  $F_z^{-1}(p)$  will depend upon the specific parametric estimates for the distribution of  $z_t$ . Alternatively, the standardized in-sample residuals based on the simpler-to-implement QMLE for the parameters, say  $\hat{z}_t \equiv \hat{\sigma}_{t|t-1}^{-1}(y_t - \hat{\mu}_{t|t-1})$ , may be used in nonparametrically estimating the distribution of  $z_t$ , and in turn the quantiles,  $\hat{F}_z^{-1}(p)$ .

The procedures discussed above generally work well in approximating VARs within the main range of support of the distribution, say  $0.01 < p < 0.99$ . However, for quantiles in the very far left or right tail, it is not possible to meaningfully estimate  $F_z^{-1}(p)$  without imposing some additional structure on the problem. Extreme Value Theory (EVT) provides a framework for doing so. In particular, it follows from EVT that under general conditions the tails of any admissible distribution must behave like those of the Generalized Pareto class of distributions. Hence, provided that  $z_t$  is i.i.d., the extreme quantiles in  $f(y_{t+1} | \mathcal{F}_t)$  may be inferred exactly as above, using only the  $[rT]$  smallest (largest) values of  $\hat{z}_t$  in actually estimating the parameters of the corresponding extreme value distribution used in calculating  $\hat{F}_z^{-1}(p)$ . The fraction  $r$  of the full sample  $T$  used in this estimation dictates where the tails, and consequently the extreme value distribution, begin. In addition to standard MLE techniques, a number of simplified procedures, including the popular Hill estimator, are also available for estimating the required tail parameters.

The calculation of multi-period forecast distributions is more complicated. To facilitate the presentation, suppose that the information set defining the conditional one-step-ahead distribution,  $f(y_{t+1} | \mathcal{F}_t)$ , and consequently the conditional mean and variance,  $\mu_{t+1|t}$  and  $\sigma_{t+1|t}^2$ , respectively, is restricted to current and past values of  $y_t$ . The multi-period-ahead predictive distribution is then formally defined by the convolution of the corresponding  $h$  one-step-ahead distributions,

$$f(y_{t+h} | \mathcal{F}_t) = \iiint \dots \int f(y_{t+h} | \mathcal{F}_{t+h-1}) f(y_{t+h-1} | \mathcal{F}_{t+h-2}) \dots f(y_{t+1} | \mathcal{F}_t) dy_{t+h-1} dy_{t+h-2} \dots dy_{t+1}. \quad (3.26)$$

This multi-period mixture distribution generally has fatter tails than the underlying one-step-ahead distributions. In particular, assuming that the one-step-ahead distributions are conditionally normal as in (3.23) then, if the limiting value exists, the unconditional

distribution,  $f(y_t) = \lim_{h \rightarrow \infty} f(y_t | \mathcal{F}_{t-h})$ , will be leptokurtic relative to the normal. This is, of course, entirely consistent with the unconditional distribution of most speculative returns having fatter tails than the normal. It is also worth noting that even though the conditional one-step-ahead predictive distributions,  $f(y_{t+1} | \mathcal{F}_t)$ , may be symmetric, if the conditional variance depends on the past values of  $y_t$  in an asymmetric fashion, as in the GJR, AGARCH or EGARCH models, the multi-step-ahead distribution,  $f(y_{t+h} | \mathcal{F}_t)$ ,  $h > 1$ , will generally be asymmetric. Again, this is directly in line with the negative skewness observed in the unconditional distribution of most equity index return series.

Despite these general results, analytical expressions for the multi-period predictive density in (3.26) are not available in closed-form. However, numerical techniques may be used in recursively building up an estimate for the predictive distribution, by repeatedly drawing future values for  $y_{t+j} = \mu_{t+j|t+j-1} + \sigma_{t+j|t+j-1}z_{t+j}$  based on the assumed parametric distribution  $f_z(z_t)$ , or by bootstrapping  $z_{t+j}$  from the in-sample distribution of the standardized residuals.

Alternatively,  $f(y_{t+h} | \mathcal{F}_t)$  may be approximated by a time-invariant parametric or nonparametrically estimated distribution with conditional mean and variance,  $\mu_{t+h|t} \equiv E(y_{t+j} | \mathcal{F}_t)$  and  $\sigma_{t+h|t}^2 \equiv \text{Var}(y_{t+j} | \mathcal{F}_t)$ , respectively. The multi-step conditional variance is readily calculated along the lines of the recursive prediction formulas discussed in the preceding sections. This approach obviously neglects any higher order dependencies implied by the convolution in (3.26). However, in contrast to the common approach of scaling which, as illustrated in Figure 5, may greatly exaggerate the volatility-of-the-volatility, the use of the correct multi-period conditional variance means that this relatively simple-to-implement approach for calculating multi-period predictive distributions usually works very well in practice.

The preceding discussion has focused on one or multi-period forecast distributions spanning the identical unit time interval as in the underlying GARCH model. However, as previously noted, in financial applications the forecast distribution of interest often involves the sum of  $y_{t+j}$  over multiple periods corresponding to the distribution of continuously compounded multi-period returns, say  $y_{t:t+h} \equiv y_{t+h} + y_{t+h-1} + \dots + y_{t+1}$ . The same numerical techniques used in approximating  $f(y_{t+h} | \mathcal{F}_t)$  by Monte Carlo simulations discussed above may, of course, be used in approximating the corresponding distribution of the sum,  $f(y_{t:t+h} | \mathcal{F}_t)$ .

Alternatively, assuming that the  $y_{t+j}$ 's are serially uncorrelated, as would be approximately true for most speculative returns over daily or weekly horizons, the conditional variance of  $y_{t:t+h}$  is simply equal to the sum of the corresponding  $h$  variance forecasts,

$$\text{Var}(y_{t:t+h} | \mathcal{F}_t) \equiv \sigma_{t:t+h|t}^2 = \sigma_{t+h|t}^2 + \sigma_{t+h-1|t}^2 + \dots + \sigma_{t+1|t}^2. \tag{3.27}$$

Thus, in this situation the conditional distribution of  $y_{t:t+h}$  may be estimated on the basis of the corresponding in-sample standardized residuals,  $\hat{z}_{t:t+h} \equiv \hat{\sigma}_{t:t+h|t}^{-1}(y_{t:t+h} - \hat{\mu}_{t:t+h|t})$ . Now, if the underlying GARCH process for  $y_t$  is covariance stationary, we have  $\lim_{h \rightarrow \infty} h^{-1}\mu_{t:t+h} = E(y_t)$  and  $\lim_{h \rightarrow \infty} h^{-1}\sigma_{t:t+h}^2 = \text{Var}(y_t)$ . Moreover, as shown by Diebold (1988), it follows from a version of the standard Central Limit

1 Theorem that  $z_{t:t+h} \Rightarrow N(0, 1)$ . Thus, volatility clustering disappears under tempo- 1  
 2 ral aggregation, and the unconditional return distributions will be increasingly better 2  
 3 approximated by a normal distribution the longer the return horizons. This suggests that 3  
 4 for longer-run forecasts, or moderately large values of  $h$ , the distribution of  $z_{t:t+h}$  will be 4  
 5 approximately normal. Consequently, the calculation of longer-run multi-period VARs 5  
 6 may reasonably rely on the conventional quantiles from a standard normal probability 6  
 7 table in place of  $F_z^{-1}(p)$  in the formula  $\text{VaR}_{t:t+h|t}^P = \mu_{t:t+h|t} + \sigma_{t:t+h|t} F_z^{-1}(p)$ . 7

### 8 3.7. Further reading 9

10  
 11 The ARCH and GARCH class of models have been extensively surveyed elsewhere; 11  
 12 see, e.g., review articles by Andersen and Bollerslev (1998b), Bollerslev, Chou and Kro- 12  
 13 ner (1992), Bollerslev, Engle and Nelson (1994), Diebold (2004), Diebold and Lopez 13  
 14 (1995), Engle (2001, 2004), Engle and Patton (2001), Pagan (1996), Palm (1996), and 14  
 15 Shephard (1996). The models have now also become part of the standard toolbox dis- 15  
 16 cussed in econometrics and empirical oriented finance textbooks; see, e.g., Hamilton 16  
 17 (1994), Mills (1993), Franses and van Dijk (2000), Gouriéroux and Jasiak (2001), 17  
 18 Alexander (2001), Brooks (2002), Chan (2002), Tsay (2002), Christoffersen (2003), 18  
 19 Enders (2004), and Taylor (2004). A series of the most influential early ARCH papers 19  
 20 have been collected in Engle (1995). A fairly comprehensive list as well as forecast 20  
 21 comparison of the most important parametric formulations are also provided in Hansen 21  
 22 and Lunde (2005). 22

23 Several different econometric and statistical software packages are available for es- 23  
 24 timating all of the most standard univariate GARCH models, including EViews, PC- 24  
 25 GIVE, Limdep, Microfit, RATS, S+, SAS, SHAZAM, and TSP. The open-ended ma- 25  
 26 trix programming environments GAUSS, Matlab, and Ox also offer easy add-ons for 26  
 27 GARCH estimation, while the NAG library and the UCSD Department of Economics 27  
 28 website provide various Fortran based procedures and programs. Partial surveys and 28  
 29 comparisons of some of these estimation packages and procedures are given in Brooks 29  
 30 (1997), Brooks, Burke and Persaud (2001), and McCullough and Renfro (1998). 30

31 The asymmetry, or “leverage” effect, directly motivating a number of the alternative 31  
 32 GARCH formulations were first documented empirically by Black (1976) and Christie 32  
 33 (1982). In addition to the papers by Nelson (1991), Engle and Ng (1993), Glosten, 33  
 34 Jagannathan and Runkle (1993), and Zakoian (1994) discussed in Section 3.3, other im- 34  
 35 portant studies on modeling and understanding the volatility asymmetry in the GARCH 35  
 36 context include Campbell and Hentschel (1992), Hentschel (1995), and Bekaert and Wu 36  
 37 (2000), while Engle (2001) provides an illustration of the importance of incorporating 37  
 38 asymmetry in GARCH-based VaR calculations. 38

39 The long-memory FIGARCH model of Baillie, Bollerslev and Mikkelsen (1996) in 39  
 40 Section 3.4 may be seen as a special case of the ARCH( $\infty$ ) model in Robinson (1991). 40  
 41 The FIGARCH model also encompasses the IGARCH model of Engle and Bollerslev 41  
 42 (1986) for  $d = 1$ . However, even though the approach discussed here affords a con- 42  
 43 venient framework for generating point forecasts with long-memory dependencies, when 43

viewed as a model the unconditional variance does not exist, and the FIGARCH class of models has been criticized accordingly by Giraitis, Kokoszka and Leipus (2000), among others. An alternative formulation which breaks the link between the conditions for second-order stationarity and long-memory dependencies have been proposed by Davidson (2004). Alternative long-memory GARCH formulations include the FIE-GARCH model of Bollerslev and Mikkelsen (1996), and the model in Ding and Granger (1996) based on the superposition of an infinite number of ARCH models. In contrast, the component GARCH model in Engle and Lee (1999) and the related developments in Gallant, Hsu and Tauchen (1999) and Müller et al. (1997), is based on the mixture of only a few components; see also the earlier related results on modeling and forecasting long-run dynamic dependencies in the mean by O'Connell (1971) and Tiao and Tsay (1994). Meanwhile, Bollerslev and Mikkelsen (1999) have argued that when pricing very long-lived financial contracts, the fractionally integrated volatility approach can result in materially different prices from the ones implied by the more standard GARCH models with exponential decay. The multifractal models recently advocated by Calvet and Fisher (2002, 2004) afford another approach for incorporating long memory into volatility forecasting.

Long memory also has potential links to regimes and structural break in volatility. Diebold and Inoue (2001) argue that the apparent finding of long memory could be due to the existence of regime switching. Mikosch and Starica (2004) explicitly uses non-stationarity as a source of long memory in volatility. Structural breaks in volatility is considered by Andreou and Ghysels (2002), Lamoureux and Lastrapes (1990), Pastor and Stambaugh (2001), and Schwert (1989). Hamilton and Lin (1996) and Perez-Quiros and Timmermann (2000) study volatility across business cycle regimes. The connections between long memory and structural breaks are reviewed in Banerjee and Urga (2005); see also Chapter 12 by Clements and Hendry in this Handbook.

Early contributions concerning the probabilistic and statistical properties of GARCH models, as well as the MLE and QMLE techniques discussed in Section 3.5, include Bollerslev and Wooldridge (1992), Lee and Hansen (1994), Lumsdaine (1996), Nelson (1990), and Weiss (1986); for a survey of this literature see also Li, Ling and McAleer (2002). Bollerslev (1986) discusses conditions for existence of the second moment in the specific context of the GARCH model. Loretan and Phillips (1994) contains a more general discussion on the issue of covariance stationarity. Bayesian methods for estimating ARCH models were first implemented by Geweke (1989a) and they have since be developed further in Bauwens and Lubrano (1998, 1999). The GARCH- $t$  model discussed in Section 3.5 was first introduced by Bollerslev (1987), while Nelson (1991) suggested the so-called Generalized Error Distribution (GED) for better approximating the distribution of the standardized innovations. Engle and Gonzalez-Rivera (1991) first proposed the use of kernel-based methods for nonparametrically estimating the conditional distribution, whereas McNeil and Frey (2000) relied on Extreme Value Theory (EVT) for estimating the uppermost tails in the conditional distribution; see also Embrechts, Klüppelberg and Mikosch (1997) for a general discussion of extreme value theory.

1 As discussed in Section 3.6, even if the one-step-ahead conditional distribution is 1  
2 known (by assumption), the corresponding multi-period distributions are not avail- 2  
3 able in closed-form and are generally unknown. Some of the complications that arise 3  
4 in this situation have been discussed in Baillie and Bollerslev (1992), who also con- 4  
5 sider the use of a Cornish–Fisher expansion for approximating specific quantiles in the 5  
6 multi-step-ahead predictive distributions. Numerical techniques for calculating the pre- 6  
7 dictive distributions based on importance sampling schemes were first implemented by 7  
8 Geweke (1989b). Other important results related to the distribution of temporally ag- 8  
9 gregated GARCH models include Drost and Nijman (1993), Drost and Werker (1996), 9  
10 and Meddahi and Renault (2004). 10  
11  
12

#### 13 4. Stochastic volatility 13 14

15 This section introduces the general class of models labeled Stochastic Volatility (SV). 15  
16 In the widest sense of the term, SV models simply allow for a stochastic element in 16  
17 the time series evolution of the conditional variance process. For example, GARCH 17  
18 models are SV models. The more meaningful categorization, which we adopt here, is 18  
19 to contrast ARCH type models with *genuine* SV models. The latter explicitly includes 19  
20 an unobserved (nonmeasurable) shock to the return variance into the characterization of 20  
21 the volatility dynamics. In this scenario, the variance process becomes inherently latent 21  
22 so that – even conditional on all past information and perfect knowledge about the data 22  
23 generating process – we cannot recover the exact value of the current volatility state. 23  
24 The technical implication is that the volatility process is not measurable with respect 24  
25 to observable (past) information. Hence, the assessment of the volatility state at day  $t$  25  
26 changes as contemporaneous or future information from days  $t + j$ ,  $j \geq 0$ , is incorpo- 26  
27 rated into the analysis. This perspective renders estimation of latent variables from past 27  
28 data alone (filtering) as well as from all available, including future, data (smoothing) 28  
29 useful. In contrast, GARCH models treat the conditional variance as observable given 29  
30 past information and, as discussed above, typically applies (quasi-) maximum likelihood 30  
31 techniques for inference, so smoothing has no role in that setting. 31

32 Despite these differences, the two model classes are closely related, and we consider 32  
33 them to be complementary rather than competitors. In fact, from a practical forecasting 33  
34 perspective it is hard to distinguish the performance of standard ARCH and SV mod- 34  
35 els. Hence, even if one were to think that the SV framework is appealing, the fact that 35  
36 ARCH models typically are easier to estimate explains practitioners reliance on ARCH 36  
37 as the volatility forecasting tool of choice. Nonetheless, the development of power- 37  
38 ful method of simulated moments, Markov Chain Monte Carlo (MCMC) and other 38  
39 simulation based procedures for estimation and forecasting of SV models may render 39  
40 them competitive with ARCH over time. Moreover, the development of the concept of 40  
41 realized volatility and the associated use of intraday data for volatility measurement, 41  
42 discussed in the next section, is naturally linked to the continuous-time SV framework 42  
43 of financial economics. 43

The literature on SV models is vast and rapidly growing, and excellent surveys are already available on the subject, e.g., Ghysels, Harvey and Renault (1996) and Shephard (1996, 2004). Consequently, we focus on providing an overview of the main approaches with particular emphasis on the generation of volatility forecasts within each type of model specification and inferential technique.

#### 4.1. Model specification

Roughly speaking, there are two main perspectives behind the SV paradigm when used in the context of modeling financial rate of returns. Although both may be adapted to either setting, there are precedents for one type of reasoning to be implemented in discrete time and the other to be cast in continuous time. The first centers on the Mixture of Distributions Hypothesis (MDH), where returns are governed by an event time process that represents a transformation of the time clock in accordance with the intensity of price relevant news, dating back to Clark (1973). The second approach stems from financial economics where the price and volatility processes often are modeled separately via continuous sample path diffusions governed by stochastic differential equations. We briefly introduce these model classes and point out some of the similarities to ARCH models in terms of forecasting procedures. However, the presence of a latent volatility factor renders both the estimation and forecasting problem more complex for the SV models. We detail these issues in the following subsections.

##### 4.1.1. The mixture-of-distributions hypothesis

Adopting the rational perspective that asset prices reflect the discounted value of future expected cash flows, such prices should react almost continuously to the myriad of news that arrive on a given trading day. Assuming that the number of news arrival is large, one may expect a central limit theory to apply and financial returns should be well approximated by a conditional normal distribution with the conditioning variable corresponding to the number of relevant news events. More generally, a number of other variables associated with the overall activity of the financial market such as the daily number of trades, the daily cumulative trading volume or the number of quotes may well be similarly related to the information flow in the market. These considerations inspire the following type of representation,

$$y_t = \mu_y s_t + \sigma_y s_t^{1/2} z_t, \quad (4.1)$$

where  $y_t$  is the market “activity” variable under consideration,  $s_t$  is the strictly positive process reflecting the intensity of relevant news arrivals,  $\mu_y$  represents the mean response of the variable per news event,  $\sigma_y$  is a scale parameter, and  $z_t$  is i.i.d.  $N(0, 1)$ . Equivalently, this relationship may be written as

$$y_t | s_t \sim N(\mu_y s_t, \sigma_y^2 s_t). \quad (4.2)$$

1 This formulation constitutes a normal mixture model. If the  $s_t$  process is time-varying 1  
 2 it induces a fat-tailed unconditional distribution, consistent with stylized facts for most 2  
 3 return and trading volume series. Intuitively, days with high information flow display 3  
 4 more price fluctuations and activity than days with fewer news releases. Moreover, if 4  
 5 the  $s_t$  process is positively correlated, then shocks to the conditional mean and variance 5  
 6 process for  $y_t$  will be persistent. This is consistent with the observed activity clustering 6  
 7 in financial markets, where return volatility, trading volume, the number of transactions 7  
 8 and quotes, the number of limit orders submitted to the market, etc., all display pro- 8  
 9 nounced serial dependence. 9

10 The specification in (4.1) is analogous to the one-step-ahead decomposition given 10  
 11 in Equation (3.5). The critical difference is that the formulation is endowed with a 11  
 12 structural interpretation, implying that the mean and variance components cannot be 12  
 13 observed prior to the trading day as the number of news arrivals is inherently random. 13  
 14 In fact, it is usually assumed that the  $s_t$  process is unobserved by the econometrician, 14  
 15 even during period  $t$ , so that the true mean and variance series are both latent. From a 15  
 16 technical perspective this implies that we must distinguish between the full information 16  
 17 set ( $s_t \in \mathcal{F}_t$ ) and observable information ( $s_t \notin \mathfrak{S}_t$ ). The latter property is a defining 17  
 18 feature of the genuine volatility class. The inability to observe this important component of 18  
 19 the MDH model complicates inference and forecasting procedures as discussed below. 19

20 In the case of short horizon return series,  $\mu_y$  is close to negligible and can reasonably 20  
 21 be ignored or simply fixed at a small constant value. Furthermore, if the mixing variable 21  
 22  $s_t$  is latent then the scaling parameter,  $\sigma_y$ , is not separately identified and may be fixed 22  
 23 at unity. This produces the following return (innovation) model, 23

$$24 \quad r_t = s_t^{1/2} z_t, \quad (4.3) \quad 24$$

25 implying a simple normal-mixture representation, 25

$$26 \quad r_t | s_t \sim N(0, s_t). \quad (4.4) \quad 26$$

27 Both univariate models for returns of the form (4.4) or multivariate systems including 27  
 28 a return variable along with other related market activity variables, such as trading 28  
 29 volume or the number of transactions, are referred to as derived from the Mixture-of- 29  
 30 Distributions Hypothesis (MDH). 30  
 31 31

32 The representation in (4.3) is of course directly comparable to that for the return 32  
 33 innovation in Equation (3.5). It follows immediately that volatility forecasting is related 33  
 34 to forecasts of the latent volatility factor given the observed information, 34  
 35 35

$$36 \quad \text{Var}(r_{t+h} | \mathfrak{S}_t) = E(s_{t+h} | \mathfrak{S}_t). \quad (4.5) \quad 36$$

37 If some relevant information is not observed and thus not included in  $\mathfrak{S}_t$ , then the ex- 37  
 38 pression in (4.5) will generally not represent the actual conditional return variance, 38  
 39  $E(s_{t+h} | \mathcal{F}_t)$ . This point is readily seen through a specific example. 39

40 In particular, Taylor (1986) first introduced the log-SV model by adopting an autore- 40  
 41 gressive parameterization of the latent log-volatility (or information flow) variable, 41  
 42 42

$$43 \quad \log s_{t+1} = \eta_0 + \eta_1 \log s_t + u_t, \quad u_t \sim \text{i.i.d.} (0, \sigma_u^2), \quad (4.6) \quad 43$$



1 where the disturbance term may be correlated with the innovation in the return equation, 1  
 2 that is,  $\rho = \text{corr}(u_t, z_t) \neq 0$ . This particular representation, along with a Gaussian 2  
 3 assumption on  $u_t$ , has been so widely adopted that it has come to be known as *the* 3  
 4 stochastic volatility model. Note that, if  $\rho$  is negative, there is an asymmetric return- 4  
 5 volatility relationship present in the model, akin to the “leverage effect” in the GJR and 5  
 6 EGARCH models discussed in Section 3.3, so that negative return shocks induce higher 6  
 7 future volatility than similar positive shocks. In fact, it is readily seen that the log-SV 7  
 8 formulation in (4.6) generalizes the EGARCH(1, 1) model by considering the case, 8  
 9

$$10 \quad u_t = \alpha(|z_t| - E|z_t|) + \gamma z_t, \quad (4.7) \quad 10$$

11 where the parameters  $\eta_0$  and  $\eta_1$  correspond to  $\omega$  and  $\beta$  in Equation (3.15), respectively. 12  
 13 Under the null hypothesis of EGARCH(1, 1), the information set,  $\mathfrak{S}_t$ , includes past 13  
 14 asset returns, and the idiosyncratic return innovation series,  $z_t$ , is effectively observable 14  
 15 so likelihood based analysis is straightforward. However, if  $u_t$  is not (only) a function 15  
 16 of  $z_t$ , i.e., Equation (4.7) no longer holds, then there are two sources of error in the 16  
 17 system. In this more general case it is no longer possible to separately identify the 17  
 18 underlying innovations to the return and volatility processes, nor the true underlying 18  
 19 volatility state. 19

20 This above example illustrates both how any ARCH model may be seen as a special 20  
 21 case of a corresponding SV model and how the defining feature of the genuine SV 21  
 22 model may complicate forecasting, as the volatility state is unobserved. Obviously, in 22  
 23 representations like (4.6), the current state of volatility is a critical ingredient for fore- 23  
 24 casts of future volatility. We expand on the tasks confronting estimation and volatility 24  
 25 forecasting in this setting in Section 4.1.3. 25

26 There are, of course, an unlimited number of alternative specifications that may be 26  
 27 entertained for the latent volatility process. However, Stochastic Autoregressive Volatil- 27  
 28 ity (SARV) of Andersen (1994) has proven particularly convenient. The representation is 28  
 29 again autoregressive, 29

$$30 \quad v_t = \omega + \beta v_{t-1} + [\gamma + \alpha v_{t-1}] u_t, \quad (4.8) \quad 31$$

32 where  $u_t$  denotes an i.i.d. sequence, and  $s_t = g(v_t)$  links the dynamic evolution of the 32  
 33 state variable to the stochastic variance factor in Equation (4.3). For example, for the 33  
 34 log-SV model,  $g(v_t) = \exp(v_t)$ . Likewise, SV generalizations of the GARCH(1, 1) 34  
 35 may be obtained via  $g(v_t) = v_t$  and an SV extension of a GARCH model for the 35  
 36 conditional standard deviation is produced by letting  $g(v_t) = v_t^{1/2}$ . Depending upon 36  
 37 the specific transformation  $g(\cdot)$  it may be necessary to impose additional (positivity) 37  
 38 constraints on the innovation sequence  $u_t$ , or the parameters in (4.8). Even if inference 38  
 39 on parameters can be done, moment based procedures do not produce estimates of the 39  
 40 latent volatility process, so from a forecasting perspective the analysis must necessarily 40  
 41 be supplemented with some method of approximating the sample path realization of the 41  
 42 underlying state variables. 42  
 43

### 4.1.2. Continuous-time stochastic volatility models

The modeling of asset returns in continuous time stems from the financial economics literature where early contributions to portfolio selection by Merton (1969) and option pricing by Black and Scholes (1973) demonstrated the analytical power of the diffusion framework in handling dynamic asset allocation and pricing problems. The idea of casting these problems in a continuous-time diffusion context also has a remarkable precedent in Bachelier (1900).

Under weak regularity conditions, the general representation of an arbitrage-free asset price process is

$$dp(t) = \mu(t) dt + \sigma(t) dW(t) + j(t) dq(t), \quad t \in [0, T], \quad (4.9)$$

where  $\mu(t)$  is a continuous, locally bounded variation process, the volatility process  $\sigma(t)$  is strictly positive,  $W(t)$  denotes a standard Brownian motion,  $q(t)$  is a jump indicator taking the values zero (no jump) or unity (jump) and, finally, the  $j(t)$  represents the size of the jump if one occurs at time  $t$ . [See, e.g., Andersen, Bollerslev and Diebold (2003a) for further discussion.] The associated one-period return is

$$\begin{aligned} r(t) &= p(t) - p(t-1) \\ &= \int_{t-1}^t \mu(\tau) d\tau + \int_{t-1}^t \sigma(\tau) dW(\tau) + \sum_{t-1 \leq \tau < t} \kappa(\tau), \end{aligned} \quad (4.10)$$

where the last sum simply cumulates the impact of the jumps occurring over the period, as we define  $\kappa(t) = j(t) \cdot \mathbf{I}(q(t) = 1)$ , so that  $\kappa(t)$  is zero everywhere except when a discrete jump occurs.

In this setting a formal ex-post measure of the return variability, derived from the theory of quadratic variation for semi-martingales, may be defined as

$$QV(t) \equiv \int_{t-1}^t \sigma^2(s) ds + \sum_{t-1 < s \leq t} \kappa^2(s). \quad (4.11)$$

In the special case of a pure SV diffusion, the corresponding quantity reduces to the integrated variance, as already defined in Equation (1.11) in Section 1,

$$IV(t) \equiv \int_{t-1}^t \sigma^2(s) ds. \quad (4.12)$$

These return variability measures are naturally related to the return variance. In fact, for a pure SV diffusion (without jumps) where the volatility process,  $\sigma(\tau)$ , is independent of the Wiener process,  $W(\tau)$ , we have

$$r(t) | \{ \mu(\tau), \sigma(\tau); t-1 \leq \tau \leq t \} \sim N \left( \int_{t-1}^t \mu(\tau) d\tau, \int_{t-1}^t \sigma^2(\tau) d\tau \right), \quad (4.13)$$

so the integrated variance is the true measure of the actual (ex-post) return variance in this context. Of course, if the conditional variance and mean processes evolve stochastically we cannot perfectly predict the future volatility, and we must instead form expectations based on the current information. For short horizons, the conditional mean variation is negligible and we may focus on the following type of forecasts, for a positive integer  $h$ ,

$$\text{Var}(r(t+h) \mid \mathfrak{F}_t) \approx E \left[ \int_{t+h-1}^{t+h} \sigma^2(\tau) d\tau \mid \mathfrak{F}_t \right] \equiv E[\text{IV}(t+h) \mid \mathfrak{F}_t]. \tag{4.14}$$

The expressions in (4.13) and (4.14) generalize the corresponding equations for discrete-time SV models in (4.4) and (4.5), respectively. Of course, the return variation arising from the conditional mean process may need to be accommodated as well over longer horizons. Nonetheless, the dominant term in the return variance forecast will invariably be associated with the expected integrated variance or, more generally, the expected quadratic variation. In simple continuous-time models, we may be able to derive closed-form expressions for these quantities, but in empirically realistic settings they are typically not available in analytic form and alternative procedures must be used. We discuss these issues in more detail below.

The initial diffusion models explored in the literature were not genuine SV diffusions but rather, with a view toward tractability, cast as special cases of the constant elasticity of variance (CEV) class of models,

$$dp(t) = (\mu - \phi[p(t) - \mu]) dt + \sigma p(t)^\gamma dW(t), \quad t \in [0, T], \tag{4.15}$$

where  $\phi \geq 0$  determines the strength of mean reversion toward the unconditional mean  $\mu$  in the log-price process, while  $\gamma \geq 0$  allows for conditional heteroskedasticity in the return process. Popular representations are obtained by specific parameter restrictions, e.g., the Geometric Brownian motion for  $\phi = 0$  and  $\gamma = 0$ , the Vasicek model for  $\gamma = 0$ , and the Cox-Ingersoll and Ross (CIR) or square-root model for  $\gamma = \frac{1}{2}$ . These three special cases allow for a closed-form characterization of the likelihood, so the analysis is straightforward. Unfortunately, they are also typically inadequate in terms of capturing the volatility dynamics of asset returns. A useful class of extensions have been developed from the CIR model. In this model the instantaneous mean and variance processes are both affine functions of the log price. The affine model class extends the above representation with  $\gamma = \frac{1}{2}$  to a multivariate setting with general affine conditional mean and variance specifications. The advantage is that a great deal of analytic tractability is retained while allowing for more general and empirically realistic dynamic features.

Many genuine SV representations of empirical interest fall outside of the affine class, however. For example, Hull and White (1987) develop a theory for option pricing under stochastic volatility using a model much in the spirit of Taylor’s discrete-time log SV in Equation (4.6). With only a minor deviation from their representation, we may write it, for  $t \in [0, T]$ ,

$$dp(t) = \mu(t) dt + \sigma(t) dW(t), \quad (4.16)$$

$$d \log \sigma^2(t) = \beta(\alpha - \log \sigma^2(t)) dt + v dW_\sigma(t).$$

The strength of the mean reversion in (log) volatility is given by  $\beta$  and the volatility is governed by  $v$ . Positive but low values of  $\beta$  induces a pronounced volatility persistence, while large values of  $v$  increase the idiosyncratic variation in the volatility series. Furthermore, the log transform implies that the volatility of volatility rises with the level of volatility, even if  $v$  is time invariant. Finally, a negative correlation,  $\rho < 0$ , between the Wiener processes  $W(t)$  and  $W_\sigma(t)$  will induce an asymmetric return-volatility relationship in line with the leverage effect discussed earlier. As such, these features allow the representation in (4.16) to capture a number of stylized facts about asset return series quite parsimoniously.

Another popular nonaffine specification is the GARCH diffusion analyzed by [Drost and Werker \(1996\)](#). This representation can formally be shown to induce a GARCH type behavior for any discretely sampled price series and it is therefore a nice framework for eliciting and assessing information about the volatility process through data gathered at different sampling frequencies. This is also the process used in the construction of [Figure 1](#). It takes the form

$$dp(t) = \mu dt + \sigma(t) dW(t), \quad (4.17)$$

$$d\sigma^2(t) = \beta(\alpha - \sigma^2(t)) dt + v\sigma^2(t) dW_\sigma(t),$$

where the two Wiener processes are now independent.

The SV diffusions in (4.16) and (4.17) are but simple examples of the increasingly complex multi-factor (affine as well as nonaffine) jump-diffusions considered in the literature. Such models are hard to estimate by standard likelihood or method of moments techniques. This renders their use in forecasting particularly precarious. There is a need for both reliable parameter estimates and reliable extraction of the values for the underlying state variables. In particular, the current value of the state vector (and thus volatility) constitutes critical conditioning information for volatility prediction. The usefulness of such specifications for volatility forecasting is therefore directly linked to the availability of efficient inference methods for these models.

#### 4.1.3. Estimation and forecasting issues in SV models

The incorporation of a latent volatility process in SV models has two main consequences. First, estimation cannot be performed through a direct application of maximum likelihood principles. Many alternative procedures will involve an efficiency loss relative to this benchmark so model parameter uncertainty may then be larger. Since forecasting is usually made conditional on point estimates for the parameters, this will tend to worsen the predictive ability of model based forecasts. Second, since the current state for volatility is not observed, there is an additional layer of uncertainty surrounding forecasts made conditional on the estimated state of volatility. We discuss these issues

below and the following sections then review two alternative estimation and forecasting procedures developed, in part, to cope with these challenges.

Formally, the SV likelihood function is given as follows. Let the vector of return (innovations) and volatilities over  $[0, T]$  be denoted by  $\underline{r} = (r_1, \dots, r_T)$  and  $\underline{s} = (s_1, \dots, s_T)$ , respectively. Collecting the parameters in the vector  $\theta$ , the probability density for the data given  $\theta$  may then be written as

$$\begin{aligned} f(\underline{r}; \theta) &= \int f(\underline{r}, \underline{s}; \theta) d\underline{s} = \prod_{t=1}^T f(r_t | \mathfrak{F}_{t-1}; \theta) \\ &= \prod_{t=1}^T \int f(r_t | s_t; \theta) f(s_t | \mathfrak{F}_{t-1}; \theta) ds_t. \end{aligned} \quad (4.18)$$

For parametric discrete-time SV models, the conditional density  $f(r_t | s_t, \theta)$  is typically known in closed form, but  $f(s_t | \mathfrak{F}_{t-1}; \theta)$  is not available. Without being able to utilize this decomposition, we face an integration over the full unobserved volatility vector which is a  $T$ -dimensional object and generally not practical to compute given the serial dependence in the latent volatility process.

The initial response to these problems was to apply alternative estimation procedures. In his original treatment Taylor (1986) uses moment matching. Later, Andersen (1994) shows that it is feasible to estimate a broad class of discrete-time SV models through standard GMM procedures. However, this is not particularly efficient as the unconditional moments that may be expressed in closed form are quite different from the (efficient) score moments associated with the (infeasible) likelihood function. Another issue with GMM estimates is the need to extract estimates of the state variables if it is to serve as a basis for volatility forecasting. GMM does not provide any direct identification of the state variables, so this must be addressed in a second step. In that regard, the Kalman filter was often used. This technique allows for sequential estimation of parameters and latent state variables. As such, it provides a conceptual basis for the analysis, even if the basic Kalman filter is inadequate for general nonlinear and non-Gaussian SV models.

Nelson (1988) first suggested casting the SV estimation problem in a state space setting. We illustrate the approach for the simplest version of the log-SV model without a leverage effect, that is,  $\rho = 0$  in (4.4) and (4.6). Now, squaring the expression in (4.3), taking logs and assuming Gaussian errors in the transition equation for the volatility state in Equation (4.6), it follows that

$$\begin{aligned} \log r_t^2 &= \log s_t + \log z_t^2, \quad z_t \sim \text{i.i.d. } N(0, 1), \\ \log s_{t+1} &= \eta_0 + \eta_1 \log s_t + u_t, \quad u_t \sim \text{i.i.d. } N(0, \sigma_u^2). \end{aligned}$$

To conform with standard notation, it is useful to consolidate the constant from the transition equation into the measurement equation for the log-squared return residual. Defining  $h_t \equiv \log s_t$ , we have

$$\begin{aligned} \log r_t^2 &= \omega + h_t + \xi_t, \quad \xi_t \sim \text{i.i.d. } (0, 4.93), \\ h_{t+1} &= \eta h_t + u_t, \quad u_t \sim \text{i.i.d. } N(0, \sigma_u^2), \end{aligned} \tag{4.19}$$

where  $\omega = \eta_0 + E(\log z_t^2) = \eta_0 - 1.27$ ,  $\eta = \eta_1$ , and  $\xi_t$  is a demeaned  $\log \chi^2$  distributed error term. The system in (4.19) is given in the standard linear state space format. The top equation provides the measurement equation where the squared return is linearly related to the latent underlying volatility state and an i.i.d. skewed and heavy tailed error term. The bottom equation provides the transition equation for the model and is given as a first-order Gaussian autoregression.

The Kalman filter applies directly to (4.19) by assuming Gaussian errors; see, e.g., Harvey (1989, 2006). However, the resultant estimators of the state variables and the future observations are only minimum mean-squared error for estimators that are linear combinations of past  $\log r_t^2$ . Moreover, the non-Gaussian errors in the measurement equation implies that the exact likelihood cannot be obtained from the associated prediction errors. Nonetheless, the Kalman filter may be used in the construction of QMLEs of the model parameters for which asymptotically valid inference is available, even if these estimates generally are fairly inefficient. Arguably, the most important insight from the state space representation is instead the inspiration it has provided for the development of more efficient estimation and forecasting procedures through nonlinear filtering techniques.

The state space representation directly focuses attention on the task of making inference regarding the latent state vector, i.e., for SV models the question of what we can learn about the current state of volatility. A comprehensive answer is provided by the solution to the *filtering problem*, i.e., the distribution of the state vector given the current information set,  $f(s_t | \mathfrak{S}_t; \theta)$ . Typically, this distribution is critical in obtaining the one-step-ahead volatility forecast,

$$f(s_t | \mathfrak{S}_{t-1}; \theta) = \int f(s_t | s_{t-1}; \theta) f(s_{t-1} | \mathfrak{S}_{t-1}; \theta) ds_{t-1}, \tag{4.20}$$

where the first term in the integral is obtained directly from the transition equation in the state space representation. Once the one-step-ahead distribution has been determined, the task of constructing multiple-step-ahead forecasts is analogous to the corresponding problem under ARCH models where multi-period forecasts also generally depend upon the full distributional characterization of the model. A unique feature of the SV model is instead the *smoothing problem*, related to ex-post inference regarding the in-sample volatility given the set of observed returns over the full sample,  $f(s_t | \mathfrak{S}_T; \theta)$ , where  $t \leq T$ . At the end of the sample, either the filtering or smoothing solution can serve as the basis for out-of-sample volatility forecasts (for  $h$  a positive integer),

$$f(s_{T+h} | \mathfrak{S}_T; \theta) = \int f(s_{T+h} | s_T; \theta) f(s_T | \mathfrak{S}_T; \theta) ds_T, \tag{4.21}$$

where, again, given the solution for  $h = 1$ , the problem of determining the multi-period forecasts is analogous to the situation with multi-period ARCH-based forecasts discussed in Section 3.6.

As noted, all of these conditional volatility distributions may in theory be derived in closed form under the linear Gaussian state space representation via the Kalman filter. Unfortunately, even the simplest SV model contains some non-Gaussian and/or nonlinear elements. Hence, standard filtering methods provide, at best, approximate solutions and they have generally been found to perform poorly in this setting, in turn necessitating alternative more specialized filtering and smoothing techniques. Moreover, we have deliberately focused on the discrete-time case above. For the continuous-time SV models, the complications are more profound as even the discrete one-period return distribution conditional on the initial volatility state typically is not known in closed form. Hence, not only is the last term on the extreme right of Equation (4.18) unknown, but the first term is also intractable, further complicating likelihood-based analysis. We next review two recent approaches that promise efficient inference more generally and also provide ways of extracting reliable estimates of the latent volatility state needed for forecasting purposes.

#### 4.2. *Efficient method of simulated moments procedures for inference and forecasting*

The Efficient Method of Moments (EMM) procedure is the prime example of a Method of Simulated Moments (MSM) approach that has the potential to deliver efficient inference and produce credible volatility forecasting for general SV models. The intuition behind EMM is that, by traditional likelihood theory, the scores (the derivative of the log likelihood with respect to the parameter vector) provide efficient estimating moments. In fact, maximum likelihood is simply a just-identified GMM estimator based on the score (moment) vector. Hence, intuitively, from an efficiency point of view, one would like to approximate the score vector when choosing the GMM moments. Since the likelihood of SV models is intractable, the approach is to utilize a semi-nonparametric approximation to the log likelihood estimated in a first step to produce the moments. Next, one seeks to match the approximating score moments with the corresponding moments from a long simulation of the SV model. Thus, the main requirement for applicability of EMM is that the model can be simulated effectively and the system is stationary so that the requisite moments can be computed by simple averaging from a simulation of the system. Again, this idea, like the MCMC approach discussed in the next section, is, of course, applicable more generally, but for concreteness we will focus on estimation and forecasting with SV models for financial rate of returns.

More formally, let the sample of discretely observed returns be given by  $\underline{r} = (r_1, r_2, \dots, r_T)$ . Moreover, let  $x_{t-1}$  denote the vector of relevant conditioning variables for the log-likelihood function at time  $t$ , and let  $\underline{x} = (x_0, x_1, \dots, x_{T-1})$ . For simplicity, we assume a long string of prior return observations are the only components of  $\underline{x}$ , but other predetermined variables from an extended dynamic representation of the system may be incorporated as well. In the terminology of Equation (4.18), the complication is that the likelihood contribution from the  $t$ th return is not available, that is,  $f(r_t | \mathfrak{F}_{t-1}; \theta) \equiv f(r_t | x_{t-1}; \theta)$  is unknown. The proposal is to instead approximate this density by a flexible semi-nonparametric (SNP) estimate using the full data

sample. Without going into specifics, an appropriate procedure may be developed to obtain a close approximation to this conditional density within a class of SNP densities which are analytically tractable and allow for explicit computation of the associated score vector. The leading term will typically consist of a GARCH type model. Essentially, the information regarding the probabilistic structure available from the data is being encoded into an empirically tractable SNP representation, so that, for a large enough sample, we have

$$g(r_t | x_{t-1}; \hat{\eta}_T) \approx f(r_t | \mathfrak{S}_{t-1}; \theta_0), \quad (4.22)$$

where  $g(r_t | x_{t-1}; \hat{\eta}_T)$  denotes the fitted SNP density evaluated at the (pseudo) maximum likelihood estimate  $\hat{\eta}_T$ , and  $\theta_0$  denotes the true (unknown) parameter vector of the model generating the data under the null hypothesis. In general, the functional form of  $g$  is entirely different from the unknown  $f$ , and hence there is no direct compatibility between the two parameter vectors  $\eta$  and  $\theta$ , although we require that the dimension of  $\eta$  is at least as large as that of  $\theta$ . Notice how this SNP representation sidesteps the lack of a tractable expression for the likelihood contribution as given by the middle term in the likelihood expression in (4.18). Although the SNP density is not used for formal likelihood estimation, it is used to approximate the “efficient” score moments.

By construction,  $\hat{\eta}_T$  satisfies a set of first order conditions for the pseudo log-likelihood function under the empirical measure induced by the data, that is, letting  $\underline{r}_t = (r_t, x_{t-1})$ , it holds that

$$\frac{1}{T} \sum_{t=1}^T \frac{\partial}{\partial \eta} \log g(r_t | x_{t-1}; \hat{\eta}_T) \equiv \frac{1}{T} \sum_{t=1}^T \psi_T(\underline{r}_t) = 0. \quad (4.23)$$

It is clear that (4.23) takes the form of (pseudo) score moments. This representation of the data through a set of (efficient) moment conditions is the key part of the “projection step” of EMM. The data structure has effectively been projected onto an analytically tractable class of SNP densities augmented, as appropriate, by a leading dynamic (GARCH) term.

Since we are working under the assumption that we have a good approximation to the underlying true conditional density, we would intuitively expect that, for  $T$  large,

$$E_{\theta_0}[\psi_T(\tilde{r})] \approx \frac{1}{M} \sum_{i=1}^M \psi_T(\tilde{r}_i) \approx \frac{1}{T} \sum_{t=1}^T \psi_T(\underline{r}_t) = 0, \quad (4.24)$$

and any large artificial sample,  $\tilde{r} = (\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_N, \tilde{x}_0, \tilde{x}_1, \dots, \tilde{x}_{M-1})$ , generated by the same assumed (true) data generating process,  $f(r_t | \mathfrak{S}_{t-1}; \theta_0)$ , that is behind the observed return data,  $\underline{r}$ . These conjectures are formalized by Gallant and Tauchen (1996), who show how the pseudo score moments obtained in (4.23) by fixing  $\hat{\eta}_T$  can serve as valid (and efficient) moment conditions for estimating the parameter of interest,  $\theta$ . Since no analytic expression for the expectation on the extreme left in (4.24) is available, they propose a simulation estimator where the expectation is approximated arbitrarily well



by a very large simulated sample moment ( $M \gg T$ ) from the true underlying model. The ability to practically eliminate the simulation error renders the EMM estimator (in theory) independent of simulation size,  $M$ , but the uncertainty associated with the projection step, for which the sample size is constrained by the actual data, remains and the estimator,  $\hat{\theta}_T$ , is asymptotically normal with standard errors that reflects the estimation uncertainty in (4.23).

An obvious attraction of the EMM technique, beyond the potential for efficient inference, is that there are almost no restrictions on the underlying parametric model apart from stationarity and the ability to be able to simulate effectively from the model. This implies that the procedure can be used for continuous-time processes, even if we only observe a set of discretely sampled data. A seemingly important drawback, however, is the lack of any implied estimates of the underlying latent state variables which are critical for successful forecasting. Gallant and Tauchen (1998) provides a solution within the EMM setting through the so-called reprojection technique, but the procedure can be used more widely in parametric dynamic latent variable model estimated by other means as well.

Reprojection takes the parameter estimate of the system as given, i.e., the EMM estimator for  $\theta$  in the current context. It is then feasible to generate arbitrarily long simulated series of observable and latent variables. These simulated series can be used for estimation of the conditional density via a SNP density function approximation as under the projection step described above. In other words, the identical procedure is exploited but now for a long simulated series from the null model rather than for the observed data sample. For illustration, let  $\tilde{r} = (\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_N, \tilde{x}_0, \tilde{x}_1, \dots, \tilde{x}_{M-1})$  be a long simulated series from the null model,  $f(r_i | \mathcal{F}_{i-1}; \hat{\theta}_T)$ , where we condition on the EMM estimate. We may then utilize the SNP density estimate based on the simulated sample,  $g(\tilde{r}_i | \tilde{x}_{i-1}; \tilde{\eta})$ , in lieu of the unknown density for practical calculations, where the point estimate,  $\tilde{\eta}$ , is treated as independent of the sample size  $M$  since the estimation error is negligible for a sufficiently large simulated sample. In effect, the simulations integrate out the latent variables in the representation (4.5). Given the tractability of the SNP densities, we can now evaluate the one-step-ahead conditional mean and variance (or any other moments of interest) directly as a function of any observed history  $x_{t-1}$  by simply plugging into the SNP density estimate and perform the integration analytically – this is the reprojection step of recombining the SNP density with the actual data. Clearly, the corresponding multi-step ahead conditional density estimates can be constructed in an analogous fashion. Moreover, since the simulations also generate contemporaneous values for the latent state vectors we may similarly represent the conditional distributions of future latent state variables given the current and past observable variables through the SNP density approximation strategy,

$$f(\tilde{s}_{t+j} | \tilde{x}_t; \hat{\theta}_T) \approx g(\tilde{s}_{t+j} | \tilde{x}_t; \tilde{\eta}), \quad j \geq 0. \quad (4.25)$$

This allows for direct forecasts of conditional volatility and associated quantities in a genuine SV setting. As such, reprojection may be interpreted as a numerically intensive,

1 simulation-based, nonlinear Kalman filtering technique, providing a practical solution  
2 to the filtering and forecasting problems in Equations (4.20) and (4.21). 2

### 3 4.3. Markov Chain Monte Carlo (MCMC) procedures for inference and forecasting 3

4 The MCMC method represents a Bayesian approach to the high-dimensional inference 4  
5 problem implicit in the expression for the likelihood given in Equation (4.18). The 5  
6 approach was advocated as particularly well suited for analysis of the discrete SV model 6  
7 by Jacquier, Polson and Rossi (1994). Beyond the standard Bayesian perspective of 7  
8 treating the model parameters as random variables rather than fixed coefficients, the 8  
9 main conceptual shift is that the entire latent state vector is treated as additional param- 9  
10 eters. Hence, the main focus is on the joint distribution of the parameters and the vector 10  
11 of state variables,  $\psi = (\theta, \underline{s})$ , conditional on the data,  $f(\psi | \underline{r})$ , termed the posterior 11  
12 distribution. This density is extremely high dimensional and analytically intractable. 12  
13 The MCMC approach instead exploits that the joint distribution can be characterized 13  
14 fully through a set of associated conditional distributions where the density for a group 14  
15 of parameters, or even a single parameter, is expressed conditional on the remaining 15  
16 parameters. Concretely, let  $\psi_i$  denote the  $i$ th group of coefficients in  $\psi$ , and  $\psi_{-i}$  be 16  
17 the vector obtained from  $\psi$  by excluding the  $i$ th group of coefficients. The so-called 17  
18 Clifford–Hammersley theorem then implies that the following set of conditional distri- 18  
19 butions determines  $f(\psi | \underline{r})$ : 19  
20 21

$$22 \quad f(\psi_1 | \psi_{-1}, \underline{r}), f(\psi_2 | \psi_{-2}, \underline{r}), \dots, f(\psi_k | \psi_{-k}, \underline{r}), \quad (4.26) \quad 22$$

23 where, as described above,  $\psi = (\psi_1, \psi_2, \dots, \psi_k)$  is treated as  $k$  exclusive subsets of 23  
24 parameters. 24

25 The MCMC procedure starts by initializing  $\psi = (\theta, \underline{s})$  through conditioning on the 25  
26 observed data,  $\underline{r}$ , and drawing  $\psi$  from the assumed prior distribution. Next, by combin- 26  
27 ing the current draw for the parameter vector with the specified SV model dynamics and 27  
28 the observed returns, it is often feasible to draw the (group of) parameters sequentially 28  
29 conditional on the remainder of the system and cycle through the conditional densities 29  
30 in (4.26). A full run through the parameter vector is termed a sweep of the MCMC sam- 30  
31 pler. Some of these distributions may not be given in closed form and the draws may 31  
32 need to be extended through an accept–reject procedure termed a Metropolis–Hastings 32  
33 algorithm to ensure that the resulting Markov chain produces draws from the invariant 33  
34 joint posterior target distribution. If all the conditional distributions can be sampled 34  
35 directly we have a Gibbs sampler, but SV models often call for the two techniques to 35  
36 be used at different stages of the sweep, resulting in a hybrid MCMC algorithm. Typi- 36  
37 cally, a large number of sweeps is necessary to overcome the serial dependence inherent 37  
38 in draws of any parameter from subsequent sweeps of the sampler. Once a long sample 38  
39 from the joint posterior distribution has been generated, inference on individual param- 39  
40 eters and latent state variables can be done via the mode, mean and standard deviation of 40  
41 the posterior distribution, for example. Likewise, we can analyze properties of functions 41  
42 of the state variables directly using the posterior distribution. 42  
43 43

A key advantage of the MCMC procedure is that the distribution of the latent state vector is obtained as an inherent part of the estimation. Moreover, the inference automatically accounts for the uncertainty regarding model parameters,  $\theta$ . The resulting chain produces an elegant solution to the smoothing problem of determining  $f(\underline{s} | \underline{r})$ . Of course, from a forecasting perspective, the interest is in determining  $f(s_{t+j} | \underline{x}_t)$ , where the integer  $j \geq 0$  and  $\underline{x}_t = (r_1, r_2, \dots, r_t)$ , rather than  $f(s_{t+j} | \underline{x}_T)$  which is generated by the MCMC procedure. Unfortunately, the filter related distribution,  $f(s_{t+1} | \underline{x}_t)$ , corresponds to the intractable term in Equation (4.18) that renders the likelihood estimation impractical for genuine SV models. The MCMC inference procedure succeeds by sidestepping the need to compute this quantity. However, given the economic import of the issue, recent research is actively seeking new effective ways for better handling the filtering problem within the MCMC framework.

For a discrete-time SV model, the possibility of filtering as well as sequential one-step-ahead volatility forecasting is linked to the feasibility of providing an effective scheme to generate a random sample from  $f(s_{t+1} | \underline{x}_t, \theta)$  given an existing set of draws (or *particles*),  $s_t^1, s_t^2, \dots, s_t^N$ , from the preceding distribution  $f(s_t | \underline{x}_{t-1}, \theta)$ . Such an algorithm is termed a *particle filter*. In order to recognize the significance of the particle filter, note that by Bayes' law,

$$f(s_{t+1} | \underline{x}_{t+1}, \theta) \propto f(r_{t+1} | s_{t+1}, \theta) f(s_{t+1} | \underline{x}_t, \theta). \tag{4.27}$$

The first distribution on the right-hand side is typically specified directly by the SV model, so the issue of determining the filtering distribution on the left-hand side is essentially equivalent to the task of obtaining the predictive distribution of the state variable on the extreme right. But given a large set of particles we can approximate the latter term in straightforward fashion,

$$\begin{aligned} f(s_{t+1} | \underline{x}_t, \theta) &= \int f(s_{t+1} | s_t, \theta) f(s_t | \underline{x}_t, \theta) ds_t \\ &\approx \frac{1}{M} \sum_{j=1}^M f(s_{t+1} | s_t^j, \theta). \end{aligned} \tag{4.28}$$

This provides a direct solution to the latent state vector forecasting problem, that in turn can be plugged into (4.27) to provide a sequential update to the particle filter. This in essence is the MCMC answer to the filtering and out-of-sample forecasting problems in Equations (4.20) and (4.21). The main substantive problem is how to best sample from the last distribution in (4.28), as schemes which may appear natural can be very inefficient; see, e.g., the discussion and suggestions in [Kim, Shephard and Chib \(1998\)](#).

In summary, the MCMC approach works well for many problems of significant interest, but there are serious issues under scrutiny concerning the use of the technique for more complex settings. When applicable, it has some unique advantages such as providing a complete solution to the smoothing problem and accounting for inherent parameter estimation uncertainty. On the other hand, there are systems that are more

1 amenable to analysis under EMM and the associated diagnostic tools and general repro- 1  
2 jection procedures under EMM render it a formidable contender. It is remarkable that 2  
3 the issues of efficient forecasting and filtering within genuine SV models now has two 3  
4 attractive, albeit computationally intensive, solutions whereas just a few years ago no 4  
5 serious approach to the problem existed. 5  
6

#### 7 4.4. Further reading 7 8

9 The formal distinction between *genuine* stochastic volatility and ARCH models is de- 9  
10 veloped in Andersen (1992); see also Fleming and Kirby (2003). An early advocate for 10  
11 the Mixture-of-Distributions-Hypothesis (MDH), beyond Clark (1973), is Praetz (1972) 11  
12 who shows that an i.i.d. mixture of a Gaussian term and an inverted Gamma distribu- 12  
13 tion for the variance will produce Student  $t$  distributed returns. However, if the variance 13  
14 mixture is not linked to observed variables, the i.i.d. mixture is indistinguishable from 14  
15 a standard fat-tailed error distribution and the associated volatility process is not part of 15  
16 the genuinely stochastic volatility class. 16

17 Many alternative representations of the driving process  $s_t$  have been proposed. Clark 17  
18 (1973) observes that trading volume is highly correlated with return volatility and sug- 18  
19 gest that volume may serve as a good proxy for the “activity variable”,  $s_t$ . Moreover, he 19  
20 finds volume to be approximately lognormal (unconditionally), suggesting a lognormal- 20  
21 normal mixture for the return distribution. One drawback of this formulation is that daily 21  
22 trading volume is assumed i.i.d. Not only is this counterfactual for trading volume, but 22  
23 it also implies that the return process is i.i.d. This is at odds with the strong empirical 23  
24 evidence of pronounced temporal dependence in return volatility. A number of nat- 24  
25 ural extensions arise from the simple MDH. Tauchen and Pitts (1983) provide a more 25  
26 structural interpretation, as they develop a characterization of the joint distribution of 26  
27 the daily return and volume relationship governed by the underlying latent information 27  
28 flow  $s_t$ . However, partially for tractability, they retain the temporal independence of the 28  
29 information flow series. For early tests of the MDH model using high-frequency data 29  
30 [see, e.g., Harris (1986, 1987)], while the early return-volume literature is surveyed by 30  
31 Karpoff (1987). Gallant, Rossi and Tauchen (1992) provides an extensive study of the 31  
32 joint conditional distribution without imposing any MDH restrictions. Direct studies of 32  
33 the MDH include Lamoureux and Lastrapes (1994) and Richardson and Smith (1994). 33  
34 While the latter strongly rejects the standard MDH formulation, Andersen (1996) devel- 34  
35 ops an alternative structurally based version of the hypothesis and finds the “modified” 35  
36 MDH to perform much better. Further refinements in the specification have been pur- 36  
37 sued by, e.g., Liesenfeld (1998, 2001) and Bollerslev and Jubinsky (1999). In principle, 37  
38 the use of additional nonreturn variables along with return data should enhance estima- 38  
39 tion efficiency and allow for a better assessment of current market conditions. On the 39  
40 other hand, it is far from obvious that structural modeling of complicated multivariate 40  
41 models will prove useful in a prediction context as even minor misspecification of the 41  
42 additional series in the system may impede forecast performance. In fact, there is no 42  
43 credible evidence yet that these models help improve volatility forecast performance, 43

1 even if they have importantly enhanced our understanding of the qualitative functioning 1  
2 of financial markets. 2

3 SV diffusion models of the form analyzed by Hull and White (1987) were also pro- 3  
4 posed concurrently by Johnson and Shanno (1987), Scott (1987), and Wiggins (1987). 4  
5 An early specification and exploration of a pure jump continuous-time model is Merton 5  
6 (1976). Melino and Turnbull (1990) were among the first to estimate SV models via 6  
7 GMM. The log-SV model from (4.2)–(4.3) has emerged as a virtual testing ground for 7  
8 alternative inference procedures in this context. Andersen and Sørensen (1996) provide 8  
9 a systematic study of the choice of moments and weighting matrix for this particu- 9  
10 lar model. The lack of efficiency is highlighted in Andersen, Chung and Sørensen (?) 10  
11 where the identical model is estimated through the scores of an auxiliary model devel- 11  
12 oped in accordance with the efficient method of moments (EMM) procedure. Another 12  
13 useful approach is to apply GMM to moment conditions in the spectral domain; see, 13  
14 e.g., Singleton (2001), Jiang and Knight (2002), and Chacko and Viceira (2003). Within 14  
15 the QMLE Kalman filter based approach, a leverage effect may be incorporated and al- 15  
16 lowance for the idiosyncratic return error to be conditionally Student  $t$  distributed can be 16  
17 made, as demonstrated by Harvey, Ruiz and Shephard (1994) and Harvey and Shephard 17  
18 (1996). Andersen and Sørensen (1997) provides an extensive discussion of the relative 18  
19 efficiency of QMLE and GMM for estimation of the discrete-time log-SV model. The 19  
20 issue of asymptotically optimal moment selection for GMM estimation from among 20  
21 absolute or log squared returns in the log-SV model has received a near definitive treat- 21  
22 ment in Dhaene and Vergote (2004). The standard log-SV model has also been estimated 22  
23 through a number of other techniques by among others Danielsson and Richard (1993), 23  
24 Danielsson (1994), Fridman and Harris (1998), Monfardini (1998), and Sandmann and 24  
25 Koopman (1998). Long memory in volatility as discussed in Section 3.4 can be simi- 25  
26 larly accommodated within an SV setting; see, e.g., Breidt, Crato and de Lima (1998), 26  
27 Harvey (1998), Comte and Renault (1998), and Deo and Hurvich (2001). Duffie, Pan 27  
28 and Singleton (2000) is a good reference for a general treatment of modeling with the 28  
29 so-called affine class of models, while Piazzesi (2003) provides a recent survey of these 29  
30 models with a view toward term structure applications. 30

31 EMM may be seen as a refinement of the Method of Simulated Moments (MSM) 31  
32 of Duffie and Singleton (1993), representing a particular choice of indirect inference 32  
33 criterion, or binding function, in the terminology of Gouriéroux, Monfort and Renault 33  
34 (1993). The approach also has precedents in Smith (1990, 1993). An early application of 34  
35 EMM techniques to the discrete-time SV model is Gallant, Hsieh and Tauchen (1997). 35  
36 Among the earliest papers using EMM for stochastic volatility models are Andersen and 36  
37 Lund (1997) and Gallant and Tauchen (1997). Extensions of the EMM approach to SV 37  
38 jump-diffusions are found in Andersen, Benzoni and Lund (2002) and Chernov et al. 38  
39 (2003). As a starting point for implementations of the EMM procedure, one may access 39  
40 general purpose EMM and SNP code from a web site maintained by A. Ronald Gallant 40  
41 and George E. Tauchen at Duke University at the link <ftp.econ.duke.edu> in the direc- 41  
42 tories /pub/get/emm and /pub/arg/snp, respectively. In practical applications, it is of- 42  
43 ten advantageous to further refine the SNP density approximations through specifically 43

1 designed leading GARCH terms which parsimoniously capture the dependency structure in the specific data under investigation. The benefits of doing so is further discussed in Andersen and Lund (1997) and Andersen, Benzoni and Lund (2002).

2  
3  
4 The particle filter discussed above for the generation of filter estimates for the latent variables of interest within the standard SV model arguably provides a more versatile approach than the alternative importance sampling methods described by, e.g., Danielsson (1994) and Sandmann and Koopman (1998). The extension of the MCMC inference technique to a continuous-time setting is discussed in Elerian, Chib and Shephard (2001) and Eraker (2001). The latter also provides one of the first examples of MCMC estimation of an SV diffusion model, while Eraker, Johannes and Polson (2003) further introduces jumps in both prices and volatility. Johannes and Polson (2003) offer a recent comprehensive survey of the still ongoing research on the use of the MCMC approach in the general nonlinear jump-diffusion SV setting.

## 16 5. Realized volatility

17  
18 The notion of realized volatility has at least two key distinct implications for practical volatility estimation and forecasting. The first relates to the measurement of realizations of the latent volatility process without the need to rely on an explicit model. As such, the realized volatility provides *the* natural benchmark for forecast evaluation purposes. The second relates to the possibility of modeling volatility directly through standard time series techniques with discretely sampled daily observations, while effectively exploiting the information in intraday high-frequency data.

### 27 5.1. The notion of realized volatility

28  
29 The most fundamental feature of realized volatility is that it provides a consistent non-parametric estimate of the price variability that has transpired over a given discrete interval. Any log-price process subject to a no-arbitrage condition and weak auxiliary assumptions will constitute a semi-martingale that may be decomposed into a locally predictable mean component and a martingale with finite second moments. Within this class, there is a unique measure for the realized sample-path variation termed the quadratic variation. By construction the quadratic variation cumulates the intensity of the unexpected price changes over the specific horizon and it is thus a prime candidate for a formal volatility measure.

30  
31  
32  
33  
34  
35  
36  
37  
38 The intuition behind the use of realized volatility as a return variation measure is most readily conveyed within the popular continuous-time diffusion setting (4.9) obtained by ruling out jumps and thus reducing to the representation (1.7), reproduced here for convenience,

$$39 \quad dp(t) = \mu(t) dt + \sigma(t) dW(t), \quad t \in [0, T]. \quad (5.1)$$

Applying a discretization of the process as in Section 1, we have for small  $\Delta > 0$ , that

$$r(t, \Delta) \equiv p(t) - p(t - \Delta) \simeq \mu(t - \Delta)\Delta + \sigma(t - \Delta) \Delta W(t), \tag{5.2}$$

where  $\Delta W(t) \equiv W(t) - W(t - \Delta) \sim N(0, \Delta)$ .

Over short intervals the squared return and the squared return innovation are closely related as both are largely determined by the idiosyncratic return component,

$$r^2(t, \Delta) \simeq \mu^2(t - \Delta)\Delta^2 + 2\Delta\mu(t - \Delta)\sigma(t - \Delta) \Delta W(t) + \sigma^2(t - \Delta)(\Delta W(t))^2. \tag{5.3}$$

In particular, the return variance is (approximately) equal to the expected squared return innovation,

$$\text{Var}[r(t, \Delta) \mid \mathcal{F}_{t-\Delta}] \simeq E[r^2(t, \Delta) \mid \mathcal{F}_{t-\Delta}] \simeq \sigma^2(t - \Delta)\Delta. \tag{5.4}$$

This suggests that we may be able to measure the return volatility directly from the squared return observations. However, this feature is not of much direct use as the high-frequency returns have a large idiosyncratic component that induces a sizeable measurement error into the actual squared return relative to the underlying variance. Up to the dominant order in  $\Delta$ ,

$$\text{Var}[r^2(t, \Delta) \mid \mathcal{F}_{t-\Delta}] \simeq 2\sigma^4(t - \Delta)\Delta^2, \tag{5.5}$$

where terms involving higher powers of  $\Delta$  are ignored as they become negligible for small values of  $\Delta$ . Thus, it follows that the “noise-to-signal” ratio in squared returns relative to the underlying volatility is of the same order as volatility itself,

$$\frac{\text{Var}[r^2(t, \Delta) \mid \mathcal{F}_{t-\Delta}]}{E[r^2(t, \Delta) \mid \mathcal{F}_{t-\Delta}]} \simeq 2E[r^2(t, \Delta) \mid \mathcal{F}_{t-\Delta}]. \tag{5.6}$$

This relationship cannot be circumvented when only one (squared) return observation is used as a volatility proxy. Instead, by exploiting the fact that return innovations, under a no-arbitrage (semi-martingale) assumption, are serially uncorrelated to construct volatility measures for lower frequency returns we find, to dominant order in  $\Delta$ ,

$$\sum_{j=1}^{1/\Delta} E[r^2(t - 1 + j \cdot \Delta, \Delta) \mid \mathcal{F}_{t-1+j \cdot \Delta}] \simeq \sum_{j=1}^{1/\Delta} \sigma^2(t - 1 + j \cdot \Delta) \cdot \Delta \simeq \int_{t-1}^t \sigma^2(s) ds, \tag{5.7}$$

where the last approximation stems from the sum converging to the corresponding integral as the size of  $\Delta$  shrinks toward zero. Equation (5.7) generalizes (5.4) to the multi-period setting with the second approximation in (5.7) only being meaningful for  $\Delta$  small.

The advantage of (5.7) is that the uncorrelated “measurement errors” have been effectively smoothed away to generate a much better noise-to-signal ratio. The expression

in (5.5) may be extended in a similar manner to yield

$$\begin{aligned} \sum_{j=1}^{1/\Delta} \text{Var}[r^2(t-1+j\cdot\Delta, \Delta) \mid \mathcal{F}_{t-1+j\cdot\Delta}] &\simeq 2 \sum_{j=1}^{1/\Delta} \sigma^4(t-1+j\cdot\Delta) \cdot \Delta^2 \\ &\simeq 2\Delta \int_{t-1}^t \sigma^4(s) ds. \end{aligned} \quad (5.8)$$

Consequently,

$$\begin{aligned} \frac{\sum_{j=1}^{1/\Delta} \text{Var}[r^2(t-1+j\cdot\Delta, \Delta) \mid \mathcal{F}_{t-1+j\cdot\Delta}]}{\sum_{j=1}^{1/\Delta} E[r^2(t-1+j\cdot\Delta, \Delta) \mid \mathcal{F}_{t-1+j\cdot\Delta}]} &\simeq 2\Delta \frac{\int_{t-1}^t \sigma^4(s) ds}{\int_{t-1}^t \sigma^2(s) ds} \\ &\equiv 2\Delta \frac{\text{IQ}(t)}{\text{IV}(t)}, \end{aligned} \quad (5.9)$$

where the integrated quarticity is defined through the identity on the right-hand side of (5.9), with the integrated variance,  $\text{IV}(t)$ , having previously been defined in (4.12).

The fact that the “noise-to-signal” ratio in (5.9) shrinks to zero with  $\Delta$  suggests that high-frequency returns may be very useful for estimation of the underlying (integrated) volatility process. The notion of realized volatility is designed to take advantage of these features. Formally, realized volatility is defined as

$$\text{RV}(t, \Delta) = \sum_{j=1}^{1/\Delta} r^2(t-1+j\cdot\Delta, \Delta). \quad (5.10)$$

Equation (5.8) suggests that realized volatility is consistent for the integrated volatility in the sense that finer and finer sampling of the intraday returns,  $\Delta \rightarrow 0$ , ultimately will annihilate the measurement error and, in the limit, realized volatility measures the latent integrated volatility perfectly, that is,

$$\text{RV}(t, \Delta) \rightarrow \text{IV}(t), \quad (5.11)$$

as  $\Delta \rightarrow 0$ . These arguments may indeed be formalized; see, e.g., the extended discussion in Andersen, Bollerslev and Diebold (2003a). In reality, there is a definite lower bound on the return horizon that can be used productively for computation of the realized volatility, both because we only observe discretely sampled returns and, more important, market microstructure features such as discreteness of the price grid and bid–ask spreads induce gross violations of the semi-martingale property at the very highest return frequencies. This implies that we typically will be sampling returns at an intraday frequency that leaves a nonnegligible error term in the estimate of integrated volatility. It is natural to conjecture from (5.9) that asymptotically, as  $\Delta \rightarrow 0$ ,

$$\sqrt{1/\Delta}[\text{RV}(t, \Delta) - \text{IV}(t)] \sim N(0, 2 \cdot \text{IQ}(t)), \quad (5.12)$$

which turns out to be true under quite general assumptions. Of course, the  $\text{IQ}(t)$  measure must be estimated as well for the above result to provide a practical tool for inference.



The distributional result in (5.12) and a feasible consistent estimator for  $IQ(t)$  based purely on intraday data have been provided by Barndorff-Nielsen and Shephard (2002, 2004b). It may further be shown that these measurement errors are approximately uncorrelated across consecutive periods which has important simplifying implications for time series modeling.

The consistency result in (5.11) extends to the general semi-martingale setting where the price path may display discontinuities due to jumps, as specified in Equation (4.9). The realized volatility will still converge in the continuous-record limit ( $\Delta \rightarrow 0$ ) to the period-by-period quadratic variation of the semi-martingale. However, the quadratic variation is no longer identical to the integrated volatility but will also include the cumulative squared jumps,

$$RV(t, \Delta) \rightarrow QV(t) = \int_{t-1}^t \sigma^2(s) ds + \sum_{t-1 < s \leq t} \kappa^2(s). \quad (5.13)$$

A few comments are in order. First,  $QV(t)$  is best interpreted as the actual return variation that transpired over the period, and as such it is the natural target for realized volatility measurement. Second,  $QV(t)$  is the realization of a random variable which generally cannot be forecasted with certainty at time  $t - 1$ . But it does represent the future realization that volatility forecasts for time  $t$  should be compared against. In other words, the quadratic variation constitutes *the* quantity of interest in volatility measurement and forecasting. Since the realizations of  $QV(t)$  are latent, it is natural to use the observed  $RV(t, \Delta)$  as a feasible proxy. Third, financial decision making is concerned with forecasts of volatility (or quadratic variation) rather than the  $QV(t)$  directly. Fourth, the identification of forecasts of return volatility with forecasts of quadratic variation is only approximate as it ignores variation in the process induced by innovations in the conditional mean process. Over short horizons the distinction is negligible, but for longer run volatility prediction (quarterly or annual) one may need to pay some attention to the discrepancy between the two quantities, as discussed at length in Andersen, Bollerslev and Diebold (2003a).

The distribution theory for quadratic variation under the continuous sample path assumption has also been extended to cover cumulative absolute returns raised to an arbitrary power. The leading case involves cumulating the high-frequency absolute returns. These quantities display improved robustness properties relative to realized volatility as the impact of outliers are mitigated. Although the limiting quantity – the power variation – is not directly linked to the usual volatility measure of interest in finance, this concept has inspired further theoretical developments that has led to intriguing new non-parametric tests for the presence of jumps and the identification of the associated jump sizes; see, e.g., Barndorff-Nielsen and Shephard (2004a). Since the jumps may have very different intertemporal persistence characteristics than the diffusion volatility, explicit disentangling of the components of quadratic variation corresponding to jumps versus diffusion volatility can have important implications for volatility forecasting.

In summary, the notion of realized volatility represents a model-free approach to (continuous-record) consistent estimation of the quadratic return variation under general

assumptions based primarily upon arbitrage-free financial markets. As such it allows us to harness the information inherent in high-frequency returns for assessment of lower frequency return volatility. It is thus *the* natural approach to measuring actual (ex-post) realized return variation over a given horizon. This perspective has now gained widespread acceptance in the literature, where alternative volatility forecast models are routinely assessed in terms of their ability to explain the distribution of subsequent realized volatility, as defined above.

## 5.2. Realized volatility modeling

The realized volatility is by construction an observed proxy for the underlying quadratic variation and the associated (measurement) errors are uncorrelated. This suggests a straightforward approach where the temporal features of the series are modeled through standard time series techniques, letting the data guide the choice of the appropriate distributional assumptions and the dynamic representation. This is akin to the standard procedure for modeling macroeconomic data where the underlying quantities are measured (most likely with a substantial degree of error) and then treated as directly observed variables.

The strategy of estimating time series models directly for realized volatility is advocated in a sequence of papers by Andersen et al. (2001a, 2001b, 2003). A striking finding is that the realized volatility series share fundamental statistical properties across different asset classes, time periods, and countries. The evidence points strongly toward a long-memory type of dependency in volatility. Moreover, the logarithmic realized volatility series is typically much closer to being homoskedastic and approximately unconditionally Gaussian. These features are readily captured through an ARFIMA( $p, d, 0$ ) representation of the logarithmic realized volatility,

$$\Phi(L)(1-L)^d(\log RV(t, \Delta) - \mu_0) = u_t, \quad t = 1, 2, \dots, T, \quad (5.14)$$

where  $(1-L)^d$  denotes the fractional differencing operator,  $\Phi(L)$  is a polynomial lag operator accounting for standard autoregressive structure,  $\mu_0$  represents the unconditional mean of the logarithmic realized volatility, and  $u_t$  is a white noise error term that is (approximately) Gaussian. The coefficient  $d$  usually takes a value around 0.40, consistent with a stationary but highly persistent volatility process for which shocks only decay at a slow hyperbolic rate rather than the geometric rate associated with standard ARMA models or GARCH models for the conditional variance. Finally, the volatility of volatility is strongly increasing in the level of volatility as log-realized volatility is approximately homoskedastic. This is, of course, reminiscent of the log-SV and the EGARCH models.

A number of practical modeling issues have been sidestepped above. One is the choice of the sampling frequency at which the realized volatility measures are constructed. The early literature focused primarily on determining the highest intraday frequency at which the underlying returns satisfy the maintained semi-martingale assumption of being approximately uncorrelated. An early diagnostic along these lines

1 termed the “volatility signature plot” was developed by Andersen et al. (1999, 2000), as 1  
2 discussed further in Section 7 below. A simple alternative is to apply standard ARMA 2  
3 filtering to the high-frequency returns in order to strip them of any “artificial” serial cor- 3  
4 relation induced by the market microstructure noise, and then proceed with the filtered 4  
5 uncorrelated returns in lieu of the raw high-frequency returns. While none of these pro- 5  
6 cedures are optimal in a formal statistical sense, they both appear to work reasonable 6  
7 well in many practical situations. Meanwhile, a number of alternative more efficient 7  
8 sampling schemes under various assumptions about the market microstructure compli- 8  
9 cations have recently been proposed in a series of interesting papers, and this is still 9  
10 very much ongoing research. 10

11 A second issue concerns the potential separation of jumps and diffusive volatility 11  
12 components in the realized volatility process. The theoretical basis for these proce- 12  
13 dures and some initial empirical work is presented in Barndorff-Nielsen and Shephard 13  
14 (2004a). The issue has been pursued empirically by Andersen, Bollerslev and Diebold 14  
15 (2003b), who find compelling evidence that the diffusive volatility is much more per- 15  
16 sistent than the jump component. In fact, the jumps appear close to i.i.d., although the 16  
17 jumps in equity indices display some clustering, especially in the size of the jumps. This 17  
18 points to potentially important improvements in modeling and forecasting from this type 18  
19 of separation of the realized volatility into sudden discrete shifts in prices versus more 19  
20 permanent fluctuations in the intensity of the regular price movements. Empirically, 20  
21 this is in line with the evidence favoring non-Gaussian fat-tailed return innovations in 21  
22 ARCH models. 22

23 A third issue is the approach used to best accommodate the indications of “long 23  
24 memory”. An alternative to fractional integration is to introduce several autoregressive 24  
25 volatility components into the model. As discussed in the context of the GARCH class 25  
26 of models in Section 3.4, if the different components display strong, but varying, degrees 26  
27 of persistence they may combine to produce a volatility dependency structure that is 27  
28 indistinguishable from long memory over even relatively long horizons. 28

### 29 5.3. Realized volatility forecasting 29

30 Forecasting is straightforward once the realized volatility has been cast within the 30  
31 traditional time series framework and the model parameters have been estimated. Since 31  
32 the driving variable is the realized volatility we no longer face a latent variable issue. 32  
33 This implies that standard methods for forecasting a time series within the ARFIMA 33  
34 framework is available; see, e.g., Beran (1994) for an introduction to models incorporat- 34  
35 ing long-memory features. One-step-ahead minimum mean-squared error forecasts are 35  
36 readily produced, and within the linear Gaussian setting it is then legitimate to further 36  
37 condition on the forecast in order to iterate forward and produce multiple-step-ahead 37  
38 forecasts. There are a couple of caveats, however. First, as with most other volatility 38  
39 forecasting procedures, the forecasts are, of course, conditional on the point estimate 39  
40 for the model parameters. Second, if the model is formulated in terms of the logarithmic 40  
41 volatility then it is also log volatility that is being predicted through the usual forecast 41  
42 42  
43 43

1 procedures. There is a practical problem of converting the forecast for log volatility into 1  
2 a “pure” volatility forecast as the expected value of the transformed variable depends 2  
3 not only on the expected log volatility, but on the entire multiple-step-ahead conditional 3  
4 distribution of log volatility. For short horizons this is not an issue as the requisite cor- 4  
5 rection term usually is negligible, but for longer horizons adjustments may be necessary. 5  
6 This is similar to the issue that arise in the construction of forecast form the EGARCH 6  
7 model. As discussed in Section 3.6, the required correction term may be constructed by 7  
8 simulation based methods, but the preferred approach will depend on the application at 8  
9 hand and the distributional characteristics of the model. For additional inspiration on 9  
10 how to address such issues consult, e.g., Chapter 6 on ARMA forecasting methods by 10  
11 Lütkepohl (2006) in this handbook. 11

12 A few additional comments are in order. First, the evidence in Andersen, Bollerslev 12  
13 and Diebold (2003a) indicates that the above approach has very good potential. The as- 13  
14 sociated forecasts for foreign exchange rate volatility outperform a string of alternative 14  
15 candidate models from the literature. This is not a tautology as it should be preferable 15  
16 to generate the forecasts from the true underlying model rather than an ad hoc time se- 16  
17 ries model estimated from period-by-period observations of realized volatility. In other 17  
18 words, if a GARCH diffusion is the true model then optimal forecasts would incorporate 18  
19 the restrictions implied by this model. However, the high-frequency volatility process is 19  
20 truly complex, possessing several periodic components, erratic short run dynamics and 20  
21 longer run persistence features that combined appear beyond reach of simple parametric 21  
22 models. The empirical evidence suggests that daily realized volatility serves as a simple, 22  
23 yet effective, aggregator of the volatility information inherent in the intraday data. 23

24 Second, there is an issue of how to compute realized volatility for a calendar period 24  
25 when the trading day is limited by an official closing. This problem is minor for the 25  
26 over-the-counter foreign exchange market where 24-hour trading is observed, but this 26  
27 is often not the case for equity or bond markets. For example, for a one-month-ahead 27  
28 equity volatility forecast there may only be twenty-two trading days with about six-and- 28  
29 a-half hours of trading per day. But the underlying price process is not stalled while 29  
30 the markets are closed. Oftentimes there will be substantial changes in prices between one 30  
31 market close and the subsequent opening, reflecting return volatility overnight and over 31  
32 the weekend. One solution is to simply rely on the intraday returns for a realized volatil- 32  
33 ity measure over the trading day and then scale this quantity up by a factor that reflects 33  
34 the average ratio of volatility over the calendar day versus the trading day. This may 34  
35 work quite satisfactorily in practice, but it obviously ignores the close-to-open return 35  
36 for a given day entirely in constructing the realized volatility for that calendar day. Al- 36  
37 ternatively, the volatility of the close-to-open return may be modeled by a conventional 37  
38 GARCH type model. 38

39 Third, we have not discussed the preferred sampling frequency of intraday returns in 39  
40 situations where the underlying asset is relatively illiquid. If updated price observations 40  
41 are only available intermittently throughout the trading day, many high-frequency re- 41  
42 turns may have to be computed from prices or quotes earlier in the day. This brings up a 42  
43 couple of issues. One, the effective sampling frequency is lower than the one that we are 43

1 trying to use for the realized volatility computation. Two, illiquid price series also tend 1  
2 to have larger bid–ask spreads and be more sensitive to random fluctuations in order 2  
3 flow, implying that the associated return series will contain a relatively large amount of 3  
4 noise. A simple response that will help alleviate both issues is to lower the sampling 4  
5 frequency. However, with the use of less intraday returns comes a larger measurement 5  
6 error in realized volatility, as evidenced by Equation (5.12). Nonetheless, for an illiquid 6  
7 asset it may only be possible to construct meaningful weekly rather than daily realized 7  
8 volatility measures from say half-hourly or hourly return observations rather than five- 8  
9 minute returns. Consequently, the intertemporal fluctuations are smoothed out so that 9  
10 the observed measure carries less information about the true state of the volatility at the 10  
11 end of the period. This, of course, can be critically important for accurate forecasting. 11

12 In sum, the use of the realized volatility measures for forecasting is still in its infancy 12  
13 and many issues must be explored in future work. However, it is clear that the use of 13  
14 intraday information has large potential to improve upon the performance of standard 14  
15 volatility forecast procedures based only on daily or lower frequency data. The real- 15  
16 ized volatility approach circumvents the need to model the intraday data directly and 16  
17 thus provides a great deal of simplification. Importantly, it seems to achieve this ob- 17  
18 jective without sacrificing a lot of efficiency. For example, Andersen, Bollerslev and 18  
19 Meddahi (2004) find the time series approach built directly from the realized volatil- 19  
20 ity measures to be very good approximations to the theoretically optimal procedures in 20  
21 a broad class of SV diffusion models that can be analyzed analytically through newly 21  
22 developed tools associated with the so-called Eigenfunction SV models of Meddahi 22  
23 (2001). Nonetheless, if the objective exclusively is volatility forecasting, some very re- 23  
24 cent work suggests that alternative intraday measures may carry even more empirically 24  
25 relevant information regarding future volatility, including the power variation measures 25  
26 constructed from cumulative absolute returns; see, e.g., Ghysels, Santa-Clara and Valka- 26  
27 nov (2004). This likely reflects superior robustness features of absolute versus squared 27  
28 intraday returns, but verification of such conjectures awaits future research. The conflu- 28  
29 ence of compelling empirical performance, novel econometric theory, the availability of 29  
30 ever more high-frequency data and computational power, and the importance of forecast 30  
31 performance for decision making render this approach fertile ground for new research. 31

#### 32 33 5.4. Further reading 33 34

35 The realized volatility approach has a precedent in the use of cumulative daily squared 35  
36 returns as monthly volatility measures; see, e.g., French, Schwert and Stambaugh (1987) 36  
37 and Schwert (1989). Hsieh (1989) was among the first to informally apply this same 37  
38 procedure with high-frequency intraday returns, while Zhou (1996) provides one of the 38  
39 earliest formal assessments of the relationship between cumulative squared intraday 39  
40 returns and the underlying return variance, albeit in a highly stylized setting. The pio- 40  
41 neering work by Olsen & Associates on the use of high-frequency data, as summarized 41  
42 in Dacorogna et al. (2001), also importantly paved the way for many of the more recent 42  
43 empirical developments in the realized volatility area. 43

1 The use of component structures and related autoregressive specifications for ap- 1  
2 proximating long-memory dependencies within the realized volatility setting has been 2  
3 explored by Andersen, Bollerslev and Diebold (2003b), Barndorff-Nielsen and Shep- 3  
4 hard (2001), Bollerslev and Wright (2001), and Corsi (2003), among others. The finite 4  
5 sample performance of alternative nonparametric tests for jumps based on the bipower 5  
6 variation measure introduced by Barndorff-Nielsen and Shephard (2004a) have been 6  
7 extensively analyzed by Huang and Tauchen (2004). Andersen, Bollerslev and Diebold 7  
8 (2003b) demonstrate the importance of disentangling the components of quadratic 8  
9 variation corresponding to jumps versus diffusion volatility for volatility forecasting. 9  
10 The complexities involved in a direct high-frequency characterization of the volatility 10  
11 process is also illustrated by Andersen and Bollerslev (1998c). 11

12  
13 Ways of incorporating noisy overnight returns into the daily realized volatility mea- 13  
14 sure are discussed in Fleming, Kirby and Ostdiek (2003) and Hansen and Lunde 14  
15 (2004a). The related issue of measuring the integrated variance in the presence of mar- 15  
16 ket microstructure noise and how to best use all of the available high frequency data 16  
17 has been addressed in a rapidly growing recent literature. Corsi et al. (2001) argue for 17  
18 the use of exponential moving average filtering, similar to a standard MA(1) filter for 18  
19 the high-frequency returns, while other more recent procedures, including sub-sampling 19  
20 and ways of choosing the “optimal” sampling frequency, have been suggested and ana- 20  
21 lyzed empirically by, e.g., Aït-Sahalia, Mykland and Zhang (2005), Bandi and Russell 21  
22 (2004), Barucci and Reno (2002), Bollen and Inder (2002), Curci and Corsi (2004), and 22  
23 Hansen and Lunde (2004b), among others. Some of these issues are discussed further 23  
24 in Section 7 below, where we also consider the robust alternative range based volatil- 24  
25 ity estimator recently explored by Alizadeh, Brandt and Diebold (2002) for dynamic 25  
26 volatility modeling and forecasting. 26  
27

28 Implied volatility provides yet another forward looking volatility measure. Implied 28  
29 volatilities are based on the market’s forecasts of future volatilities extracted from the 29  
30 prices of options written on the asset of interest. As discussed in Section 2.2.4 above, 30  
31 using a specific option pricing formula, one may infer the expected integrated volatil- 31  
32 ity of the underlying asset over the remaining time-to-maturity of the option. The main 32  
33 complication associated with the use of these procedures lies in the fact that the op- 33  
34 tion prices also generally reflect a volatility risk premium in the realistic scenario where 34  
35 the volatility risk cannot be perfectly hedged; see, e.g., the discussion in Bollerslev and 35  
36 Zhou (2005). Nonetheless, many studies find options implied volatilities to provide use- 36  
37 ful information regarding the future volatility of the underlying asset. At the same time, 37  
38 the results pertaining to the forecast performance of implied volatilities are somewhat 38  
39 mixed, and there is still only limited evidence regarding the relative predictive power of 39  
40 implied volatilities versus the realized volatility procedures discussed above. Another 40  
41 issue is that many assets of interest do not have sufficiently active options markets that 41  
42 reliable implied volatilities can be computed on, say, a daily basis. 42  
43

## 6. Multivariate volatility

The discussion in the preceding three sections has been focused almost exclusively on univariate forecasts. Yet, as discussed in Section 2, in many practical situations covariance and/or correlation forecasting plays an equal, if not even more important, role in the uses of volatility forecasts. Fortunately, many of the same ideas and procedures discussed in the context of univariate forecasts are easily adapted to the multivariate setting. However, two important complications arise in this setting, namely the imposition of sufficient conditions to ensure that the forecasts for the covariance matrix remain positive definite for all forecasting horizons, and, second, maintaining an empirically realistic yet parsimoniously parameterized model. We will organize our discussion of the various multivariate approaches with these key concerns in mind.

Before turning to this discussion, it is worth noting that in many situations, multivariate volatility modeling and forecasting may be conveniently sidestepped through the use of much-simpler-to-implement univariate procedures for appropriately transformed series. In particular, in the context of financial market volatility forecasting, consider the leading case involving the variance of a portfolio made up of  $N$  individual assets. In the notation of Section 2.2.1 above,

$$r_{w,t+1} = \sum_{i=1}^N w_{i,t} r_{i,t+1} \equiv w'_t R_{t+1}. \quad (6.1)$$

The conditional one-step-ahead variance of the portfolio equals

$$\sigma_{w,t+1|t}^2 = \sum_{i=1}^N \sum_{j=1}^N w_{i,t} w_{j,t} \{\Omega_{t+1|t}\}_{i,j} = w'_t \Omega_{t+1|t} w_t, \quad (6.2)$$

where  $\Omega_{t+1|t}$  denotes the  $N \times N$  covariance matrix for the returns. A forecast for the portfolio return variance based upon this representation therefore requires the construction of multivariate forecasts for the  $\frac{1}{2}N(N+1)$  unique elements in the covariance matrix for the assets in the portfolio. Alternatively, define the univariate time series of artificial historical portfolio returns constructed on the basis of the weights for the current portfolio in place,

$$r_{w,\tau}^t \equiv w'_t R_\tau, \quad \tau = 1, 2, \dots, t. \quad (6.3)$$

A univariate forecast for the variance of the returns on this artificially constructed portfolio indirectly ensures that the covariances among the individual assets receive exactly the same weight as in Equation (6.2). Note, that unless the portfolio weights for the actual portfolio in place are constantly rebalanced, the returns on this artificially constructed portfolio will generally differ from the actual portfolio returns, that is  $r_{w,\tau}^t \equiv w'_t R_\tau \neq w'_\tau R_\tau \equiv r_{w,\tau}$  for  $\tau \neq t$ . As such, the construction of the variance forecasts for  $r_{w,\tau}^t$  requires the estimation of a new (univariate) model each period to properly reflect the relevant portfolio composition in place at time  $t$ . Nonetheless,

1 univariate volatility models are generally much easier to implement than their multi- 1  
 2 variate counterparts, so that this approach will typically be much less computationally 2  
 3 demanding than the formulation of a satisfactory full scale multivariate volatility model 3  
 4 for  $\Omega_{t+1|t}$ , especially for large values of  $N$ . Moreover, since the relative changes in 4  
 5 the actual portfolio weights from one period to the next are likely to be small, good 5  
 6 starting values for the parameters in the period-by-period univariate models are readily 6  
 7 available from the estimates obtained in the previous period. Of course, this simplified 7  
 8 approach also requires that historical returns for the different assets in the portfolio are 8  
 9 actually available. If that is not the case, artificial historical prices could be constructed 9  
 10 from a pricing model, or by matching the returns to those of other assets with similar 10  
 11 characteristics; see, e.g., Andersen et al. (2005) for further discussion along these lines. 11

12 Meanwhile, as discussed in Sections 2.2.2 and 2.2.3, there are, of course, many situa- 12  
 13 tions in which forecasts for the covariances and/or correlations play a direct and 13  
 14 important role in properly assessing and comparing the risks of different decisions or in- 14  
 15 vestment opportunities. We next turn to a discussion of some of the multivariate models 15  
 16 and forecasting procedures available for doing so. 16

### 17 6.1. Exponential smoothing and RiskMetrics 17

18 The exponentially weighted moving average filter, championed by RiskMetrics, is ar- 18  
 19 guable the most commonly applied approach among finance practitioners for estimating 19  
 20 time-varying covariance matrices. Specifically, let  $Y_t \equiv R_t$  denote the  $N \times 1$  vector of 20  
 21 asset returns. The estimate for the current covariance matrix is then defined by 21  
 22 22  
 23 23  
 24 24

$$25 \hat{\Omega}_t = \gamma Y_t Y_t' + (1 - \gamma) \hat{\Omega}_{t-1} \equiv \gamma \sum_{i=1}^{\infty} (1 - \gamma)^{i-1} Y_t Y_t'. \quad (6.4) \quad 25$$

26 This directly parallels the earlier univariate definition in Equation (3.2), with the ad- 26  
 27 ditional assumption that the mean of all the elements in  $Y_t$  is equal to zero. As in the 27  
 28 univariate case, practical implementation is typically done by truncating the sum at 28  
 29  $I = t - 1$ , scaling the finite sum by  $1/[1 - (1 - \gamma)^I]$ . This approach is obviously very 29  
 30 simple to implement in any dimension  $N$ , involving only a single tuning parameter,  $\gamma$ , 30  
 31 or by appealing to the values advocated by RiskMetrics (0.06 and 0.04 in the case of 31  
 32 daily and monthly returns, respectively) no unknown parameters whatsoever. Moreover, 32  
 33 the resulting covariance matrix estimates are guaranteed to be positive definite. 33  
 34 34  
 35 35

36 The simple one-parameter filter in (6.4) may, of course, be further refined by allowing 36  
 37 for different decay rates for the different elements in  $\hat{\Omega}_t$ . Specifically, by using a smaller 37  
 38 value of  $\gamma$  for the off-diagonal, or covariance, terms in  $\hat{\Omega}_t$ , the corresponding time- 38  
 39 varying correlations, 39  
 40 40

$$41 \hat{\rho}_{ij,t} \equiv \frac{\{\hat{\Omega}_t\}_{ij}}{\{\hat{\Omega}_t\}_{ii}^{1/2} \{\hat{\Omega}_t\}_{jj}^{1/2}}, \quad (6.5) \quad 41$$

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43



will exhibit more persistent dynamic dependencies. This slower rate of decay for the correlations often provide a better characterization of the dependencies across assets.

Meanwhile, the  $h$ -period-ahead forecasts obtained by simply equating the future conditional covariance matrix with the current filtered estimate,

$$\text{Var}(Y_{t+h} | \mathcal{F}_t) \equiv \Omega_{t+h|t} \approx \hat{\Omega}_t, \quad (6.6)$$

are plagued by the same counterfactual implications highlighted in the context of the corresponding univariate filter in Sections 3.1 and 3.2. In particular, assuming that the one-period returns are serially uncorrelated so that the forecast for the covariance matrix of the multi-period returns equals the sum of the successive one-period covariance forecasts,

$$\text{Var}(Y_{t+k} + Y_{t+k-1} + \cdots + Y_{t+1} | \mathcal{F}_t) \equiv \Omega_{t:t+k|t} \approx k\hat{\Omega}_t, \quad (6.7)$$

the multi-period covariance matrix scales with the forecast horizon,  $k$ , rather than incorporating empirically more realistic mean-reversion. Moreover, it is difficult to contemplate the choice of the tuning parameter(s),  $\gamma$ , for the various elements in  $\hat{\Omega}_t$  without a formal model. The multivariate GARCH class of models provides an answer to these problems by formally characterizing the temporal dependencies in the forecasts for the individual variances and covariances within a coherent statistical framework.

## 6.2. Multivariate GARCH models

The multivariate GARCH class of models was first introduced and estimated empirically by Bollerslev, Engle and Wooldridge (1988). Denoting the one-step-ahead conditional mean vector and covariance matrix for  $Y_t$  by  $M_{t|t-1} \equiv E(Y_t | \mathcal{F}_{t-1})$  and  $\Omega_{t|t-1} \equiv \text{Var}(Y_t | \mathcal{F}_{t-1})$ , respectively, the multivariate version of the decomposition in (3.5) may be expressed as

$$Y_t = M_{t|t-1} + \Omega_{t|t-1}^{1/2} Z_t, \quad Z_t \sim \text{i.i.d.}, \quad E(Z_t) = 0, \quad \text{Var}(Z_t) = I, \quad (6.8)$$

where  $Z_t$  now denotes a vector white noise process with unit variances. The square root of the  $\Omega_{t|t-1}$  matrix is not unique, but any operator satisfying the condition that  $\Omega_{t|t-1}^{1/2} \cdot \Omega_{t|t-1}^{1/2} \equiv \Omega_{t|t-1}$  will give rise to the same conditional covariance matrix.

The multivariate counterpart to the successful univariate GARCH(1, 1) model in (3.6) is now naturally defined by

$$\text{vech}(\Omega_{t|t-1}) = C + A \text{vech}(e_{t-1} e'_{t-1}) + B \text{vech}(\Omega_{t-1|t-2}), \quad (6.9)$$

where  $e_t \equiv \Omega_{t|t-1}^{1/2} Z_t$ ,  $\text{vech}(\cdot)$  denotes the operator that stacks the  $\frac{1}{2}N(N+1)$  unique elements in the lower triangular part of a symmetric matrix into a  $\frac{1}{2}N(N+1) \times 1$  vector, and the parameter matrices  $C$ ,  $A$ , and  $B$ , are of dimensions  $\frac{1}{2}N(N+1) \times 1$ ,  $\frac{1}{2}N(N+1) \times \frac{1}{2}N(N+1)$ , and  $\frac{1}{2}N(N+1) \times \frac{1}{2}N(N+1)$ , respectively. As in the univariate case, the GARCH(1, 1) model in (6.9) is readily extended to higher order



In the diagonal vech model the  $A$  and  $B$  matrices are both assumed to be diagonal, so that a particular element in the conditional covariance matrix only depends on its own lagged value and the corresponding cross product of the innovations. This model may alternatively be written in terms of Hadamard products, or element-by-element multiplication, as

$$\Omega_{t|t-1} = C + A \circ (e_{t-1}e'_{t-1}) + B \circ \Omega_{t-1|t-2}, \quad (6.11)$$

where  $C$ ,  $A$ , and  $B$  now denote symmetric positive definite matrices of dimension  $N \times N$ . This model greatly reduces the number of free parameters to  $3(N^2 + N)/2$ , and, importantly, covariance matrix forecasts generated from this model according to the recursions in (6.10) are guaranteed to be positive definite. However, the model remains prohibitively “expensive” in terms of parameters in large dimensions. For instance, for  $N = 100$  there are still 15,150 free parameters in the unrestricted diagonal vech model.

A further dramatic simplification is obtained by restricting all of the elements in the  $A$  and  $B$  matrices in (6.11) to be the same,

$$\Omega_{t|t-1} = C + \alpha(e_{t-1}e'_{t-1}) + \beta\Omega_{t-1|t-2}. \quad (6.12)$$

This scalar diagonal multivariate GARCH representation mimics the RiskMetrics exponential smoother in Equation (6.4), except for the positive definite  $C$  matrix intercept, and the one additional smoothing parameter. Importantly however, provided that  $\alpha + \beta < 1$ , the unconditional covariance matrix implied by the model in (6.12) equals  $\Omega = (1 - \alpha - \beta)^{-1}C$ , and in parallel to the expression for the univariate GARCH(1, 1) model in Equation (3.9), the  $h$ -period forecasts mean reverts to  $\Omega$  according to the formula,

$$\Omega_{t+h|t} = \Omega + (\alpha + \beta)^{h-1}(\Omega_{t+1|t} - \Omega).$$

This contrasts sharply with the RiskMetrics forecasts, which as previously noted show no mean reversion, with the counterfactual implication that the multi-period covariance forecasts for (approximately) serially uncorrelated returns scale with the forecast horizon. Of course, the scalar model in (6.12) could easily be refined to allow for different (slower) decay rates for the covariances by adding just one or two additional parameters to describe the off-diagonal elements. Still, the model is arguably too simplistic from an empirical perspective, and we will discuss other practically feasible multivariate models and forecasting procedures in the subsequent sections. Before doing so, however, we briefly discuss some of the basic principles and ideas involved in the estimation of multivariate GARCH models.

### 6.3. Multivariate GARCH estimation

Estimation and inference for multivariate GARCH models may formally proceed along the same lines as for the univariate models discussed in Section 3.5. In particular, assume that the conditional distribution of  $Y_t$  is multivariate normal with mean,  $M_{t|t-1}$ ,

and covariance matrix,  $\Omega_{t|t-1}$ . The log-likelihood function is given by the sum of the corresponding  $T$  logarithmic conditional normal densities,

$$\begin{aligned} \log L(\theta; Y_T, \dots, Y_1) &= -\frac{TN}{2} \log(2\pi) \\ &\quad - \frac{1}{2} \sum_{t=1}^T [\log \Omega_{t|t-1}(\theta) - (Y_t - M_{t|t-1}(\theta))' \Omega_{t|t-1}(\theta)^{-1} (Y_t - M_{t|t-1}(\theta))], \end{aligned} \quad (6.13)$$

where we have highlighted the explicit dependence on the parameter vector,  $\theta$ . Provided that the assumption of conditional normality is true and the parametric models for the mean and covariance matrices are correctly specified, the resulting estimates, say  $\hat{\theta}_T$ , will satisfy the usual optimality conditions associated with maximum likelihood. Moreover, even if the conditional normality assumption is violated, the resulting estimates may still be given a QMLE interpretation, with robust parameter inference based on the “sandwich-form” of the covariance matrix estimator, as discussed in Section 3.5.

Meanwhile, as discussed in Section 2, when constructing interval or VaR type forecasts, the whole conditional distribution becomes important. Thus, in parallel to the discussion in Sections 3.5 and 3.6, other multivariate conditional distributions may be used in place of the multivariate normal distributions underlying the likelihood function in (6.13). Different multivariate generalizations of the univariate fat-tailed Student  $t$  distribution in (3.24) have proved quite successful for many daily and weekly financial rate of returns.

The likelihood function in (6.13), or generalizations allowing for conditionally non-normal innovations, may in principle be maximized by any of a number of different numerical optimization techniques. However, even for moderate values of  $N$ , say  $N \geq 5$ , the dimensionality of the problem for the general model in (6.9) or the diagonal vech model in (6.11) renders the computations hopelessly demanding from a practical perspective. As previously noted, this lack of tractability motivates the more parsimonious parametric specifications discussed below.

An alternative approach for circumventing the curse-of-dimensionality within the context of the diagonal vech model has recently been advocated by Ledoit, Santa-Clara and Wolf (2003). Instead of estimating all of the elements in the  $C$ ,  $A$  and  $B$  matrices jointly, inference in their Flex GARCH approach proceed by estimating separate bivariate models for all of the possible pairwise combinations of the  $N$  elements in  $Y_t$ . These individual matrix estimates are then “pasted” together to a full-dimensional model in such a way that the resulting  $N \times N$  matrices in (6.11) are ensured to be positive definite.

Another practical approach for achieving more parsimonious and empirically meaningful multivariate GARCH forecasting models rely on so-called variance targeting techniques. Specifically, consider the general multivariate formulation in (6.9) obtained

1 by replacing  $C$  with

$$2 \quad C = (I - A - B) \text{vech}(V), \quad (6.14) \quad 3$$

4 where  $V$  denotes a positive definite matrix. Provided that the norm of all the eigen- 4  
 5 values for  $A + B$  are less than unity, so that the inverse of  $(I - A - B)$  exists, this 5  
 6 reparameterization implies that the long-run forecasts for  $\Omega_{t+h|t}$  will converge to  $V$  for 6  
 7  $h \rightarrow \infty$ . As such, variance targeting can help ensure that the long-run forecasts are 7  
 8 well behaved. Of course, this doesn't reduce the number of unknown parameters in the 8  
 9 model per se, as the long-run covariance matrix,  $V$ , must now be determined. However, 9  
 10 an often employed approach is to fix  $V$  at the unconditional sample covariance matrix, 10

$$11 \quad \hat{V} = \frac{1}{T} \sum_{t=1}^T (Y_t - \hat{M}_{t|t-1})(Y_t - \hat{M}_{t|t-1})', \quad 12$$

13 where  $\hat{M}_{t|t-1}$  denotes some first-stage estimate for the conditional mean. This esti- 14  
 15 mation of  $V$  obviously introduces an additional source of parameter estimation error 15  
 16 uncertainty, although the impact of this is typically ignored in practice when conducting 16  
 17 inference about the other parameters entering the equation for the conditional covari- 17  
 18 ance matrix. 18  
 19

#### 20 6.4. Dynamic conditional correlations 21

22 One commonly applied approach for large scale dynamic covariance matrix model- 23  
 24 ing and forecasting is the Constant Conditional Correlation (CCC) model of [Bollerslev](#) 24  
 25 (1990). Specifically, let  $D_{t|t-1}$  denote the  $N \times N$  diagonal matrix with the conditional 25  
 26 standard deviations, or the square root of the diagonal elements in  $\Omega_{t|t-1} \equiv \text{Var}(Y_t |$  26  
 27  $\mathcal{F}_{t-1})$ , along the diagonal. The conditional covariance matrix may then be uniquely 27  
 28 expressed in terms of the decomposition, 28

$$29 \quad \Omega_{t|t-1} = D_{t|t-1} \Gamma_{t|t-1} D_{t|t-1}, \quad (6.15) \quad 30$$

31 where  $\Gamma_{t|t-1}$  denote the  $N \times N$  matrix of conditional correlations. Of course, this de- 31  
 32 composition does not result in any immediate simplifications from a modeling perspec- 32  
 33 tive, as the conditional correlation matrix must now be estimated. However, following 33  
 34 [Bollerslev \(1990\)](#) and assuming that the temporal variation in the covariances are driven 34  
 35 solely by the temporal variation in the corresponding conditional standard deviations, 35  
 36 so that the conditional correlations are constant, 36

$$37 \quad \Gamma_{t|t-1} \equiv \Gamma, \quad (6.16) \quad 38$$

39 dramatically reduces the number of parameters in the model relative to the linear vech 39  
 40 specifications discussed above. Moreover, this assumption also greatly simplifies the 40  
 41 multivariate estimation problem, which may now proceed in two steps. In the first step 41  
 42  $N$  individual univariate GARCH models are estimated for each of the series in  $Y_t$ , re- 42  
 43 sulting in an estimate for the diagonal matrix,  $\hat{D}_{t|t-1}$ . Then defining the  $N \times 1$  vector 43

of standardized residuals for each of the univariate series,

$$\hat{\varepsilon}_t \equiv \hat{D}_{t|t-1}^{-1}(Y_t - \hat{M}_{t|t-1}), \tag{6.17}$$

the elements in  $\Gamma$  may simply be estimated by the corresponding sample analogue,

$$\hat{\Gamma} = \frac{1}{T} \sum_{t=1}^T \hat{\varepsilon}_t \hat{\varepsilon}_t'. \tag{6.18}$$

Importantly, this estimate for  $\Gamma$  is guaranteed to be positive definite with ones along the diagonal and all of the other elements between minus one and one. In addition to being simple to implement, this approach therefore has the desirable feature that as long as the individual variances in  $\hat{D}_{t|t-1}$  are positive, the resulting covariance matrices defined by (6.15) are guaranteed to be positive definite.

While the assumption of constant conditional correlations may often be a reasonable simplification over shorter time periods, it is arguable too simplistic in many situations of practical interest. To circumvent this, while retaining the key features of the decomposition in (6.15), Engle (2002) and Tse and Tsui (2002) have recently proposed a convenient framework for directly modeling any temporal dependencies in the conditional correlations. In the most basic version of the Dynamic Conditional Correlation (DCC) model of Engle (2002), the temporal variation in the conditional correlation is characterized by a simple scalar GARCH(1, 1) model, along the lines of (6.12), with the covariance matrix for the standardized residuals targeted at their unconditional value in (6.18). That is,

$$Q_{t|t-1} = (1 - \alpha - \beta)\hat{\Gamma} + \alpha(\hat{\varepsilon}_{t-1}\hat{\varepsilon}_{t-1}') + \beta Q_{t-1|t-2}. \tag{6.19}$$

Although this recursion guarantees that the  $Q_{t|t-1}$  matrices are positive definite, the individual elements are not necessarily between minus one and one. Thus, in order to arrive at an estimate for the conditional correlation matrix, the elements in  $Q_{t|t-1}$  must be standardized, resulting in the following estimate for the  $ij$ th correlation:

$$\hat{\rho}_{ij,t} \equiv \{\hat{\Gamma}_{t|t-1}\}_{ij} = \frac{\{Q_t\}_{ij}}{\{Q_t\}_{ii}^{1/2}\{Q_t\}_{jj}^{1/2}}. \tag{6.20}$$

Like the CCC model, the DCC model is also relatively simple to implement in large dimensions, requiring only the estimation of  $N$  univariate models along with a choice of the two exponential smoothing parameters in (6.19).

Richer dynamic dependencies in the correlations could be incorporated in a similar manner, although this immediately raises some of the same complications involved in directly parameterizing  $\Omega_{t|t-1}$ . However, as formally shown in Engle and Sheppard (2001), the parameters in (6.19) characterizing the dynamic dependencies in  $Q_{t|t-1}$ , and in turn  $\Gamma_{t|t-1}$ , may be consistently estimated in a second step by maximizing the partial log-likelihood function,

$$\log L(\theta; Y_T, \dots, Y_1)^* = -\frac{1}{2} \sum_{t=1}^T [\log |\Gamma_{t|t-1}(\theta)| - \hat{\varepsilon}_t' \Gamma_{t|t-1}(\theta)^{-1} \hat{\varepsilon}_t],$$

where  $\hat{\varepsilon}_t$  refers to the first step estimates defined in (6.17). Of course, the standard errors for the resulting correlation parameter estimates must be adjusted to take account of the first stage estimation errors in  $\hat{D}_{t|t-1}$ . Extensions of the basic DCC structure in (6.19) and (6.20) along these lines allowing for greater flexibility in the dependencies in the correlations across different types of assets, asymmetries in the way in which the correlations respond to past negative and positive return innovations, regime switches in the correlations, to name but a few, are currently being actively explored by a number of researchers.

### 6.5. Multivariate stochastic volatility and factor models

An alternative approach for achieving a more manageable and parsimonious multivariate volatility forecasting model entails the use of factor structures. Factor structures are, of course, central to the field of finance, and the Arbitrage Pricing Theory (APT) in particular. Multivariate factor GARCH and stochastic volatility models were first analyzed by Diebold and Nerlove (1989) and Engle, Ng and Rothschild (1990). To illustrate, consider a simple one-factor model in which the commonality in the volatilities across the  $N \times 1$   $R_t$  vector of asset returns is driven by a single scalar factor,  $f_t$ ,

$$R_t = a + bf_t + e_t, \quad (6.21)$$

where  $a$  and  $b$  denote  $N \times 1$  parameter vectors, and  $e_t$  is assumed to be i.i.d. through time with covariance matrix  $\Lambda$ . This directly captures the idea that variances (and covariances) generally move together across assets. Now, assuming that the factor is conditionally heteroskedastic, with conditional variance denoted by  $\sigma_{t|t-1}^2 \equiv \text{Var}(f_t | \mathcal{F}_{t-1})$ , the conditional covariance matrix for  $R_t$  takes the form

$$\Omega_{t|t-1} \equiv \text{Var}(R_t | \mathcal{F}_{t-1}) = bb' \sigma_{t|t-1}^2 + \Lambda. \quad (6.22)$$

Compared to the unrestricted GARCH models discussed in Section 6.2, the factor GARCH representation greatly reduces the number of free parameters. Moreover, the conditional covariance matrix in (6.22) is guaranteed to be positive definite.

To further appreciate the implications of the factor representation, let  $b_i$  and  $\lambda_{ij}$  denote the  $i$ th and  $ij$ th element in  $b$  and  $\Lambda$ , respectively. It follows then directly from the expression in (6.22) that the conditional correlation between the  $i$ th and the  $j$ th observation is given by

$$\rho_{ij,t} \equiv \frac{b_i b_j \sigma_{t|t-1}^2 + \lambda_{ij}}{(b_i^2 \sigma_{t|t-1}^2 + \lambda_{ii})^{1/2} (b_j^2 \sigma_{t|t-1}^2 + \lambda_{jj})^{1/2}}. \quad (6.23)$$

Thus, provided that the corresponding factor loadings are of the same sign, or  $b_i b_j > 0$ , the conditional correlation implied by the model will increase toward unity as the volatility of the factor increases. That is, there is an empirically realistic built-in volatility-in-correlation effect.

Importantly, multivariate conditional covariance matrix forecasts are also readily constructed from forecasts for the univariate factor variance. In particular, assuming that

the vector of returns is serially uncorrelated, the conditional covariance matrix for the  $k$ -period continuously compounded returns is simply given by

$$\Omega_{t:t+k|t} \equiv \text{Var}(R_{t+k} + \cdots + R_{t+1} | \mathcal{F}_t) = bb' \sigma_{t:t+k|t}^2 + k \Lambda, \quad (6.24)$$

where  $\sigma_{t:t+k|t}^2 \equiv \text{Var}(f_{t+k} + \cdots + f_{t+1} | \mathcal{F}_t)$ . Further assuming that the factor is directly observable and that the conditional variance for  $f_t$  is specified in terms of the observable information set,  $\mathcal{F}_{t-1}$ , the forecasts for  $\sigma_{t:t+k|t}^2$  may be constructed along the lines of the univariate GARCH class of models discussed in Section 3. If, on the other hand, the factor is latent or if the conditional variance for  $f_t$  is formulated in terms of unobservable information,  $\mathfrak{F}_{t-1}$ , one of the more intensive numerical procedures for the univariate stochastic volatility class of models discussed in Section 4 must be applied in calculating  $\sigma_{t:t+k|t}^2$ . Of course, the one-factor model in (6.21) could easily be extended to allow for multiple factors, resulting in obvious generalizations of the expressions in (6.22) and (6.24). As long as the number of factors remain small, the same appealing simplifications hold true.

Meanwhile, an obvious drawback from an empirical perspective to the simple factor model in (6.21) with homoskedastic innovations concerns the lack of heteroskedasticity in certain portfolios. Specifically, let  $\Psi = \{\psi \mid \psi' b = 0, \psi \neq 0\}$  denote the set of  $N \times 1$  vectors orthogonal to the vector of factor loadings,  $b$ . Any portfolio constructed from the  $N$  original assets with portfolio weights,  $w = \psi / (\psi_1 + \cdots + \psi_N)$  where  $\psi \in \Psi$ , will then be homoskedastic,

$$\text{Var}(r_{w,t} | \mathcal{F}_{t-1}) \equiv \text{Var}(w' R_t | \mathcal{F}_{t-1}) = w' b b' w \sigma_{t|t-1}^2 + w' \Lambda w = w' \Lambda w. \quad (6.25)$$

Similarly, the corresponding multi-period forecasts defined in (6.24) will also be time invariant. Yet, in applications with daily or weekly returns it is almost always impossible to construct portfolios which are void of volatility clustering effects. The inclusion of additional factors does not formally solve the problem. As long as the number of factors is less than  $N$ , the corresponding null-set  $\Psi$  is not empty. Of course, allowing the covariance matrix of the idiosyncratic innovations to be heteroskedastic would remedy this problem, but that then raises the issue of how to model the temporal variation in the  $(N \times N)$ -dimensional  $\Lambda_t$  matrix. One approach would be to include enough factors so that the  $\Lambda_t$  matrix may be assumed to be diagonal, only requiring the estimation of  $N$  univariate volatility models for the elements in  $e_t$ .

Whether the rank deficiency in the forecasts of the conditional covariance matrices from the basic factor structure and the counterfactual implication of no volatility clustering in certain portfolios discussed above should be a cause for concern ultimately depends upon the uses of the forecasts. However, it is clear that the reduction in the dimension of the problem to a few systematic risk factors may afford great computational simplifications in the context of large scale covariance matrix modeling and forecasting.



### 6.6. Realized covariances and correlations

The high-frequency data realized volatility approach for measuring, modeling and forecasting univariate volatilities outlined in Section 5 may be similarly adapted to modeling and forecasting covariances and correlations. To set out the basic idea, let  $R(t, \Delta)$  denote the  $N \times 1$  vector of logarithmic returns over the  $[t - \Delta, t]$  time interval,

$$R(t, \Delta) \equiv P(t) - P(t - \Delta). \quad (6.26)$$

The  $N \times N$  realized covariation matrix for the unit time interval,  $[t - 1, t]$ , is then formally defined by

$$\text{RCOV}(t, \Delta) = \sum_{j=1}^{1/\Delta} R(t - 1 + j \cdot \Delta, \Delta) R(t - 1 + j \cdot \Delta, \Delta)'. \quad (6.27)$$

This directly parallels the univariate definition in (5.10). Importantly, the realized covariation matrix is symmetric by construction, and as long as the returns are linearly independent and  $N < 1/\Delta$ , the matrix is guaranteed to be positive definite.

In order to more formally justify the realized covariation measure, suppose that the evolution of the  $N \times 1$  vector price process may be described by the  $N$ -dimensional continuous-time diffusion,

$$dP(t) = M(t) dt + \Sigma(t) dW(t), \quad t \in [0, T], \quad (6.28)$$

where  $M(t)$  denotes the  $N \times 1$  instantaneous drifts,  $\Sigma(t)$  refer to the  $N \times N$  instantaneous diffusion matrix, and  $W(t)$  now denotes an  $(N \times 1)$ -dimensional vector of independent standard Brownian motions. Intuitively, for small values of  $\Delta > 0$ ,

$$R(t, \Delta) \equiv P(t) - P(t - \Delta) \simeq M(t - \Delta) \Delta + \Sigma(t - \Delta) \Delta W(t), \quad (6.29)$$

where  $\Delta W(t) \equiv W(t) - W(t - \Delta) \sim N(0, \Delta I_N)$ . Of course, this latter expression directly mirrors the univariate equation (5.2). Now, using similar arguments to the ones in Section 5.1, it follows that the multivariate realized covariation in (6.27) will converge to the corresponding multivariate integrated covariation for finer and finer sampled high-frequency returns, or  $\Delta \rightarrow 0$ ,

$$\text{RCOV}(t, \Delta) \rightarrow \int_{t-1}^t \Sigma(s) \Sigma(s)' ds \equiv \text{ICOV}(t). \quad (6.30)$$

Again, by similar arguments to the ones in Section 5.1, the multivariate integrated covariation defined by the right-hand side of Equation (6.30) provides the true measure for the actual return variation and *covariation* that transpired over the  $[t - 1, t]$  time interval. Also, extending the univariate results in (5.12), Barndorff-Nielsen and Shephard (2004b) have recently shown that the multivariate realized volatility errors,  $\sqrt{1/\Delta}[\text{RCOV}(t, \Delta) - \text{ICOV}(t)]$ , are approximately serially uncorrelated and asymptotically (for  $\Delta \rightarrow 0$ ) distributed as a mixed normal with a random covariance matrix

1 that may be estimated. Moreover following (5.13), the consistency of the realized covariation measure for the true quadratic covariation carries over to situations in which  
 2 the vector price process contains jumps. As such, these theoretical results set the stage  
 3 for multivariate volatility modeling and forecasting based on the realized covariation  
 4 measures along the same lines as the univariate discussion in Sections 5.2 and 5.3.

5 In particular, treating the  $\frac{1}{2}N(N + 1) \times 1$  vector,  $\text{vech}[\text{RCOV}(t, \Delta)]$ , as a direct  
 6 observation (with uncorrelated measurement errors) on the unique elements in the covariation  
 7 matrix of interest, standard multivariate time series techniques may be used in  
 8 jointly modeling the variances and the off-diagonal covariance elements. For instance, a  
 9 simple VAR(1) forecasting model, analogues to the GARCH(1, 1) model in (6.9), may  
 10 be specified as

$$11 \quad \text{vech}[\text{RCOV}(t, \Delta)] = C + A \text{vech}[\text{RCOV}(t - 1, \Delta)] + u_t, \quad (6.31) \quad 12$$

13 where  $u_t$  denotes an  $N \times 1$  vector white noise process. Of course, higher order dynamic  
 14 dependencies could be included in a similar manner. Indeed, the results in Andersen  
 15 et al. (2001b, 2003), suggest that for long-run forecasting it may be important to incor-  
 16 porate long-memory type dependencies in both variances and covariances. This could  
 17 be accomplished through the use of a true multivariate fractional integrated model, or  
 18 as previously discussed an approximating component type structure.

19 Even though  $\text{RCOV}(t, \Delta)$  is positive definite by construction, nothing guarantees  
 20 that the forecasts from an unrestricted multivariate time series model along the lines  
 21 of the VAR(1) in (6.31) will result in positive definite covariance matrix forecasts.  
 22 Hence, it may be desirable to utilize some of the more restrictive parameterizations  
 23 for the multivariate GARCH models discussed in Section 6.2, to ensure positive defi-  
 24 nite covariance matrix forecasts. Nonetheless, replacing  $\Omega_{t|t-1}$  with the directly ob-  
 25 servable  $\text{RCOV}(t, \Delta)$ , means that the parameters in the corresponding models may  
 26 be estimated in a straightforward fashion using simple least squares, or some other  
 27 easy-to-implement estimation method, rather than the much more numerically intensive  
 28 multivariate MLE or QMLE estimation schemes.

29 Alternatively, an unrestricted model for the  $\frac{1}{2}N(N + 1)$  nonzero elements in the  
 30 Cholesky decomposition, or lower triangular matrix square-root, of  $\text{RCOV}(t, \Delta)$ , could  
 31 also be estimated. Of course, the nonlinear transformation involved in such a decompo-  
 32 sition means that the corresponding matrix product of the forecasts from the model will  
 33 generally not be unbiased for the elements in the covariation matrix itself. Following  
 34 Andersen et al. (2003), sometimes it might also be possible to infer the covariances of  
 35 interest from the variances of different cross-rates or portfolios through appropriately  
 36 defined arbitrage conditions. In those situations forecasts for the covariances may there-  
 37 fore be constructed from a set of forecasting models for the corresponding variances, in  
 38 turn avoiding directly modeling any covariances.

39 The realized covariation matrix in (6.27) may also be used in the construction of re-  
 40 alized correlations, as in Andersen et al. (2001a, 2001b). These realized correlations  
 41 could be modeled directly using standard time series techniques. However, the corre-  
 42 lations are, of course, restricted to lie between minus one and one. Thus, to ensure  
 43

1 that this constraint is not violated, it might be desirable to use the Fisher transform, 1  
2  $z = 0.5 \cdot \log[(1 + \rho)/(1 - \rho)]$ , or some other similar transformation, to convert the sup- 2  
3 port of the distribution for the correlations from  $[-1, 1]$  to the whole real line. This is 3  
4 akin to the log transformation for the univariate realized volatilities employed in Equa- 4  
5 tion (5.14). Meanwhile, there is some evidence that the dynamic dependencies in the 5  
6 correlations between many financial assets and markets are distinctly different from 6  
7 that of the corresponding variances and covariances, exhibiting occasional “correla- 7  
8 tion breakdowns”. These types of dependencies may best be characterized by regime 8  
9 switching type models. Rather than modeling the correlations individually, the realized 9  
10 correlation matrix could also be used in place of  $\hat{\epsilon}_t \hat{\epsilon}'_t$  in the DCC model in (6.19), or 10  
11 some generalization of that formulation, in jointly modeling all of the elements in the 11  
12 conditional correlation matrix. 12

13 The realized covariation and correlation measures discussed above are, of course, 13  
14 subject to the same market microstructure complications that plague the univariate re- 14  
15 alized volatility measures discussed in Section 5. In fact, some of the problems are 15  
16 accentuated with the possibility of nonsynchronous observations in two or more mar- 16  
17 kets. Research on this important issues is still very much ongoing, and it is too early 17  
18 to draw any firm conclusions about the preferred method or sampling scheme to em- 18  
19 ploy in the practical implementation of the realized covariation measures. Nonetheless, 19  
20 it is clear that the realized volatility approach afford a very convenient and powerful ap- 20  
21 proach for effectively incorporating high-frequency financial data into both univariate 21  
22 and multivariate volatility modeling and forecasting. 22

### 23 6.7. Further reading 23

24 The use of historical pseudo returns as a convenient way of reducing the multivariate 24  
25 modeling problem to a univariate setting, as outlined in Section 6.1, is discussed at some 25  
26 length in Andersen et al. (2005). This same study also discusses the use of a smaller set 26  
27 of liquid base assets along with a factor structure as another computationally conven- 27  
28 ient way of reducing the dimension of time-varying covariance matrix forecasting for 28  
29 financial rate of returns. 29  
30 30

31 The RiskMetrics, or exponential smoothing approach, for calculating covariances and 31  
32 associated Value-at-Risk measures is discussed extensively in Christoffersen (2003), 32  
33 Jorion (2000), and Zaffaroni (2004) among others. Following earlier work by DeSantis 33  
34 and Gerard (1997), empirically more realistic slower decay rates for the covariances in 34  
35 the context of exponential smoothing has been successfully implemented by DeSantis 35  
36 et al. (2003). 36  
37 37

38 In addition to the many ARCH and GARCH survey papers and book treatments 38  
39 listed in Section 3, the multivariate GARCH class of models has recently been surveyed 39  
40 by Bauwens, Laurent and Rombouts (2005). A comparison of some of the available 40  
41 commercial software packages for the estimation of multivariate GARCH models is 41  
42 available in Brooks, Burke and Persaud (2003). Conditions for the covariance matrix 42  
43 forecasts for the linear formulations discussed in Section 6.2 to be positive definite was 43

1 first established by Engle and Kroner (1995), who also introduced the so-called BEKK 1  
2 parameterization. Asymmetries, or leverage effects, within this same class of models 2  
3 were subsequently studied by Kroner and Ng (1998). The bivariate EGARCH model of 3  
4 Braun, Nelson and Sunier (1995) and the recent matrix EGARCH model of Kawakatsu 4  
5 (2005) offer alternative ways of doing so. The multivariate GARCH QMLE procedures 5  
6 outlined in Section 6.3 were first discussed by Bollerslev and Wooldridge (1992), while 6  
7 Ling and McAleer (2003) provide a more recent account of some of the subsequent 7  
8 important theoretical developments. The use of a fat tailed multivariate Student  $t$  distri- 8  
9 bution in the estimation of multivariate GARCH models was first considered by Harvey, 9  
10 Ruiz and Sentana (1992); see also Bauwens and Laurent (2005) and Fiorentini, Sentana 10  
11 and Calzolari (2003) for more recent applications of alternative multivariate nonnormal 11  
12 distributions. Issues related to cross-sectional and temporal aggregation of multivariate 12  
13 GARCH and stochastic volatility models have been studied by Nijman and Sentana 13  
14 (1996) and Meddahi and Renault (2004). 14

15 Several empirical studies have documented important temporal dependencies in asset 15  
16 return correlations, including early contributions by Erb, Harvey and Viskanta (1994) 16  
17 and Longin and Solnik (1995) focusing on international equity returns. More recent 17  
18 work by Ang and Chen (2002) and Cappiello, Engle and Sheppard (2004) have em- 18  
19 phasized the importance of explicitly incorporating asymmetries in the way in which 19  
20 the correlations respond to past negative and positive return shocks. Along these lines, 20  
21 Longin and Solnik (2001) report evidence in support of more pronounced dependen- 21  
22 cies following large (extreme) negative return innovations. A test for the assumption of 22  
23 constant conditional correlations underlying the CCC model discussed in Section 6.4 23  
24 has been derived by Bera and Kim (2002). Recent work on extending the DCC model 24  
25 to allow for more flexible dynamic dependencies in the correlations, asymmetries in 25  
26 the responses to past negative and positive returns, as well as switches in the corre- 26  
27 lations across regimes, include Billio, Caporin and Gobbo (2003), Cappiello, Engle 27  
28 and Sheppard (2004), Franses and Hafner (2003), and Pelletier (2005). Guidolin and 28  
29 Timmermann (2005b) find large variations in the correlation between stock and bond 29  
30 returns across different market regimes defined as crash, slow growth, bull and recovery. 30  
31 Sheppard (2004) similarly finds evidence of business cycle frequency dynamics in 31  
32 conditional covariances. 32

33 The factor ARCH models proposed by Diebold and Nerlove (1989) and Engle, 33  
34 Ng and Rothschild (1990) have been used by Ng, Engle and Rothschild (1992) and 34  
35 Bollerslev and Engle (1993), among others, in modeling common persistence in condi- 35  
36 tional variances and covariances. Harvey, Ruiz and Sheppard (1994) and King, Sentana 36  
37 and Wadhvani (1994) were among the first to estimate multivariate stochastic volatili- 37  
38 ty models. More recent empirical studies and numerically efficient algorithms for the 38  
39 estimation of latent multivariate volatility structures include Aguilar and West (2000), 39  
40 Fiorentini, Sentana and Sheppard (2004), and Liesenfeld and Richard (2003). Issues 40  
41 related to identification within heteroskedastic factor models have been studied by 41  
42 Sentana and Fiorentini (2001). A recent insightful discussion of the basic features of 42  
43 multivariate stochastic volatility factor models, along with a discussion of their ori- 43

gins, is provided in Shephard (2004). The multivariate Markov-switching multifractal model of Calvet, Fisher and Thompson (2005) may also be interpreted as a latent factor stochastic volatility model with a closed form likelihood. Other related relatively easy-to-implement multivariate approaches include the two-step Orthogonal GARCH model of Alexander (2001), in which the conditional covariance matrix is determined by univariate models for a (small) set of the largest (unconditional) principal components.

The realized volatility approach discussed in Section 6.6 affords a simple practically feasible way for covariance and correlation forecasting in situations when high-frequency data is available. The formal theory underpinning this approach in the multivariate setting has been spelled out in Andersen et al. (2003) and Barndorff-Nielsen and Shephard (2004b). A precursor to some of these results is provided by the alternative double asymptotic rolling regression based framework in Foster and Nelson (1996). The benefits of the realized volatility approach versus more conventional multivariate GARCH based forecasts in the context of asset allocation have been forcefully demonstrated by Fleming, Kirby and Ostdiek (2003). Meanwhile, the best way of actually implementing the realized covariation measures with high-frequency financial data subject to market microstructure frictions still remains very much of an open research question. In a very early paper, Epps (1979) first observed a dramatic drop in high-frequency based sample correlations among individual stock returns as the length of the return interval approached zero; see also Lundin, Dacorogna and Müller (1998). In addition to the many mostly univariate studies noted in Section 4, Martens (2003) provides a recent assessment and comparison of some of the existing ways for best alleviating the impact of market microstructure frictions in the multivariate setting, including the covariance matrix estimator of De Jong and Nijman (1997), the lead-lag adjustment of Scholes and Williams (1977), and the range-based covariance measure of Brandt and Diebold (2005).

The multivariate procedures discussed in this section are (obviously) not exhaustive of the literature. Other recent promising approaches for covariance and correlation forecasting include the use of copulas for conveniently linking univariate GARCH [e.g., Jondeau and Rockinger (2005) and Patton (2004)] or realized volatility models; the use of shrinkage to ensure positive definiteness in the estimation and forecasting of very large-dimensional covariance matrices [e.g., Jagannathan and Ma (2003) and Ledoit and Wolf (2003)]; and forecast model averaging techniques [e.g., Pesaran and Zaffaroni (2004)]. It remains to be seen which of these, if any, will be added to the standard multivariate volatility forecasting toolbox.

## 7. Evaluating volatility forecasts

The discussion up until this point has surveyed the most important univariate and multivariate volatility models and forecasting procedures in current use. This section gives an overview of some of the most useful methods available for volatility forecast evaluation. The methods introduced here can either be used by an external evaluator or by

the forecaster him/herself as a diagnostic tool on the forecasting model. A more general discussion and review of the forecast evaluation literature can be found in [Diebold and Lopez \(1996\)](#) and [Chapter 3](#) by West in this Handbook.

Below, we will first introduce a general loss function framework and then highlight the particular issues involved when forecasting volatility itself is the direct object of interest. We then discuss several other important forecasting situations where volatility dynamics are crucial, including Value-at-Risk, probability, and density forecasting.

### 7.1. Point forecast evaluation from general loss functions

Consider the general forecast loss function,  $L(y_{t+1}, \hat{y}_{t+1|t})$ , discussed in [Section 2](#), in which the arguments are the univariate discrete-time real-valued stochastic variable,  $y_{t+1}$ , as well as its forecast,  $\hat{y}_{t+1|t}$ . From the optimization problem solved by the optimal forecast,  $\hat{y}_{t+1|t}$  must solve the generic first order condition

$$E_t \left[ \frac{\partial L(y_{t+1}, \hat{y}_{t+1|t})}{\partial \hat{y}} \right] = 0. \quad (7.1)$$

The partial derivative of the loss function – the term inside the conditional expectation – is sometimes referred to as the generalized forecast error. Realizations of this partial derivative should fluctuate unpredictably around zero, directly in line with the standard optimality condition that regular forecasts display uncorrelated prediction errors.

Specifically, consider the situation in which we observe a sequence of out-of-sample forecasts and subsequent realizations,  $\{y_{t+1}, \hat{y}_{t+1|t}\}_{t=1}^T$ . A natural diagnostic on [\(7.1\)](#) is then given by the simple regression version of the conditional expectation, that is

$$\frac{\partial L(y_{t+1}, \hat{y}_{t+1|t})}{\partial \hat{y}} = a + b'x_t + \varepsilon_{t+1}, \quad (7.2)$$

where  $x_t$  denotes a vector of candidate explanatory variables in the time  $t$  information set observed by the forecaster,  $\mathcal{F}_t$ , and  $b$  is a vector of regression coefficients. An appropriately calibrated forecast should then have  $a = b = 0$ , which can be tested using standard  $t$ - and  $F$ -tests properly robustified to allow for heteroskedasticity in the regression errors,  $\varepsilon_{t+1}$ . Intuitively, if a significant coefficient is obtained on a forecasting variable, which the forecaster should reasonably have known at time  $t$ , then the forecasting model is not optimal, as the variable in question could and should have been used to make the generalized forecast error variance lower than it actually is.

If the forecasts arise from a known well-specified statistical model with estimated parameters then the inherent parameter estimation error should ideally be accounted for. This can be done using the asymptotic results in [West and McCracken \(1998\)](#) or the finite sample Monte Carlo tests in [Dufour \(2004\)](#). However, external forecast evaluators may not have knowledge of the details of the underlying forecasting model (if one exists) in which case parameter estimation error uncertainty is not easily accounted for. Furthermore, in most financial applications the estimation sample is typically fairly large rendering the parameter estimation error relatively small compared with other

potentially more serious model specification errors. In this case standard (heteroskedasticity robust)  $t$ -tests and  $F$ -tests may work reasonably well. Note also that in the case of, say,  $h$ -day forecasts from a daily model, the horizon overlap implies that the first  $h - 1$  autocorrelations will not be equal to zero, and this must be allowed for in the regression.

As an example of the general framework in (7.2), consider the case of quadratic loss,  $L(y_{t+1}, \hat{y}_{t+1|t}) = (y_{t+1} - \hat{y}_{t+1|t})^2$ . In this situation

$$\frac{\partial L(y_{t+1}, \hat{y}_{t+1|t})}{\partial \hat{y}} = -2(y_{t+1} - \hat{y}_{t+1|t}), \quad (7.3)$$

which suggests the forecast evaluation regression

$$(y_{t+1} - \hat{y}_{t+1|t}) = a + b'x_t + \varepsilon_{t+1}. \quad (7.4)$$

While the choice of information variables to include in  $x_t$  is somewhat arbitrary, one obvious candidate does exist, namely the time  $t$  forecast itself. Following this idea and letting  $x_t = \hat{y}_{t+1|t}$ , results in the so-called [Mincer and Zarnowitz \(1969\)](#) regression, which can thus be viewed as a test of forecast optimality relative to a limited information set. We write

$$(y_{t+1} - \hat{y}_{t+1|t}) = a + b\hat{y}_{t+1|t} + \varepsilon_{t+1},$$

or equivalently

$$y_{t+1} = a + (b + 1)\hat{y}_{t+1|t} + \varepsilon_{t+1}. \quad (7.5)$$

Clearly the ex-ante forecast should not be able to explain the ex-post forecast error. For example, if  $b$  is significantly negative, and thus  $(b + 1) < 1$ , then the forecast is too volatile relative to the subsequent realization and the forecast should be scaled down.

It is often of interest to compare forecasts from different models, or forecasters. This is easily done by letting  $x_t = [\hat{y}_{t+1|t} \ \hat{y}_{A,t+1|t}]$ , where  $\hat{y}_{A,t+1|t}$  denotes the alternative forecast. The forecast evaluation regression then takes the form,

$$y_{t+1} = a + (b + 1)\hat{y}_{t+1|t} + b_A\hat{y}_{A,t+1|t} + \varepsilon_{t+1}, \quad (7.6)$$

where a failure to reject the hypothesis that  $b_A = 0$  implies that the additional information provided by the alternative forecast is not significant. Or, in other words, the benchmark forecast encompasses the alternative forecast.

## 7.2. Volatility forecast evaluation

The above discussion was cast at a general level. We now turn to the case in which volatility itself is the forecasting object of interest. Hence,  $y_{t+1} \equiv \sigma_{t:t+1}^2$  now refers to some form of ex-post volatility measure, while  $y_{t+1|t} \equiv \hat{\sigma}_{t:t+1|t}^2$  denotes the corresponding ex-ante volatility forecast.

The regression-based framework from above then suggests the general volatility forecast evaluation regression

$$\sigma_{t:t+1}^2 - \hat{\sigma}_{t:t+1|t}^2 = a + b'x_t + \varepsilon_{t+1}, \quad (7.7)$$

or as a special case the Mincer–Zarnowitz volatility regression

$$\sigma_{t:t+1}^2 = a + (b + 1)\hat{\sigma}_{t:t+1|t}^2 + \varepsilon_{t+1},$$

where an optimal forecast would satisfy  $a = b = 0$ . Immediately, however, the question arises of how to actually measure the ex-post variance? As discussed at some length in Sections 1 and 5, the “true” variance, or volatility, is inherently unobservable, and we are faced with the challenge of having to rely on a proxy in order to assess the forecast quality.

The simplest proxy is the squared observation of the underlying variable,  $y_{t+1}^2$ , which, when the mean is zero, has the property of being (conditionally) unbiasedness, or  $E_t[y_{t+1}^2] = \sigma_{t:t+1}^2$ . Thus, the accuracy of the volatility forecasts could be assessed by the following simple regression:

$$y_{t+1}^2 = a + (b + 1)\hat{\sigma}_{t:t+1|t}^2 + \varepsilon_{t+1}. \quad (7.8)$$

However, as illustrated by Figure 1, the squared observation typically provides a very noisy proxy for the true (latent) volatility process of interest. We are essentially estimating the variance each period using just a single observation, and the corresponding regression fit is inevitably very low, even if the volatility forecast is accurate. For instance, regressions of the form (7.8), using daily or weekly squared returns as the left-hand side independent variable, typically result in unspectacular  $R^2$ 's of around 5–10%. We are seemingly stuck with an impossible task, namely to precisely assess the forecastability of something which is itself not observed.

Fortunately, Figure 1 and the accompanying discussion in Sections 1 and 5 suggest a workable solution to this conundrum. In financial applications observations are often available at very high frequencies. For instance, even if the forecaster is only interested in predicting volatility over daily or longer horizons, observations on the asset prices are often available at much finer intradaily sampling frequencies, say  $1/\Delta \gg 1$  observations per “day” or unit time interval. Hence, in this situation following the discussion in Section 5.1, a proxy for the (latent) daily volatility may be calculated from the intradaily squared return as

$$RV(t + 1, \Delta) \equiv \sum_{j=1}^{1/\Delta} [p(t + j \cdot \Delta) - p(t + (j - 1) \cdot \Delta)]^2.$$

The resulting forecast evaluation regression thus takes the form,

$$RV(t + 1, \Delta) = a + (b + 1)\hat{\sigma}_{t:t+1|t}^2 + \varepsilon_{t+1}, \quad (7.9)$$

which coincides with (7.8) for  $\Delta = 1$ . However, in contrast to the low  $R^2$ 's associated with (7.8), Andersen and Bollerslev (1998a) find that in liquid markets the  $R^2$  of the regression in (7.9) can be as high as 50% for the very same volatility forecast that produces an  $R^2$  of only 5–10% in the former regression! In other words, even a reasonably accurate volatility forecasting model will invariably appear to have a low degree of forecastability when evaluated on the basis of a noisy volatility proxy. Equally important,



1 it will be difficult to detect a poor volatility forecasting model when a noisy volatility  
2 proxy is used.

3 Reliable high-frequency information is, of course, not available for all financial mar-  
4 kets. Still, intra-day high and low prices, or quotes, are often available over long his-  
5 torical time periods. Under idealized conditions – a Geometric Brownian motion with  
6 a constant diffusive volatility  $\sigma$  – the expected value of the log range (the difference  
7 between the high and the low logarithmic price) over the unit time interval is directly  
8 related to volatility of the process by the equation

$$9 \quad E[(\max\{p(\tau) \mid t \leq \tau < t + 1\} - \min\{p(\tau) \mid t \leq \tau < t + 1\})^2] \\ 10 \quad = 4 \log(2)\sigma^2. \quad (7.10)$$

11 Hence, a range-based proxy for the per-period volatility is naturally defined by

$$12 \quad \sigma_{r,t:t+1}^2 = \frac{1}{4 \log(2)} \\ 13 \quad \times (\max\{p(\tau) \mid t \leq \tau < t + 1\} - \min\{p(\tau) \mid t \leq \tau < t + 1\})^2. \quad (7.11)$$

14 It is readily seen that, under ideal conditions, this range-based volatility proxy is inferior  
15 to the realized variance measure constructed with a large number of intraday observa-  
16 tions, or  $\Delta \ll 1$ . However, as previously discussed, a host of market microstructure and  
17 other complications often render practical situations less than ideal. Thus, even when  
18 high-frequency data are available, the range-based volatility forecast evaluation regres-  
19 sion,

$$20 \quad \sigma_{r,t:t+1}^2 = a + (b + 1)\hat{\sigma}_{r,t+1|t}^2 + \varepsilon_{t+1}, \quad (7.12)$$

21 may still provide a useful robust alternative, or complement, to the realized volatility  
22 regression in (7.9).

23 To illustrate, consider Figure 7, which graphs a simulated geometric Brownian motion  
24 price process during a “24 hour” or “288 five-minute” period. The “fundamental”,  
25 but unobserved, price process is given by the dashed line. In practice, however, we  
26 only observe this fundamental price plus a random bid–ask spread, as indicated by the  
27 jagged solid line in the figure. The figure conveys several important insights. First, notice  
28 that the squared daily return is small (close to zero) even though there are large  
29 within-day price fluctuations. As such, the true but unobserved volatility is fairly high,  
30 and poorly estimated by the daily squared return. Second, the bid–ask bounces effect  
31 introduces artificial volatility in the observed prices. As a result, realized volatilities  
32 based on very finely sampled high-frequency squared returns produce upward biased  
33 volatility measures. As previously discussed, it is, of course, possible to adjust for this  
34 bias, and several procedures for doing so have recently been proposed in the litera-  
35 ture. Nonetheless, the figure highlights the dangers of using too small a value for  $\Delta$   
36 in the realized volatility estimation without accounting for the bid–ask spread effect.  
37  
38  
39  
40  
41  
42  
43




Figure 7. Simulated fundamental and observed intraday prices. The smooth dashed line represents the fundamental, but unobserved, simulated asset price. The jagged solid line solid represents the observed transaction prices reflecting bid or ask driven transactions. The two horizontal lines denote the min and max prices observed within the day.

Third, the bid–ask spread only affects the range-based measure (the difference between the two horizontal lines) twice as opposed to  $1/\Delta$  times for every high-frequency return entering the realized volatility calculation. As such, the range affords a more robust (to market microstructure frictions) volatility measure. Meanwhile, an obvious drawback to the range-based volatility measure is that the multiplicative adjustment in Equation (7.11) only provides an unbiased measure for integrated volatility under the ideal, and empirically unrealistic, assumption of a geometric Brownian motion, and the “right” multiplication factor is generally unknown. Moreover, extensions to multivariate settings and covariance estimation is difficult to contemplate in the context of the range.

The preceding discussion highlights the need for tools to help in choosing the value of  $\Delta$  in the realized volatility measure. To this end Andersen et al. (1999, 2000) first proposed the “volatility signature plot”, as a simple indicative graphical tool. The signature plot provides a graphical representation of the realized volatility averaged over multiple days as a function of the sampling frequency,  $\Delta$ , going from very high (say one-minute intervals) to low (say daily) frequencies. Recognizing that the bid–ask spread (and other frictions) generally bias the realized volatility measure, this suggests choosing the highest frequency possible for which the average realized volatility appears to have stabilized. To illustrate, Figure 8 shows a simulated example corresponding to the somewhat exaggerated market microstructure effects depicted in Figure 7. In this situation the plot suggests a sampling frequency of around “120 to 180 minutes” or “2 to




Figure 8. Volatility signature plot. The figure depicts the impact of the bid–ask spread for measuring realized volatility by showing the unconditional sample means for the realized volatilities as a function of the length of the return interval for the high-frequency data underlying the calculations. The simulated prices are subject to bid–ask bounce effects shown in Figure 7.

3 hours”. Meanwhile, the actual empirical evidence for a host of actively traded assets indicate that fixing  $\Delta$  somewhere between 5 and 15 minutes typically works well, but many other more refined procedures for eliminating the systematic bias in the simple realized volatility estimator are now also available.

7.3. Interval forecast and Value-at-Risk evaluation

We now discuss situations where the dynamic volatility constitutes an important part of the forecast, but the volatility itself is not the direct object of interest, leading examples of which include Value-at-Risk and probability forecasts. Specifically, consider the interval forecasts of the form discussion in Section 2,

$$\hat{y}_{t+1|t} \equiv \{ \hat{y}_{t+1|t}^L, \hat{y}_{t+1|t}^U \}, \tag{7.13}$$

where the lower and upper parts of the interval forecast are defined so that there is a  $(1 - p/2)$  probability of the ex-post realization falling below the lower interval and above the upper interval, respectively. In other words, the forecast promises that the ex-post outcome,  $y_{t+1}$ , will fall inside the ex-ante forecasted interval with conditional probability,  $p$ . This setting naturally suggests the definition of a zero–one indicator sequence taking the value one if the realization falls inside the predicted interval and zero otherwise. We denote this indicator by

$$I_{t+1} \equiv \mathbf{I}(\hat{y}_{t+1|t}^L < y_{t+1} < \hat{y}_{t+1|t}^U). \tag{7.14}$$

Thus, for a correctly specified conditional interval forecast the conditional probability satisfies

$$P(I_{t+1} | \mathcal{F}_t) = p,$$

which also equals the conditional expectation of the zero–one indicator sequence,

$$E(I_{t+1} | \mathcal{F}_t) = p \cdot 1 + (1 - p) \cdot 0 = p. \quad (7.15)$$

A general regression version of this conditional expectation is readily expressed as

$$I_{t+1} - p = a + b'x_t + \varepsilon_{t+1}, \quad (7.16)$$

where the joint hypothesis that  $a = b = 0$  would be indicative of a correctly conditionally calibrated interval forecast series.

Since the construction of the interval forecast depends crucially on the forecasts for the underlying volatility, the set of information variables,  $x_t$ , could naturally include one or more volatility forecasts. The past value of the indicator sequence itself could also be included in the regression as an even easier and potentially effective choice of information variable. If the interval forecast ignores important volatility dynamics then the ex-post observations falling outside the ex-ante interval will cluster corresponding to periods of high volatility. In turn, this will induce serial dependence in the indicator sequence leading to a significantly positive  $b$  coefficient for  $x_t = (I_t - p)$ .

As noted in Section 2, the popular Value-at-Risk forecast corresponds directly to a one-sided interval forecast, and the regression in (7.16) can similarly be used to evaluate, or backtest, VARs. The indicator sequence in this case would simply be

$$I_{t+1} = \mathbf{I}(y_{t+1} < \text{VaR}_{t+1|t}^p), \quad (7.17)$$

where  $y_{t+1}$  now refers to the ex-post portfolio return. Capturing clustering in the indicator series (and thus clustered VaR violations) is particularly important within the context of financial risk management. The occurrence of, say, three VaR violations in one week is more likely to cause financial distress than three violations scattered randomly throughout the year. Recognizing that clusters in VaR violations likely are induced by neglected volatility dynamics again highlights the importance of volatility modeling and forecasting in financial risk management.

#### 7.4. Probability forecast evaluation and market timing tests

The interval and VaR forecasts discussed above correspond to quantiles (or thresholds) in the conditional distribution for a fixed and pre-specified probability of interest,  $p$ . In Section 2 we also considered probability forecasting in which the threshold of interest is pre-specified, with the probability of the random variable exceeding the threshold being forecasted. In this case the loss function is given by

$$L(y_{t+1}, \hat{y}_{t+1|t}) = (\mathbf{I}(y_{t+1} > c) - \hat{y}_{t+1|t})^2, \quad (7.18)$$

where  $c$  denotes the threshold, and the optimal forecast equals

$$\hat{y}_{t+1|t} = P(y_{t+1} > c \mid \mathcal{F}_t).$$

The generalized forecast error follows directly from (7.3),  $-2(\mathbf{I}(y_{t+1} > c) - \hat{y}_{t+1|t})$ , resulting in the corresponding forecast evaluation regression

$$\mathbf{I}(y_{t+1} > c) - \hat{y}_{t+1|t} = a + b'x_t + \varepsilon_{t+1}, \quad (7.19)$$

where the hypothesis of probability forecast unbiasedness corresponds to  $a = 0$  and  $b = 0$ . Again, the volatility forecast as well as the probability forecast itself would both be natural candidates for the vector of information variables. Notice also the similarity between the probability forecast evaluation regression in (7.19) and the interval forecast and VaR evaluation regression in (7.16).

The probability forecast evaluation framework above is closely related to tests for market timing in empirical finance. In market timing tests,  $y_{t+1}$  is the excess return on a risky asset and interest centers on forecasting the probability of a positive excess return, thus  $c = 0$ . In this regard, money managers are often interested in the correspondence between ex-ante probability forecasts which are larger than 0.5 and the occurrence of a positive excess return ex-post. In particular, suppose that a probability forecast larger than 0.5 triggers a long position in the risky asset and vice versa. The regression

$$\mathbf{I}(y_{t+1} > 0) = a + b\mathbf{I}(\hat{y}_{t+1|t} > 0.5) + \varepsilon_{t+1} \quad (7.20)$$

then provides a simple framework for evaluating the market timing ability in the forecasting model underlying the probability forecast,  $\hat{y}_{t+1|t}$ . Based on this regression it is also possible to show that  $b = p_+ + p_- - 1$ , where  $p_+$  and  $p_-$  denote the probabilities of a correctly forecasted positive and negative return, respectively. A significantly positive  $b$  thus implies that either  $p_+$  or  $p_-$  or both are significantly larger than 0.5.

### 7.5. Density forecast evaluation

The forecasts considered so far all predict certain aspects of the conditional distribution without necessarily fully specifying the distribution over the entire support. For many purposes, however, the entire predictive density is of interest, and tools for evaluating density forecasts are therefore needed. In Section 2 we explicitly defined the conditional density forecast as

$$\hat{y}_{t+1|t} = f_{t+1|t}(y) \equiv f(y_{t+1} = y \mid \mathcal{F}_t).$$

The Probability Integral Transform (PIT), defined as the probability of obtaining a value below the actual ex-post realization according to the ex-ante density forecast,

$$u_{t+1} \equiv \int_{-\infty}^{y_{t+1}} f_{t+1|t}(s) ds, \quad (7.21)$$

provides a general framework for evaluating the predictive distribution. As the PIT variable is a probability, its support is necessarily between zero and one. Furthermore, if the

density forecast is correctly specified,  $u_{t+1}$  must be i.i.d. uniformly distributed,

$$u_{t+1} \sim \text{i.i.d. } U(0, 1). \quad (7.22)$$

Intuitively, if the density forecast on average puts too little weight, say, in the left extreme of the support then a simple histogram of the PIT variable would not be flat but rather have too many observations close to zero. Thus, the PIT variable should be uniformly distributed. Furthermore, one should not be able to forecast at time  $t$  where in the forecasted density the realization will fall at time  $t + 1$ . If one could, then that part of the density forecast is assigned too little weight at time  $t$ . Thus, the PIT variable should also be independent over time.

These considerations show that it is not sufficient to test whether the PIT variable is uniformly distributed on average. We also need conditional tests to properly assess whether the  $u_{t+1}$ 's are i.i.d. Testing for an i.i.d. uniform distribution is somewhat cumbersome due to the bounded support. Alternatively, one may more conveniently test for normality of the transformed PIT variable,

$$\tilde{u}_{t+1} \equiv \Phi^{-1}(u_{t+1}) \sim \text{i.i.d. } N(0, 1), \quad (7.23)$$

where  $\Phi^{-1}(u)$  denotes the inverse cumulative density function of a standard normal variable.

In particular, the i.i.d. normal property in (7.23) implies that the conditional moment of any order  $j$  should equal the corresponding unconditional (constant) moment in the standard normal distribution, say  $\mu_j$ . That is,

$$E[\tilde{u}_{t+1}^j \mid \mathcal{F}_t] - \mu_j = 0. \quad (7.24)$$

This in turn suggests a simple density forecast evaluation system of regressions

$$\tilde{u}_{t+1}^j - \mu_j = a_j + b_j' x_{j,t} + \varepsilon_{j,t+1}, \quad (7.25)$$

where  $j$  determines the order of the moment in question. For instance, testing the hypothesis that  $a_j = b_j = 0$  for  $j = 1, 2, 3, 4$  will assess if the first four conditional (noncentral) moments are constant and equal to their standard normal values.

Consider now the case where the density forecast specification underlying the forecast supposedly is known,

$$y_{t+1} = \mu_{t+1|t} + \sigma_{t+1|t} z_{t+1}, \quad z_{t+1} \sim \text{i.i.d. } F. \quad (7.26)$$

In this situation, it is possible to directly test the validity of the dynamic model specification for the innovations,

$$z_{t+1} = (y_{t+1} - \mu_{t+1|t})/\sigma_{t+1|t} \sim \text{i.i.d. } F. \quad (7.27)$$

The i.i.d. property is most directly and easily tested via the autocorrelations of various powers,  $j$ , of the standardized residuals, say  $\text{Corr}(z_t^j, z_{t-k}^j)$ .

In particular, under the null hypothesis that the autocorrelations are zero at all lags, the Ljung–Box statistics for up to  $K$ th order serial correlation,

$$LB^j(K) \equiv T(T+2) \sum_{k=1}^K \text{Corr}^2(z_t^j, z_{t-k}^j) / (T-k), \quad (7.28)$$

should be the realization of a chi-square distribution with  $K$  degrees of freedom. Of course, this  $K$  degree of freedom test ignores the fact that the parameters in the density forecasting model typically will have to be estimated. As noted in Section 7.1, refined test statistics as well as simulation based techniques are available to formally deal with this issue.

As previously noted, in most financial applications involving daily or weekly returns, it is reasonable to assume that  $\mu_{t+1|t} \approx 0$ , so that

$$z_{t+1}^2 \approx y_{t+1}^2 / \sigma_{t+1|t}^2.$$

Thus, a dynamic variance model can readily be thought of as removing the dynamics from the squared observations. Misspecified variance dynamics are thus likely to show up as significant autocorrelations in  $z_{t+1}^2$ . This therefore suggests setting  $j = 2$  in (7.28) and calculating the Ljung–Box test based on the autocorrelations of the squared innovations,  $\text{Corr}(z_t^2, z_{t-k}^2)$ . This same Ljung–Box test procedure can, of course, also be used in testing for the absence of dynamic dependencies in the moments of the density forecast evaluation variable from (7.23),  $\tilde{u}_{t+1}$ .

### 7.6. Further reading

This section only scratches the surface on forecast evaluation. The properties and evaluation of point forecasts from general loss functions have recently been analyzed by Patton and Timmermann (2003, 2004). The statistical comparison of competing forecasts under general loss functions has been discussed by Diebold and Mariano (1995), Giacomini and White (2004), and West (1996). Forecast evaluation under mean-squared error loss is discussed in detail by West in Chapter 3 of this Handbook. Interval, quantile and Value-at-Risk forecast evaluation is developed further in Christoffersen (1998, 2003), Christoffersen, Hahn and Inoue (2001), Christoffersen and Pelletier (2004), Engle and Manganelli (2004), and Giacomini and Komunjer (2005). The evaluation of probability forecasts, sign forecasts and market timing techniques is surveyed in Breen, Glosten and Jagannathan (1989), Campbell, Lo and MacKinlay (1997, Chapter 2), and Christoffersen and Diebold (2003). Methods for density forecast evaluation are developed in Berkowitz (2001), Diebold, Gunther and Tay (1998), Giacomini (2002), and Hong (2000), as well as in Chapter 5 by Corradi and Swanson in this Handbook.

White (2000) provides a framework for assessing if the best forecasting model from a large set of potential models outperforms a given benchmark. Building on this idea, Hansen, Lunde and Nason (2003, 2005) develop a model confidence set approach for choosing the best volatility forecasting model.

1 Moreover, combining a number of volatility forecasts may be preferable to choosing 1  
2 a single best forecast. The general topic of forecast combination is discussed in detail in 2  
3 Chapter 4 by Timmermann in this Handbook. Volatility forecast combination has been 3  
4 found to work well in practice by Hu and Tsoukalas (1999). 4

5 Further discussion of volatility forecasting and forecast evaluation based on realized 5  
6 volatility measures can be found in Andersen and Bollerslev (1998a), Andersen, Bollerslev 6  
7 and Meddahi (2004, 2005), and Patton (2005). Andersen et al. (1999), Aït-Sahalia, 7  
8 Mykland and Zhang (2005), Bandi and Russel (2003, 2004), Bollen and Inder (2002), 8  
9 Curci and Corsi (2004), Hansen and Lunde (2004b), Martens (2003), and Zhang, Aït- 9  
10 Sahalia and Mykland (2005) all analyze the important choice of sampling frequency 10  
11 and/or the use of various sub-sampling and other corrective procedures in the practical 11  
12 construction of unbiased (and efficient) realized volatility measures. Alizadeh, Brandt 12  
13 and Diebold (2002) discuss the relative merits of realized and range-based volatility. For 13  
14 early work on the properties of range-based estimates, see Feller (1951) and Parkinson 14  
15 (1980). 15

16 Testing for normality of the transformed Probability Integral Transform (PIT) variable 16  
17 can be done in numerous ways. A couple of interesting recent procedures for testing 17  
18 dynamic models for correct distributional assumptions taking into account the parameter 18  
19 estimation error uncertainty are given by Bontemps and Meddahi (2005) and Duan 19  
20 (2003). 20

21 Several important topics were not explicitly discussed in this section. In the general 21  
22 forecasting area they include covariance and correlation forecast evaluation [see, e.g., 22  
23 Brandt and Diebold (2005)], as well as related multivariate density forecast evaluation 23  
24 [see, e.g., Diebold, Hahn and Tay (1999)]. In the area of financial forecast applica- 24  
25 tions, we did not discuss the evaluation of time-varying betas [see, e.g., Ghysels (1998)], 25  
26 volatility-based asset allocation [see, e.g., Fleming, Kirby and Ostdiek (2001, 2003)], 26  
27 and option valuation models [see, e.g., Bates (2003) and Christoffersen and Jacobs 27  
28 (2004a, 2004b)], to mention some. Nonetheless, the general volatility forecast evalua- 28  
29 tion framework set out above is flexible enough so that it may easily be adapted to 29  
30 each of these more specific situations. 30  
31 31  
32 32

## 33 8. Concluding remarks 33

34 34  
35 This chapter has focused on rigorous yet practical methods for volatility modeling and 35  
36 forecasting. The literature has obviously advanced rapidly and will almost surely con- 36  
37 tinue to thrive for the foreseeable future, as key challenges remain at least partially 37  
38 open. Some of these, such as large-dimensional covariance matrix modeling and prac- 38  
39 tical ways in which to best make use of the newly available ultra-high-frequency data 39  
40 have been touched upon. 40

41 Less obviously, and beyond the narrower realm of mathematical volatility models, 41  
42 the financial–econometric volatility literature has impacted the financial landscape in 42  
43 additional and important ways. Most notably, the newly-entrenched awareness of large 43



1 time variation and high persistence in asset return volatility has led to the emergence of 1  
2 volatility as an asset class, with a variety of vehicles now available for taking positions 2  
3 exclusively in volatility. This contrasts with traditional options-based instruments, the 3  
4 value of which varies, for example, with the price of the underlying in addition to its 4  
5 volatility. The new vehicles include both exchange-traded products such as the Chicago 5  
6 Board Options Exchange's VIX volatility index, which depends directly on the one- 6  
7 month options implied volatility for the S&P500 aggregate market index, as well as 7  
8 more specialized over-the-counter volatility and covariance swaps, which are essentially 8  
9 futures contracts written on various realized volatility measures. 9

10 In addition to the obvious and traditional uses of such products, such as hedging 10  
11 volatility exposure associated with running an options book, important new uses in 11  
12 asset allocation environments are emerging, as portfolio managers add volatility to 12  
13 otherwise-standard portfolios. While large positions in volatility may be undesirable, 13  
14 because volatility reverts to a fixed mean and hence has zero expected return in the 14  
15 long-run, small positions can provide a valuable hedge against crisis episodes in which 15  
16 simultaneously plunging prices cause both correlations and volatilities to increase. This 16  
17 type of hedge, of course, can be very appealing in both private and central bank asset- 17  
18 management environments. 18

19 Although it would be an exaggeration to claim that the mathematical volatility fore- 19  
20 casting models reviewed here are solely responsible for the emergence and deepening 20  
21 of financial volatility markets, the development of the models nevertheless provided 21  
22 (and continue to provide) a major push, the effects of which we predict will continue to 22  
23 evolve and resonate with financial market practice for many years to come. 23  
24 24  
25 25

## 26 **Uncited references** 26

27  
28 [[Ruge-Murcia \(2003\)](#)] [[Schwert \(1990\)](#)] 28  
29 29  
30 30

## 31 **References** 31

- 32 32  
33 [Aguilar, O., West, M.](#) (2000). "Bayesian dynamic factor models and variance matrix discounting for portfolio 33  
34 allocation". *Journal of Business and Economic Statistics* 18, 338–357. 34  
35 [Ait-Sahalia, Y., Brandt, M.](#) (2001). "Variable selection for portfolio choice". *Journal of Finance* 56, 1297– 35  
36 1350. 36  
37 [Ait-Sahalia, Y., Mykland, P.A., Zhang, L.](#) (2005). "How often to sample a continuous-time process in the 37  
38 presence of market microstructure noise". *Review of Financial Studies*. In press. 38  
39 [Alexander, C.](#) (2001). *Market Models: A Guide to Financial Data Analysis*. Wiley, Chichester. 39  
40 [Alizadeh, S., Brandt, M.W., Diebold, F.X.](#) (2002). "Range-based estimation of stochastic volatility models". 40  
41 *Journal of Finance* 57, 1047–1092. 41  
42 [Andersen, T.G.](#) (1992). "Volatility". Working Paper, Finance Department, Kellogg School, Northwestern Uni- 42  
43 versity. 43  
44 [Andersen, T.G.](#) (1994). "Stochastic autoregressive volatility: A framework for volatility modeling". *Mathe-* 44  
45 *matical Finance* 4, 75–102. 45

- 1 Andersen, T.G. (1996). "Return volatility and trading volume: An information flow interpretation of stochastic volatility". *Journal of Finance* 51, 169–204. 1
- 2 2
- 3 Andersen, T.G., Benzoni, L., Lund, J. (2002). "An empirical investigation of continuous-time equity return models". *Journal of Finance* 57, 1239–1284. 3
- 4 Andersen, T.G., Bollerslev, T. (1997). "Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns". *Journal of Finance* 52, 975–1005. 4
- 5 Andersen, T.G., Bollerslev, T. (1998a). "Answering the skeptics: Yes, standard volatility models do provide accurate forecasts". *International Economic Review* 39, 885–905. 5
- 6 Andersen, T.G., Bollerslev, T. (1998b). "ARCH and GARCH models". In: Kotz, S., Read, C.B., Banks, D.L. (Eds.), *Encyclopedia of Statistical Sciences*, vol. II. Wiley, New York. 6
- 7 Andersen, T.G., Bollerslev, T. (1998c). "DM-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies". *Journal of Finance* 53, 219–265. 7
- 8 Andersen, T.G., Bollerslev, T., Diebold, F.X. (2003a). "Parametric and nonparametric measurements of volatility". In: Ait-Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*. North-Holland, Amsterdam. In press. 8
- 9 Andersen, T.G., Bollerslev, T., Diebold, F.X. (2003b). "Some like it smooth and some like it rough: Untangling continuous and jump components in measuring, modeling and forecasting asset return volatility". Working Paper, Northwestern University, Duke University and University of Pennsylvania. 9
- 10 Andersen, T.G., Bollerslev, T., Meddahi, N. (2004). "Analytic evaluation of volatility forecasts". *International Economic Review* 45, 1079–1110. 10
- 11 Andersen, T.G., Bollerslev, T., Meddahi, N. (2005). "Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities". *Econometrica* 73, 279–296. 11
- 12 Andersen, T.G., Lund, J. (1997). "Estimating continuous-time stochastic volatility models of the short term interest rate diffusion". *Journal of Econometrics* 77, 343–377. 12
- 13 Andersen, T.G., Sørensen, B.E. (1996). "GMM estimation of a stochastic volatility model: A Monte Carlo study". *Journal of Business and Economic Statistics* 14, 328–352. 13
- 14 Andersen, T.G., Sørensen, B.E. (1997). "GMM and QML asymptotic standard deviations in stochastic volatility models: Comments on Ruiz (1994)". *Journal of Econometrics* 76, 397–403. 14
- 15 Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (1999). "(Understanding, optimizing, using and forecasting) realized volatility and correlation". Working Paper, Northwestern University, Duke University and University of Pennsylvania. 15
- 16 Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2000). "Great realizations". *Risk Magazine* 18, 105–108. 16
- 17 Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H. (2001a). "The distribution of stock return volatility". *Journal of Financial Economics* 61, 43–76. 17
- 18 Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2001b). "The distribution of exchange rate volatility". *Journal of the American Statistical Association* 96, 42–55. 18
- 19 Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2003). "Modeling and forecasting realized volatility". *Econometrica* 71, 579–625. 19
- 20 Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X. (2005). "Practical volatility and correlation modeling for financial market risk management". In: Carey, M., Stulz, R. (Eds.), *Risks of Financial Institutions*. University of Chicago Press for National Bureau of Economic Research, Chicago. 20
- 21 Andreou, E., Ghysels, E. (2002). "Detecting multiple breaks in financial market volatility dynamics". *Journal of Applied Econometrics* 17, 579–600. 21
- 22 Ang, A., Chen, J. (2002). "Asymmetric correlation of equity portfolios". *Journal of Financial Economics* 63, 443–494. 22
- 23 Bachelier, L. (1900). "Théorie de la spéculation". *Annales de l'Ecole Normale Supérieure* 3. Gauthier-Villars, Paris. English translation in: Cootner, P.H. (Ed.), *The Random Character of Stock Market Prices*. MIT Press, Cambridge. 23
- 24 Baillie, R.T., Bollerslev, T. (1992). "Prediction in dynamic models with time dependent conditional variances". *Journal of Econometrics* 52, 91–113. 24
- 25 25
- 26 26
- 27 27
- 28 28
- 29 29
- 30 30
- 31 31
- 32 32
- 33 33
- 34 34
- 35 35
- 36 36
- 37 37
- 38 38
- 39 39
- 40 40
- 41 41
- 42 42
- 43 43

- 1 **Baillie, R.T., Bollerslev, T., Mikkelsen, H.O.** (1996). "Fractionally integrated generalized autoregressive con- 1  
2 ditional heteroskedasticity". *Journal of Econometrics* 74, 3–30. 2
- 3 **Bakshi, G., Cao, C., Chen, Z.** (1997). "Empirical performance of alternative option pricing models". *Journal* 3  
4 *of Finance* 52, 2003–2049. 4
- 5 **Bandi, F., Russell, J.R.** (2003). "Volatility or microstructure noise". Working Paper, University of Chicago. 5  
6 **Bandi, F., Russell, J.R.** (2004). "Microstructure noise, realized variance, and optimal sampling". Working 6  
7 Paper, Graduate School of Business, University of Chicago. 7
- 8 **Banerjee, A., Urga, G.** (2005). "Modelling structural breaks, long memory and stock market volatility: An 8  
9 overview". *Journal of Econometrics*. In press. 9
- 10 **Barndorff-Nielsen, O.E., Shephard, N.** (2001). "Non-Gaussian Ornstein–Uhlenbeck-based models and some 10  
11 of their uses in financial economics". *Journal of the Royal Statistical Society, Series B* 63, 167–207. 11
- 12 **Barndorff-Nielsen, O.E., Shephard, N.** (2002). "Estimating quadratic variation using realised variance". *Jour- 12  
13 nal of Applied Econometrics* 17, 457–477. 13
- 14 **Barndorff-Nielsen, O.E., Shephard, N.** (2004a). "Power and bipower variation with stochastic volatility and 14  
15 jumps". *Journal of Financial Econometrics* 2, 1–37. 15
- 16 **Barndorff-Nielsen, O.E., Shephard, N.** (2004b). "Econometric analysis of realized covariation: High fre- 16  
17 quency based covariance, regression and correlation in financial economics". *Econometrica* 72, 885–925. 17
- 18 **Barone-Adesi, G., Giannopoulos, K., Vosper, L.** (1999). "VaR without correlations for non-linear portfolios". 18  
19 *Journal of Futures Markets* 19, 583–602. 19
- 20 **Barrett, C.** (1999). "The effects of real exchange rate depreciation on stochastic producer prices in low-income 20  
21 agriculture". *Agricultural Economics* 20, 215–230. 21
- 22 **Barucci, E., Reno, R.** (2002). "On measuring volatility and the GARCH forecasting performance". *Journal of* 22  
23 *International Financial Markets, Institutions and Money* 12, 182–200. 23
- 24 **Bates, D.S.** (1996). "Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark 24  
25 options". *Review of Financial Studies* 9, 69–107. 25
- 26 **Bates, D.S.** (2003). "Empirical option pricing: A retrospection". *Journal of Econometrics* 116, 387–404. 26
- 27 **Battle, C., Barquin, J.** (2004). "Fuel prices scenario generation based on a multivariate GARCH model for 27  
28 risk analysis in a wholesale electricity market". *International Journal of Electrical Power and Energy* 28  
29 *Systems* 26, 273–280. 29
- 30 **Bauwens, L., Laurent, S.** (2005). "A new class of multivariate skew densities with application to GARCH 30  
31 models". *Journal of Business and Economic Statistics*. In press. 31
- 32 **Bauwens, L., Laurent, S., Rombouts, J.V.K.** (2005). "Multivariate GARCH models: A survey". *Journal of* 32  
33 *Applied Econometrics*. In press. 33
- 34 **Bauwens, L., Lubrano, M.** (1998). "Bayesian inference on GARCH models using the Gibbs sampler". *The* 34  
35 *Econometrics Journal* 1, 23–46. 35
- 36 **Bauwens, L., Lubrano, M.** (1999). *Bayesian Dynamic Econometrics*. Oxford University Press, Oxford. 36
- 37 **Bekaert, G., Wu, G.** (2000). "Asymmetric volatility and risk in equity markets". *Review of Financial Stud-* 37  
38 *ies* 13, 1–42. 38
- 39 **Bera, A.K., Kim, S.** (2002). "Testing constancy of correlation and other specifications of the BGARCH model 39  
40 with an application to international equity returns". *Journal of Empirical Finance* 7, 305–362. 40
- 41 **Beran, J.** (1994). *Statistics for Long-Memory Processes*. Chapman and Hall, New York. 41
- 42 **Berkowitz, J.** (2001). "Testing density forecasts with applications to risk management". *Journal of Business* 42  
43 *and Economic Statistics* 19, 465–474. 43
- 44 **Berkowitz, J., O'Brien, J.** (2002). "How accurate are the value-at-risk models at commercial banks?" *Journal* 44  
45 *of Finance* 57, 1093–1112. 45
- 46 **Billio, M., Caporin, M., Gobbo, M.** (2003). "Block dynamic conditional correlation multivariate GARCH 46  
47 models". Working Paper, Università di Venezia. 47
- 48 **Black, F.** (1976). "Studies of stock market volatility changes". *Proceedings of the American Statistical Asso-* 48  
49 *ciation, Business and Economic Statistics Section*, 177–181. 49
- 50 **Black, F., Scholes, M.** (1973). "The pricing of options and corporate liabilities". *Journal of Political Econ-* 50  
51 *omy* 81, 637–654. 51

- 1 **Bollen, B., Inder, B.** (2002). "Estimating daily volatility in financial markets utilizing intraday data". *Journal of Empirical Finance* 9, 551–562. 1
- 2 of Empirical Finance 9, 551–562. 2
- 3 **Bollerslev, T.** (1986). "Generalized autoregressive conditional heteroskedasticity". *Journal of Econometrics* 31, 307–327. 3
- 4 **Bollerslev, T.** (1987). "A conditionally heteroskedastic time series model for speculative prices and rates of return". *Review of Economics and Statistics* 69, 542–547. 4
- 5 **Bollerslev, T.** (1990). "Modeling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model". *Review of Economics and Statistics* 72, 498–505. 5
- 6 **Bollerslev, T., Chou, R.Y., Kroner, K.F.** (1992). "ARCH modeling in finance: A selective review of the theory and empirical evidence". *Journal of Econometrics* 52, 5–59. 6
- 7 **Bollerslev, T., Engle, R.F.** (1993). "Common persistence in conditional variances". *Econometrica* 61, 167–186. 7
- 8 **Bollerslev, T., Engle, R.F., Nelson, D.B.** (1994). "ARCH models". In: Engle, R.F., McFadden, D. (Eds.), *Handbook of Econometrics*, vol. IV. North-Holland, Amsterdam, pp. 2959–3038. 8
- 9 **Bollerslev, T., Engle, R.F., Wooldridge, J.M.** (1988). "A capital asset pricing model with time varying covariances". *Journal of Political Economy* 96, 116–131. 9
- 10 **Bollerslev, T., Jubinsky, P.D.** (1999). "Equity trading volume and volatility: Latent information arrivals and common long-run dependencies". *Journal of Business and Economic Statistics* 17, 9–21. 10
- 11 **Bollerslev, T., Mikkelsen, H.O.** (1996). "Modeling and pricing long memory in stock market volatility". *Journal of Econometrics* 73, 151–184. 11
- 12 **Bollerslev, T., Mikkelsen, H.O.** (1999). "Long-term equity anticipation securities and stock market volatility dynamics". *Journal of Econometrics* 92, 75–99. 12
- 13 **Bollerslev, T., Wooldridge, J.M.** (1992). "Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances". *Econometric Reviews* 11, 143–172. 13
- 14 **Bollerslev, T., Wright, J.H.** (2001). "Volatility forecasting, high-frequency data, and frequency domain inference". *Review of Economic and Statistics* 83, 596–602. 14
- 15 **Bollerslev, T., Zhou, H.** (2005). "Volatility puzzles: A simple framework for gauging return-volatility regressions". *Journal of Econometrics*. In press. 15
- 16 **Bontemps, C., Meddahi, N.** (2005). "Testing normality: A GMM approach". *Journal of Econometrics*. In press. 16
- 17 **Brandt, M.W.** (1999). "Estimating portfolio and consumption choice: A conditional Euler equations approach". *Journal of Finance* 54, 1609–1646. 17
- 18 **Brandt, M.W.** (2004). "Portfolio choice problems". In: Aït-Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*. In press. 18
- 19 **Brandt, M.W., Diebold, F.X.** (2005). "A no-arbitrage approach to range-based estimation of return covariances and correlations". *Journal of Business*. In press. 19
- 20 **Brandt, M.W., Jones, C.** (2002). "Volatility forecasting with range-based EGARCH models". Working Paper, The Wharton School, University of Pennsylvania. 20
- 21 **Braun, P.A., Nelson, D.B., Sunier, A.M.** (1995). "Good news, bad news, volatility, and betas". *Journal of Finance* 50, 1575–1603. 21
- 22 **Breen, W., Glosten, L.R., Jagannathan, R.** (1989). "Economic significance of predictable variations in stock index returns". *Journal of Finance* 44, 1177–1189. 22
- 23 **Breidt, F.J., Crato, N., de Lima, P.F.J.** (1998). "On the detection and estimation of long memory in stochastic volatility". *Journal of Econometrics* 83, 325–348. 23
- 24 **Brooks, C.** (1997). "GARCH modelling in finance: A review of the software options". *The Economic Journal* 107, 1271–1276. 24
- 25 **Brooks, C.** (2002). *Introductory Econometrics for Finance*. Cambridge University Press, Cambridge. 25
- 26 **Brooks, C., Burke, S.P., Persaud, G.** (2001). "Benchmarks and the accuracy of GARCH model estimation". *International Journal of Forecasting* 17, 45–56. 26
- 27 **Brooks, C., Burke, S.P., Persaud, G.** (2003). "Multivariate GARCH models: Software choice and estimation issues". *Journal of Applied Econometrics* 18, 725–734. 27

- 1 Buguk, C., Hudson, D., Hanson, T. (2003). "Price volatility spillover in agricultural markets: An examination of U.S. catfish markets". *Journal of Agricultural and Resource Economics* 28, 86–99.
- 2 Calvet, L., Fisher, A. (2002). "Multifractality in asset returns: Theory and evidence". *Review of Economics and Statistics* 84, 381–406.
- 3 Calvet, L., Fisher, A. (2004). "How to forecast long-run volatility: Regime switching and the estimation of multifractal processes". *Journal of Financial Econometrics* 2, 49–83.
- 4 Calvet, L.E., Fisher, A.J., Thompson, S.B. (2005). "Volatility comovement: A multifrequency approach". *Journal of Econometrics*. In press.
- 5 Campbell, J.Y. (1987). "Stock returns and the term structure". *Journal of Financial Economics* 18, 373–399.
- 6 Campbell, J.Y. (2003). "Consumption-based asset pricing". In: Constantinides, G.M., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance*, vol. 1B. North-Holland, Amsterdam, pp. 803–888.
- 7 Campbell, J.Y., Hentschel, L. (1992). "No news is good news: An asymmetric model of changing volatility in stock returns". *Journal of Financial Economics* 31, 281–318.
- 8 Campbell, J.Y., Lo, A.W., MacKinlay, A.C. (1997). *The Econometrics of Financial Markets*. Princeton University Press, Princeton.
- 9 Campbell, S., Diebold, F.X. (2005). "Weather forecasting for weather derivatives". *Journal of the American Statistical Association* 100, 6–16.
- 10 Cappiello, L., Engle, R.F., Sheppard, K. (2004). "Asymmetric dynamic in the correlations of global equity and bond returns". Working Paper, NYU Stern School of Business.
- 11 Chacko, G., Viceira, L.M. (2003). "Spectral GMM estimation of continuous-time processes". *Journal of Econometrics* 116, 259–292.
- 12 Chan, N.H. (2002). *Time Series: Applications to Finance*. Wiley, New York.
- 13 Chernov, M., Ghysels, E. (2000). "A study towards a unified framework for the joint estimation of objective and risk neutral measures for the purpose of options valuation". *Journal of Financial Economics* 56, 407–458.
- 14 Chernov, M., Gallant, A.R., Ghysels, E., Tauchen, G.E. (2003). "Alternative models for stock price dynamics". *Journal of Econometrics* 116, 225–257.
- 15 Christie, A.A. (1982). "The stochastic behavior of common stock variances: Value, leverage and interest rate effects". *Journal of Financial Economics* 10, 407–432.
- 16 Christoffersen, P.F. (1998). "Evaluating interval forecasts". *International Economic Review* 39, 841–862.
- 17 Christoffersen, P.F. (2003). *Elements of Financial Risk Management*. Academic Press, San Diego.
- 18 Christoffersen, P.F., Diebold, F.X. (1996). "Further results on forecasting and model selection under asymmetric loss". *Journal of Applied Econometrics* 11, 561–572.
- 19 Christoffersen, P.F., Diebold, F.X. (1997). "Optimal prediction under asymmetric loss". *Econometric Theory* 13, 808–817.
- 20 Christoffersen, P.F., Diebold, F.X. (2000). "How relevant is volatility forecasting for financial risk management?" *Review of Economics and Statistics* 82, 12–22.
- 21 Christoffersen, P.F., Diebold, F.X. (2003). "Financial asset returns, direction-of-change forecasting, and volatility dynamics". NBER Working Paper 10009.
- 22 Christoffersen, P.F., Hahn, J., Inoue, A. (2001). "Testing and comparing value-at-risk measures". *Journal of Empirical Finance* 8, 325–342.
- 23 Christoffersen, P.F., Jacobs, K. (2004a). "The importance of the loss function in option valuation". *Journal of Financial Economics* 72, 291–318.
- 24 Christoffersen, P.F., Jacobs, K. (2004b). "Which GARCH model for option valuation?" *Management Science* 50, 1204–1221.
- 25 Christoffersen, P.F., Pelletier, D. (2004). "Backtesting value-at-risk: A duration-based approach". *Journal of Financial Econometrics* 2, 84–108.
- 26 Clark, P.K. (1973). "A subordinated stochastic process model with finite variance for speculative prices". *Econometrica* 41, 135–156.
- 27 Cochrane, J. (2001). *Asset Pricing*. Princeton University Press, Princeton.
- 28 Comte, F., Renault, E. (1998). "Long memory in continuous-time stochastic volatility models". *Mathematical Finance* 8, 291–323.

- 1 Corsi, F. (2003). "A simple long memory model of realized volatility". Working Paper, University of Southern  
2 Switzerland.
- 3 Corsi, F., Zumbach, G., Müller, U.A., Dacorogna, M. (2001). "Consistent high-precision volatility from high-  
4 frequency data". *Economic Notes* 30, 183–204.
- 5 Curci, G., Corsi, F. (2004). "A discrete sine transform approach for realized volatility measurement". Working  
6 Paper, University of Pisa and University of Southern Switzerland.
- 7 Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V. (2001). *An Introduction to High-  
8 Frequency Finance*. Academic Press, San Diego.
- 9 Danielsson, J. (1994). "Stochastic volatility in asset prices: Estimation by simulated maximum likelihood".  
10 *Journal of Econometrics* 54, 375–400.
- 11 Danielsson, J., Richard, J.F. (1993). "Accelerated Gaussian importance sampler with application to dynamic  
12 latent variable models". *Journal of Applied Econometrics* 8, S153–S173.
- 13 Davidson, J. (2004). "Moment and memory properties of linear conditional heteroskedasticity models, and a  
14 new model". *Journal of Business and Economic Statistics* 22, 16–29.
- 15 De Jong, F., Nijman, T. (1997). "High frequency analysis of lead-lag relationships between financial markets".  
16 *Journal of Empirical Finance* 4, 259–277.
- 17 Deo, R., Hurvich, C. (2001). "On the log periodogram regression estimator of the memory parameter in long  
18 memory stochastic volatility models". *Econometric Theory* 17, 686–710.
- 19 DeSantis, G., Gerard, B. (1997). "International asset pricing and portfolio diversification with time-varying  
20 risk". *Journal of Finance* 52, 1881–1912.
- 21 DeSantis, G., Litterman, R., Vesval, A., Winkelmann, K. (2003). "Covariance matrix estimation". In: Litter-  
22 man, R. (Ed.), *Modern Investment Management: An Equilibrium Approach*. Wiley, London.
- 23 Dhaene, G., Vergote, O. (2004). "Asymptotic results for GMM estimators of stochastic volatility models".  
24 Working Paper, Department of Economics, K.U. Leuven.
- 25 Diebold, F.X. (1988). *Empirical Modeling of Exchange Rate Dynamics*. Springer-Verlag, New York.
- 26 Diebold, F.X. (2004). "The Nobel memorial prize for Robert F. Engle". *Scandinavian Journal of Eco-  
27 nomics* 106, 165–185.
- 28 Diebold, F.X., Gunther, T., Tay, A. (1998). "Evaluating density forecasts, with applications to financial risk  
29 management". *International Economic Review* 39, 863–883.
- 30 Diebold, F.X., Hahn, J., Tay, A. (1999). "Multivariate density forecast evaluation and calibration in financial  
31 risk management: High-frequency returns on foreign exchange". *Review of Economics and Statistics* 81,  
32 661–673.
- 33 Diebold, F.X., Inoue, A. (2001). "Long memory and regime switching". *Journal of Econometrics* 105, 131–  
34 159.
- 35 Diebold, F.X., Lopez, J. (1995). "Modeling volatility dynamics". In: Hoover, K. (Ed.), *Macroeconometrics:  
36 Developments, Tensions and Prospects*. Kluwer Academic, Boston, pp. 427–472.
- 37 Diebold, F.X., Lopez, J. (1996). "Forecast evaluation and combination". In: Maddala, G.S., Rao, C.R. (Eds.),  
38 *Handbook of Statistics*. North-Holland, Amsterdam, pp. 241–268.
- 39 Diebold, F.X., Mariano, R.S. (1995). "Comparing predictive accuracy". *Journal of Business and Economic  
40 Statistics* 13, 253–265.
- 41 Diebold, F.X., Nerlove, M. (1989). "The dynamics of exchange rate volatility: A multivariate latent factor  
42 ARCH model". *Journal of Applied Econometrics* 4, 1–21.
- 43 Ding, Z., Granger, C.W.J. (1996). "Modeling volatility persistence of speculative returns: A new approach".  
*Journal of Econometrics* 73, 185–215.
- Ding, Z., Granger, C.W.J., Engle, R.F. (1993). "A long memory property of stock market returns and a new  
model". *Journal of Empirical Finance* 1, 83–106.
- Drost, F.C., Nijman, T.E. (1993). "Temporal aggregation of GARCH processes". *Econometrica* 61, 909–927.
- Drost, F.C., Werker, B.J.M. (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.

- 1 Duan, J.-C. (2003). "A specification test for time series models by a normality transformation". Working Paper, Rotman School of Management, University of Toronto. 1
- 2 Duffie, D., Pan, J., Singleton, K.J. (2000). "Transform analysis and asset pricing for affine jump-diffusions". 2
- 3 *Econometrica* 68, 1343–1376. 3
- 4 Duffie, D., Singleton, K.J. (1993). "Simulated moments estimation of Markov models of asset prices". *Econo-* 4
- 5 *metrica* 61, 929–952. 5
- 6 Dufour, J.M. (2004). "Monte Carlo tests with nuisance parameters: A general approach to finite-sample in- 6
- 7 ference and nonstandard asymptotics in econometrics". Working Paper, Université de Montréal. 7
- 8 Elerian, O., Chib, S., Shephard, N. (2001). "Likelihood inference for discretely observed nonlinear diffu- 8
- 9 sions". *Econometrica* 69, 959–994. 9
- 10 Embrechts, P., Klüppelberg, C., Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. 10
- 11 Springer-Verlag, Berlin. 11
- 12 Enders, W. (2004). *Applied Econometric Time Series*. Wiley, Hoboken. 12
- 13 Engle, R.F. (1982). "Autoregressive conditional heteroskedasticity with estimates of the variance of UK infla- 13
- 14 tion". *Econometrica* 50, 987–1008. 14
- 15 Engle, R.F. (1995). *ARCH: Selected Readings*. Oxford University Press, Oxford. 15
- 16 Engle, R.F. (2001). "GARCH 101: The use of ARCH/GARCH models in applied econometrics". *Journal of* 16
- 17 *Economic Perspectives* 15, 157–168. 17
- 18 Engle, R.F. (2002). "Dynamic conditional correlation: A simple class of multivariate GARCH models". *Jour-* 18
- 19 *nal of Business and Economic Statistics* 20, 339–350. 19
- 20 Engle, R.F. (2004). "Nobel lecture. Risk and volatility: Econometric models and financial practice". *American* 20
- 21 *Economic Review* 94, 405–420. 21
- 22 Engle, R.F., Bollerslev, T. (1986). "Modeling the persistence of conditional variances". *Econometric Re-* 22
- 23 *views* 5, 1–50. 23
- 24 Engle, R.F., Gonzalez-Rivera, G. (1991). "Semiparametric ARCH models". *Journal of Business and Eco-* 24
- 25 *nomics Statistics* 9, 345–360. 25
- 26 Engle, R.F., Ito, T., Lin, W.L. (1990). "Meteor showers or heat waves? Heteroskedastic intra-daily volatility 26
- 27 in the foreign exchange market". *Econometrica* 58, 525–542. 27
- 28 Engle, R.F., Kroner, F.K. (1995). "Multivariate simultaneous generalized ARCH". *Econometric Theory* 11, 28
- 29 122–150. 29
- 30 Engle, R.F., Lee, G.G.J. (1999). "A permanent and transitory component model of stock return volatility". In: 30
- 31 Engle, R.F., White, H. (Eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive* 31
- 32 *W.J. Granger*. Oxford University Press, Oxford, pp. 475–497. 32
- 33 Engle, R.F., Lilien, D.M., Robbins, R.P. (1987). "Estimating time varying risk premia in the term structure: 33
- 34 The ARCH-M model". *Econometrica* 55, 391–407. 34
- 35 Engle, R.F., Manganelli, S. (2004). "CAViaR: Conditional autoregressive value at risk by regression quan- 35
- 36 tiles". *Journal of Business and Economic Statistics* 22, 367–381. 36
- 37 Engle, R.F., Ng, V.K. (1993). "Measuring and testing the impact of news on volatility". *Journal of Finance* 48, 37
- 38 1749–1778. 38
- 39 Engle, R.F., Ng, V.K., Rothschild, M. (1990). "Asset pricing with a factor-ARCH covariance structure: Em- 39
- 40 pirical estimates for treasury bills". *Journal of Econometrics* 45, 213–238. 40
- 41 Engle, R.F., Patton, A.J. (2001). "What good is a volatility model?" *Quantitative Finance* 1, 237–245. 41
- 42 Engle, R.F., Sheppard, K. (2001). "Theoretical and empirical properties of dynamic conditional correlation 42
- 43 multivariate GARCH". Working Paper, NYU Stern School of Business. 43
- 44 Epps, T. (1979). "Comovements in stock prices in the very short run". *Journal of the American Statistical* 44
- 45 *Association* 74, 291–298. 45
- 46 Eraker, B. (2001). "MCMC analysis of diffusions with applications to finance". *Journal of Business and* 46
- 47 *Economic Statistics* 19, 177–191. 47
- 48 Eraker, B. (2004). "Do stock prices and volatility jump? Reconciling evidence from spot and option prices". 48
- 49 *Journal of Finance* 59, 1367–1403. 49
- 50 Eraker, B., Johannes, M., Polson, N.G. (2003). "The impact of jumps in equity index volatility and returns". 50
- 51 *Journal of Finance* 58, 1269–1300. 51

- 1 Erb, C., Harvey, C., Viskanta, T. (1994). "Forecasting international equity correlations". *Financial Analysts* 1  
2 *Journal* 50, 32–45. 2
- 3 Ewing, B.T., Piette, M.J., Payne, J.E. (2003). "Forecasting medical net discount rates". *Journal of Risk and* 3  
4 *Insurance* 70, 85–95. 4
- 5 Feller, W. (1951). "The asymptotic distribution of the range of sums of random variables". *Annals of Mathe-* 5  
6 *matical Statistics* 22, 427–432. 6
- 7 Fiorentini, G., Sentana, E., Calzolari, G. (2003). "Maximum likelihood estimation and inference in multi-
- 8 variate conditionally heteroskedastic dynamic regression models with Student- $t$  innovations". *Journal of* 7  
9 *Business and Economic Statistics* 21, 532–546. 8
- 10 Fiorentini, G., Sentana, E., Shephard, N. (2004). "Likelihood-based estimation of latent generalized ARCH
- 11 structures". *Econometrica* 72, 1481–1517. 10
- 12 Fleming, J., Kirby, C. (2003). "A closer look at the relation between GARCH and stochastic autoregressive
- 13 volatility". *Journal of Financial Econometrics* 1, 365–419. 11
- 14 Fleming, J., Kirby, C., Ostdiek, B. (2001). "The economic value of volatility timing". *Journal of Finance* 56, 12  
15 329–352. 13
- 16 Fleming, J., Kirby, C., Ostdiek, B. (2003). "The economic value of volatility timing using "realized" volatil-
- 17 ity". *Journal of Financial Economics* 67, 473–509. 14
- 18 Foster, D.P., Nelson, D.B. (1996). "Continuous record asymptotics for rolling sample variance estimators". 15  
16 *Econometrica* 64, 139–174. 16
- 17 Fouque, J.-P., Papanicolaou, G., Sircar, K.R. (2000). *Derivatives in Financial Markets with Stochastic Volat-* 17  
18 *ility*. Princeton University Press, Princeton. 18
- 19 Franses, P.H., Hafner, C. (2003). "A generalized dynamic conditional correlation model for many asset re-
- 20 turns". Working Paper, Erasmus University, Rotterdam. 19
- 21 Franses, P.H., van Dijk, D. (2000). *Non-Linear Time Series Models in Empirical Finance*. Cambridge Uni-
- 22 versity Press, Cambridge. 20
- 23 French, K.R., Schwert, G.W., Stambaugh, R.F. (1987). "Expected stock returns and volatility". *Journal of* 22  
24 *Financial Economics* 19, 3–29. 23
- 25 Fridman, M., Harris, L. (1998). "A maximum likelihood approach for non-Gaussian stochastic volatility mod-
- 26 els". *Journal of Business and Economic Statistics* 16, 284–291. 24
- 27 Gallant, A.R., Hsieh, D.A., Tauchen, G.E. (1997). "Estimation of stochastic volatility models with diagnos-
- 28 tics". *Journal of Econometrics* 81, 159–192. 25
- 29 Gallant, A.R., Hsu, C.T., Tauchen, G.E. (1999). "Using daily range data to calibrate volatility diffusions and
- 30 extract the forward integrated variance". *Review of Economics and Statistics* 81, 617–631. 27
- 31 Gallant, A.R., Rossi, P.E., Tauchen, G.E. (1992). "Stock prices and volume". *Review of Financial Studies* 5, 28  
32 199–242. 29
- 33 Gallant, A.R., Tauchen, G.E. (1996). "Which moments to match". *Econometric Theory* 12, 657–681. 30
- 34 Gallant, A.R., Tauchen, G.E. (1997). "Estimation of continuous-time models for stock returns and interest
- 35 rates". *Macroeconomic Dynamics* 1, 135–168. 31
- 36 Gallant, A.R., Tauchen, G. (1998). "Reprojecting partially observed systems with application to interest rate
- 37 diffusions". *Journal of the American Statistical Association* 93, 10–24. 32
- 38 Geweke, J. (1989a). "Bayesian inference in econometric models using Monte Carlo integration". *Economet-* 33  
39 *rica* 57, 1317–1340. 34
- 40 Geweke, J. (1989b). "Exact predictive densities in linear models with ARCH disturbances". *Journal of Econo-* 35  
41 *metrics* 40, 63–86. 36
- 42 Ghysels, E. (1998). "On stable factor structures in the pricing of risk: Do time-varying betas help or hurt?" 37  
43 *Journal of Finance* 53, 549–573. 38
- 44 Ghysels, E., Harvey, A.C., Renault, E. (1996). "Stochastic volatility". In: Maddala, G.S., Rao, C.R. (Eds.), 39  
45 *Handbook of Statistics*, vol. 14. North-Holland, Amsterdam. 40
- 46 Ghysels, E., Santa-Clara, P., Valkanov, R. (2004). "Predicting volatility: Getting the most out of data sampled 41  
42 at different frequencies". Working Paper, University of North Carolina and University of California, Los 42  
43 Angeles. 43



- 1 **Giacomini, R.** (2002). "Comparing density forecasts via weighted likelihood ratio tests: Asymptotic and boot- 1  
strap methods". Boston College, Department of Economics Working Paper 583. 2
- 3 **Giacomini, R., Komunjer, I.** (2005). "Evaluation and combination of conditional quantile forecasts". *Journal* 3  
*of Business and Economic Statistics*. In press. 4
- 5 **Giacomini, R., White, H.** (2004). "Tests of conditional predictive ability". Working Paper, University of Cal- 5  
ifornia, San Diego. 6
- 7 **Giordani, P., Soderlind, P.** (2003). "Inflation forecast uncertainty". *European Economic Review* 47, 1037– 6  
1059. 7
- 8 **Giraitis, L., Kokoszka, P., Leipus, R.** (2000). "Stationary ARCH models: Dependence structure and central 8  
limit theorem". *Econometric Theory* 16, 3–22. 9
- 10 **Glosten, L.R., Jagannathan, R., Runkle, D.** (1993). "On the relation between the expected value and the 10  
volatility of the nominal excess return on stocks". *Journal of Finance* 48, 1779–1801. 11
- 12 **Gouriéroux, C., Jasiak, J.** (2001). *Financial Econometrics*. Princeton University Press, Princeton. 11
- 13 **Gouriéroux, C., Monfort, A., Renault, E.** (1993). "Indirect inference". *Journal of Applied Econometrics* 8, 12  
S85–S118. 13
- 14 **Granger, C.W.J.** (1969). "Prediction with a generalized cost of error function". *Operational Research Quar-* 14  
*terly* 20, 199–207. 15
- 16 **Granger, C.W.J.** (1980). "Long memory relationships and the aggregation of dynamic models". *Journal of* 15  
*Econometrics* 14, 227–238. 16
- 17 **Granger, C.W.J., White, H., Kamstra, M.** (1989). "Interval forecasting: An analysis based on ARCH-quantile 17  
estimators". *Journal of Econometrics* 40, 87–96. 18
- 18 **Gronke, P., Brehm, J.** (2002). "History, heterogeneity, and presidential approval: A modified ARCH ap- 18  
proach". *Electoral Studies* 21, 425–452. 19
- 20 **Guidolin, M., Timmermann, A.** (2005a). "Term structure of risk under alternative specifications". *Journal of* 20  
*Econometrics*. In press. 21
- 22 **Guidolin, M., Timmermann, A.** (2005b). "An econometric model of nonlinear dynamics in the joint distribu- 22  
tion of stock and bond returns". *Journal of Applied Econometrics*. In press. 23
- 24 **Hamilton, J.D.** (1994). *Time Series Analysis*. Princeton University Press, Princeton. 23
- 25 **Hamilton, J.D., Lin, G.** (1996). "Stock market volatility and the business cycle". *Journal of Applied Econo-* 24  
*metrics* 11, 573–593. 25
- 26 **Hansen, P.R., Lunde, A.** (2004a). "A realized variance for the whole day based on intermittent high-frequency 26  
data". Working Paper, Department of Economics, Stanford University. 27
- 28 **Hansen, P.R., Lunde, A.** (2004b). "An unbiased measure of realized variance". Working Paper, Department of 27  
Economics, Stanford University. 28
- 29 **Hansen, P.R., Lunde, A.** (2005). "A forecast comparison of volatility models: Does anything beat a 29  
GARCH(1, 1)?" *Journal of Applied Econometrics*. In press. 30
- 31 **Hansen, P.R., Lunde, A., Nason, J.M.** (2003). "Choosing the best volatility models: A model confidence set 31  
approach". *Oxford Bulletin of Economics and Statistics* 65, 839–861. 32
- 32 **Hansen, P.R., Lunde, A., Nason, J.M.** (2005). "Model confidence sets for forecasting models". Working Paper, 32  
Stanford University. 33
- 34 **Harris, L.** (1986). "Cross-security tests of the mixture of distributions hypothesis". *Journal of Financial and* 34  
*Quantitative Analysis* 21, 39–46. 35
- 36 **Harris, L.** (1987). "Transactions data tests of the mixture of distributions hypothesis". *Journal of Financial* 35  
*and Quantitative Analysis* 22, 127–141. 36
- 37 **Harvey, A.C.** (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge Univer- 37  
sity Press, Cambridge. 38
- 39 **Harvey, A.C.** (1998). "Long memory in stochastic volatility". In: Knight, J., Satchell, S. (Eds.), *Forecasting* 39  
*Volatility in Financial Markets*. Butterworth–Heinemann, Oxford, pp. 307–320. 40
- 41 **Harvey, A.C.** (2006). "Forecasting with unobserved components time series models". In: Elliott, G., Granger, 40  
C.W.J., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*. Elsevier, Amsterdam. This volume. 41
- 42 **Harvey, A.C., Ruiz, E., Sentana, E.** (1992). "Unobserved component time series models with ARCH distur- 42  
bances". *Journal of Econometrics* 52, 129–157. 43

- 1 Harvey, A.C., Ruiz, E., Shephard, E. (1994). "Multivariate stochastic variance models". *Review of Economic* 1  
2 *Studies* 61, 247–264. 2
- 3 Harvey, A.C., Shephard, E. (1996). "Estimation of an asymmetric model of asset prices". *Journal of Business* 3  
4 *and Economic Statistics* 14, 429–434. 4
- 5 Harvey, C.R. (2001). "The specification of conditional expectations". *Journal of Empirical Finance* 8, 573–  
6 637. 5
- 7 Hentschel, L. (1995). "All in the family: Nesting symmetric and asymmetric GARCH models". *Journal of* 6  
8 *Financial Economics* 39, 71–104. 7
- 9 Heston, S.L. (1993). "A closed form solution for options with stochastic volatility, with applications to bond  
10 and currency options". *Review of Financial Studies* 6, 327–343. 8
- 11 Heston, S.L., Nandi, S. (2000). "A closed-form GARCH option valuation model". *Review of Financial Stud-* 9  
12 *ies* 13, 585–625. 10
- 13 Hong, Y. (2000). "Evaluation of out-of-sample density forecasts with applications to stock prices". Working  
14 Paper, Department of Economics and Department of Statistical Science, Cornell University. 11
- 15 Hsieh, D.A. (1989). "Modeling heteroskedasticity in foreign exchange rates". *Journal of Business and Eco-* 12  
16 *nomic Statistics* 7, 307–317. 13
- 17 Hu, M., Tsoukalas, C. (1999). "Combining conditional volatility forecasts using neural networks: An appli-  
18 cation to the EMS exchange rates". *Journal of International Financial Markets, Institutions and Money* 9,  
19 407–422. 14
- 20 Huang, X., Tauchen, G.E. (2004). "The relative contribution of jumps to total price variance". Working Paper,  
21 Duke University. 15
- 22 Hull, J., White, A. (1987). "The pricing of options on assets with stochastic volatilities". *Journal of Finance* 42,  
23 281–300. 16
- 24 J.P. Morgan (1996). RiskMetrics, Technical Documents, fourth ed. New York. 17
- 25 Jacquier, E., Polson, N.G., Rossi, P.E. (1994). "Bayesian analysis of stochastic volatility models". *Journal of* 18  
26 *Business and Economic Statistics* 12, 371–389. 19
- 27 Jagannathan, R., Ma, T. (2003). "Risk reduction in large portfolios: Why imposing the wrong constraints  
28 helps". *Journal of Finance* 58, 1651–1684. 20
- 29 Jiang, G.J., Knight, J.L. (2002). "Efficient estimation of the continuous time stochastic volatility model via  
30 the empirical characteristic function". *Journal of Business and Economic Statistics* 20, 198–212. 21
- 31 Johannes, M., Polson, N.G. (2003). "MCMC methods for continuous-time financial econometrics". In: Ait-  
32 Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*. North-Holland, Amsterdam. In  
33 press. 22
- 34 Johannes, M., Polson, N., Stroud, J. (2004). "Sequential optimal portfolio performance: Market and volatility  
35 timing". Working Paper, Columbia University, University of Pennsylvania, and University of Chicago. 23
- 36 Johnson, T.D., Elashoff, R.M., Harkema, S.J.A. (2003). "Bayesian change point analysis of electromy-  
37 ographic data: Detecting muscle activation patterns and associated applications". *Biostatistics* 4, 143–164. 24
- 38 Johnson, H., Shanno, D. (1987). "Option pricing when the variance is changing". *Journal of Financial and* 25  
39 *Quantitative Analysis* 22, 143–152. 26
- 40 Jondeau, E., Rockinger, M. (2005). "The copula-GARCH model of conditional dependence: An international  
41 stock market application". *Journal of International Money and Finance*. In press. 27
- 42 Jorion, P. (2000). *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw–Hill, New York. 28
- 43 Karpoff, J.M. (1987). "The relation between price changes and trading volume: A survey". *Journal of Finan-* 29  
44 *cial and Quantitative Analysis* 22, 109–126. 30
- 45 Kawakatsu, H. (2005). "Matrix exponential GARCH". *Journal of Econometrics*. In press. 31
- 46 Kim, S., Shephard, N., Chib, S. (1998). "Stochastic volatility: Likelihood inference and comparison with  
47 ARCH models". *Review of Economic Studies* 65, 361–393. 32
- 48 King, M., Sentana, E., Wadhvani, S. (1994). "Volatility and links between national stock markets". *Econo-* 33  
49 *metrica* 62, 901–933. 34
- 50 Kroner, K.F., Ng, V.K. (1998). "Modelling asymmetric comovements of asset returns". *Review of Financial* 35  
51 *Studies* 11, 817–844. 36

- 1 Lamoureux, C.G., Lastrapes, W.D. (1990). "Persistence in variance, structural change, and the GARCH 1  
2 model". *Journal of Business and Economic Statistics* 8, 225–234. 2
- 3 Lamoureux, C.G., Lastrapes, W.D. (1994). "Endogenous trading volume and momentum in stock-return 3  
4 volatility". *Journal of Business and Economic Statistics* 14, 253–260. 4
- 5 Lastrapes, W.D. (1989). "Exchange rate volatility and US Monetary Policy: An ARCH application". *Journal 5  
6 of Money, Credit and Banking* 21, 66–77. 5
- 7 Ledoit, O., Santa-Clara, P., Wolf, M. (2003). "Flexible multivariate GARCH modeling with an application to 6  
8 international stock markets". *Review of Economics and Statistics* 85, 735–747. 7
- 9 Ledoit, O., Wolf, M. (2003). "Improved estimation of the covariance matrix of stock returns with an applica- 8  
10 tion to portfolio selection". *Journal of Empirical Finance* 10, 603–621. 9
- 11 Lee, S.W., Hansen, B.E. (1994). "Asymptotic theory for the GARCH(1, 1) quasi-maximum likelihood esti- 10  
12 mator". *Econometric Theory* 10, 29–52. 11
- 13 Lettau, M., Ludvigsson, S. (2003). "Measuring and modeling variation in the risk-return tradeoff". Working 11  
14 Paper, New York University and NBER. 12
- 15 Li, W.K., Ling, S., McAleer, M. (2002). "Recent theoretical results for time series with GARCH errors". 13  
16 *Journal of Economic Surveys* 16, 245–269. 14
- 17 Liesenfeld, R. (1998). "Dynamic bivariate mixture models: Modeling the behavior of prices and trading volu- 14  
18 me". *Journal of Business and Economic Statistics* 16, 101–109. 15
- 19 Liesenfeld, R. (2001). "A generalized bivariate mixture model for stock price volatility and trading volume". 16  
20 *Journal of Econometrics* 104, 141–178. 17
- 21 Liesenfeld, R., Richard, J.F. (2003). "Univariate and multivariate stochastic volatility models: Estimation and 18  
22 diagnostics". *Journal of Empirical Finance* 10, 505–531. 19
- 23 Ling, S., McAleer, M. (2003). "Asymptotic theory for a vector ARMA-GARCH model". *Econometric The- 19  
24 ory* 19, 280–310. 20
- 25 Linn, S.C., Zhu, Z. (2004). "Natural gas prices and the gas storage report: Public news and volatility in energy 21  
26 futures markets". *Journal of Futures Markets* 24, 283–313. 22
- 27 Longin, F., Solnik, B. (1995). "Is the correlation in international equity returns constant: 1970–1990?" *Journal 23  
28 of International Money and Finance* 14, 3–26. 24
- 29 Longin, F., Solnik, B. (2001). "Extreme correlation of international equity markets". *Journal of Finance* 56, 25  
30 649–676. 26
- 31 Loretan, M., Phillips, P.C.B. (1994). "Testing covariance stationarity under moment condition failure with an 26  
32 application to stock returns". *Journal of Empirical Finance* 1, 211–248. 27
- 33 Lumsdaine, R.L. (1996). "Consistency and asymptotic normality of the quasi-maximum likelihood estimator 28  
34 in GARCH(1, 1) and covariance stationary GARCH(1, 1) models". *Econometrica* 64, 575–596. 29
- 35 Lundin, M., Dacorogna, M.M., Müller, U.A. (1998). "Correlation of high frequency financial time series". 30  
36 In: Lequeux, P. (Ed.), *The Financial Markets Tick by Tick*. Wiley, London. 31
- 37 Lütkepohl, H. (2006). "Forecasting with VARMA models". In: Elliott, G., Granger, C.W.J., Timmermann, A. 31  
38 (Eds.), *Handbook of Economic Forecasting*. Elsevier, Amsterdam. This volume. 32
- 39 Maestas, C., Preuhs, R. (2000). "Modeling volatility in political time series". *Electoral Studies* 19, 95–110. 33
- 40 Marinova, D., McAleer, M. (2003). "Modeling trends and volatility in ecological patents in the USA". *Envi- 34  
41 ronmental Modelling and Software* 18, 195–203. 35
- 42 Markowitz, H. (1952). "Portfolio selection". *Journal of Finance* 7, 77–91. 36
- 43 Marquering, W., Verbeek, M. (2004). "The economic value of predicting stock index returns and volatility". 36  
37 *Journal of Financial and Quantitative Analysis* 39, 407–429. 37
- 38 Martens, M. (2003). "Estimating unbiased and precise realized covariances". Working Paper, University of 38  
39 Rotterdam. 39
- 40 Martin-Guerrero, J.D., Camps-Valls, G., Soria-Olivas, E., Serrano-Lopez, A.J., Perez-Ruixo, J.J., Jimenez- 40  
41 Torres, N.V. (2003). "Dosage individualization of erythropoietin using a profile dependent support vector 41  
42 regression". *IEEE Transactions on Biomedical Engineering* 50, 1136–1142. 42
- 43 McCullough, B., Renfro, C. (1998). "Benchmarks and software standards: A case study of GARCH procedu- 42  
43 res". *Journal of Economic and Social Measurement* 25, 59–71. 43

- 1 McNeil, A.J., Frey, R. (2000). "Estimation of tail-related risk measures for heteroskedastic financial time series: An extreme value approach". *Journal of Empirical Finance* 7, 271–300. 1
- 2 2
- 3 Meddahi, N. (2001). "An eigenfunction approach for volatility modeling". Working Paper, University of Montréal. 3
- 4 4
- 5 Meddahi, N., Renault, E. (2004). "Temporal aggregation of volatility models". *Journal of Econometrics* 119, 355–379. 5
- 6 6
- 7 Meghir, C., Pistaferri, L. (2004). "Income variance dynamics and heterogeneity". *Econometrica* 72, 1–32. 7
- 8 Melino, A., Turnbull, S.M. (1990). "Pricing foreign currency options with stochastic volatility". *Journal of Econometrics* 45, 239–265. 8
- 9 Merton, R.C. (1969). "Lifetime portfolio selection under uncertainty: The continuous-time case". *Review of Economics and Statistics* 51, 247–257. 9
- 10 Merton, R.C. (1976). "Option pricing when underlying stock returns are discontinuous". *Journal of Financial Economics* 3, 125–144. 10
- 11 11
- 12 Mikosch, T., Starica, C. (2004). "Nonstationarities in financial time series, the long range dependence and the IGARCH effects". *Review of Economics and Statistics* 86, 378–390. 12
- 13 13
- 14 Mills, T.C. (1993). *The Econometric Modelling of Financial Time Series*. Cambridge University Press, Cambridge. 14
- 15 15
- 16 Mincer, J., Zarnowitz, V. (1969). "The evaluation of economic forecasts". In: Mincer, J. (Ed.), *Economic Forecasts and Expectations*. National Bureau of Economic Research, New York. 16
- 17 Monfardini, C. (1998). "Estimating stochastic volatility models through indirect inference". *The Econometrics Journal* 1, C113–C128. 17
- 18 18
- 19 Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Puctet, O.V., von Weizsäcker, J. (1997). "Volatilities of different time resolutions – Analyzing the dynamics of market components". *Journal of Empirical Finance* 4, 213–239. 19
- 20 20
- 21 Nelson, D.B. (1988). "Time series behavior of stock market volatility and returns". Ph.D. dissertation, MIT. 21
- 22 Nelson, D.B. (1990). "Stationarity and persistence in the GARCH(1, 1) model". *Econometric Theory* 6, 318–334. 22
- 23 23
- 24 Nelson, D.B. (1991). "Conditional heteroskedasticity in asset returns: A new approach". *Econometrica* 59, 347–370. 24
- 25 25
- 26 Nijman, T., Sentana, E. (1996). "Marginalization and contemporaneous aggregation in multivariate GARCH processes". *Journal of Econometrics* 71, 71–87. 26
- 27 Ng, V.K., Engle, R.F., Rothschild, M. (1992). "A multi-dynamic-factor model for stock returns". *Journal of Econometrics* 52, 245–266. 27
- 28 28
- 29 O'Connell, P.E. (1971). "A simple stochastic modelling of Hurst's law". *Proceedings of International Symposium on Mathematical Models in Hydrology* 1, 169–187. 29
- 30 Pagan, A. (1996). "The econometrics of financial markets". *Journal of Empirical Finance* 3, 15–102. 30
- 31 Palm, F. (1996). "GARCH models of volatility". In: Rao, C.R., Maddala, G.S. (Eds.), *Handbook of Statistics*, vol. 14. North-Holland, Amsterdam, pp. 209–240. 31
- 32 32
- 33 Pan, J. (2002). "The jump-risk premia implicit in options: Evidence from an integrated time-series study". *Journal of Financial Economics* 63, 3–50. 33
- 34 34
- 35 Parkinson, M. (1980). "The extreme value method for estimating the variance of the rate of returns". *Journal of Business* 53, 61–65. 35
- 36 36
- 37 Pastor, L., Stambaugh, R. (2001). "The Equity premium and structural breaks". *Journal of Finance* 56, 1207–1245. 37
- 38 38
- 39 Patton, A. (2004). "Modeling asymmetric exchange rate dependence". Working Paper, London School of Economics. 39
- 40 Patton, A. (2005). "Volatility forecast evaluation and comparison using imperfect volatility proxies". Working Paper, London School of Economics. 40
- 41 41
- 42 Patton, A., Timmermann, A. (2003). "Properties of optimal forecasts". CEPR Discussion Paper 4037. 42
- 43 Patton, A., Timmermann, A. (2004). "Testable implications of forecast optimality". Working Paper, London School of Economics and University of California at San Diego. 43

- 1 Pelletier, D. (2005). "Regime switching for dynamic correlations". *Journal of Econometrics*. In press. 1
- 2 Perez-Quiros, G., Timmermann, A. (2000). "Firm size and cyclical variations in stock returns". *Journal of* 2
- 3 *Finance* 55, 1229–1262. 3
- 4 Pesaran, M.H., Zaffaroni, P. (2004). "Model averaging and value-at-risk based evaluation of large multi asset 4
- 5 volatility models for risk management". Working Paper, Department of Economics, University of Cam- 5
- 6 bridge. 6
- 7 Piazzesi, M. (2003). "Affine term structure models". In: Hansen, L.P., Ait-Sahalia, Y. (Eds.), *Handbook of* 6
- 7 *Financial Econometrics*. North-Holland, Amsterdam. In press. 7
- 8 Praetz, P.D. (1972). "The distribution of share price changes". *Journal of Business* 45, 49–55. 8
- 9 Pritsker, M. (2001). "The hidden dangers of historical simulation". Working Paper, Federal Reserve Board. 8
- 9 Ramirez, O.A., Fadiga, M. (2003). "Forecasting agricultural commodity prices with asymmetric error 9
- 10 GARCH models". *Journal of Agricultural and Resource Economics* 28, 71–85. 10
- 11 Rich, R., Tracy, J. (2004). "Uncertainty and labor contract durations". *Review of Economics and Statistics* 86, 11
- 12 270–287. 12
- 13 Richardson, M., Smith, T. (1994). "A direct test of the mixture of distributions hypothesis: Measuring the 13
- 14 daily flow of information". *Journal of Financial and Quantitative Analysis* 29, 101–116. 14
- 15 Robinson, P.M. (1991). "Testing for strong serial correlation and dynamic conditional heteroskedasticity in 15
- 16 multiple regression". *Journal of Econometrics* 47, 67–84. 16
- 17 Rossi, P.E. (1996). *Modeling Stock Market Volatility: Bridging the Gap to Continuous Time*. Academic Press, 17
- 17 San Diego. 17
- 18 Ruge-Murcia, F.J. (2003). "Inflation targeting under asymmetric preferences". *Journal of Money, Credit and* 18
- 19 *Banking* 35, 763–785. 19
- 20 Ruge-Murcia, F.J. (2004). "The inflation bias when the central bank targets the natural rate of unemployment". 20
- 21 *European Economic Review* 48, 91–107. 21
- 22 Sandmann, G., Koopman, S.J. (1998). "Estimation of stochastic volatility models through Monte Carlo max- 22
- 23 imum likelihood". *Journal of Econometrics* 87, 271–301. 23
- 24 Scholes, M., Williams, J. (1977). "Estimating betas from non-synchronous data". *Journal of Financial Eco-* 24
- 25 *nomics* 5, 309–327. 25
- 26 Schwert, G.W. (1989). "Why does stock market volatility change over time?" *Journal of Finance* 44, 1115– 26
- 27 1153. 27
- 28 Schwert, G.W. (1990). "Stock volatility and the crash of '87". *Review of Financial Studies* 3, 77–102. 28
- 29 Scott, L.O. (1987). "Option pricing when the variance changes randomly: Theory, estimation and an applica- 29
- 30 tion". *Journal of Financial and Quantitative Analysis* 22, 419–438. 30
- 31 Sentana, E., Fiorentini, G. (2001). "Identification, estimation and testing of conditionally heteroskedastic 31
- 32 factor models". *Journal of Econometrics* 102, 143–164. 32
- 33 Shanken, J.A. (1990). "Intertemporal asset pricing: An empirical investigation". *Journal of Econometrics* 45, 33
- 34 99–120. 34
- 35 Sharpe, W. (1964). "Capital asset prices – A theory of market equilibrium under conditions of risk". *Journal* 35
- 36 *of Finance* 19, 425–442. 36
- 37 Shawkly, M.A., Marathe, A., Barrett, C.L. (2003). "A first look at the empirical relation between spot and 37
- 38 futures electricity prices in the United States". *Journal of Futures Markets* 23, 931–955. 38
- 39 Shephard, N. (1996). "Statistical aspects of ARCH and stochastic volatility models". In: Cox, D.R., Hinkley, 39
- 40 D.V., Barndorff-Nielsen, O.E. (Eds.), *Time Series Models in Econometrics, Finance and Other Fields*. 40
- 41 Chapman and Hall, London, pp. 1–67. 41
- 42 Shephard, N. (2004). *Stochastic Volatility: Selected Readings*. Oxford University Press, Oxford. 42
- 43 Sheppard, K. (2004). "Economic factors and the covariance of equity returns". Working Paper, University of 43
- 44 California, San Diego. 44
- 45 Singleton, K.J. (2001). "Estimation of affine asset pricing models using the empirical characteristic function". 45
- 46 *Journal of Econometrics* 102, 111–141. 46
- 47 Smith Jr., A.A. (1990). "Three essays on the solution and estimation of dynamic macroeconomic models". 47
- 48 Ph.D. dissertation, Duke University. 48

- 1 **Smith Jr., A.A.** (1993). "Estimating nonlinear time-series models using simulated vector autoregressions". 1  
2 *Journal of Applied Econometrics* 8, S63–S84. 2
- 3 **Tauchen, G., Pitts, M.** (1983). "The price variability-volume relationship on speculative markets". *Economet-* 3  
4 *rica* 51, 485–505. 4
- 5 **Taylor, S.J.** (1986). *Modeling Financial Time Series*. Wiley, Chichester. 5
- 6 **Taylor, S.J.** (2004). *Asset Price Dynamics and Prediction*. Princeton University Press, Princeton. 6
- 7 **Taylor, J.W., Buizza, R.** (2003). "Using weather ensemble predictions in electricity demand forecasting". 7  
8 *International Journal of Forecasting* 19, 57–70. 8
- 9 **Tiao, G.C., Tsay, R.S.** (1994). "Some advances in non-linear and adaptive modeling in time series". *Journal* 9  
10 *of Forecasting* 14, 109–131. 10
- 11 **Tsay, R.S.** (2002). *Analysis of Financial Time Series*. Wiley, New York. 11
- 12 **Tse, Y.K., Tsui, A.K.C.** (2002). "A multivariate GARCH model with time-varying correlations". *Journal of* 12  
13 *Business and Economic Statistics* 20, 351–362. 13
- 14 **Tse, Y.K., Yip, P.S.L.** (2003). "The impacts of Hong Kong's currency board reforms on the interbank market". 14  
15 *Journal of Banking and Finance* 27, 2273–2296. 15
- 16 **Wang, L.** (2004). "Investing when volatility fluctuates". Working Paper, The Wharton School and Singapore 16  
17 Management University. 17
- 18 **Weiss, A.A.** (1986). "Asymptotic theory for ARCH models: Estimation and testing". *Econometric Theory* 2, 18  
19 107–131. 19
- 20 **West, K.D.** (1996). "Asymptotic inference about predictive ability". *Econometrica* 64, 1067–1084. 20
- 21 **West, K.D., Cho, D.** (1995). "The predictive ability of several models of exchange rate volatility". *Journal of* 21  
22 *Econometrics* 69, 367–391. 22
- 23 **West, K.D., McCracken, M.W.** (1998). "Regression-based tests of predictive ability". *International Economic* 23  
24 *Review* 39, 817–840. 24
- 25 **White, H.** (2000). "A reality check for data snooping". *Econometrica* 68, 1097–1127. 25
- 26 **Whitelaw, R.F.** (1997). "Time-varying Sharpe ratios and market timing". Working Paper, NYU, Stern School 26  
27 of Business. 27
- 28 **Wiggins, J.B.** (1987). "Option values under stochastic volatility: Theory and empirical estimates". *Journal of* 28  
29 *Financial Economics* 19, 351–372. 29
- 30 **Zaffaroni, P.** (2004). "Estimating and forecasting volatility with large scale models: Theoretical appraisal of 30  
31 professional practice". Working Paper, Banca d'Italia. 31
- 32 **Zakoian, J.-M.** (1994). "Threshold heteroskedastic models". *Journal of Economic Dynamics and Control* 18, 32  
33 931–955. 33
- 34 **Zhang, L., Ait-Sahalia, Y., Mykland, P.A.** (2005). "A tale of two time scales: Determining integrated volatility 34  
35 with noisy high-frequency data". *Journal of the American Statistical Association*. In press. 35
- 36 **Zhou, B.** (1996). "High-frequency data and volatility in foreign exchange rates". *Journal of Business and* 36  
37 *Economic Statistics* 14, 45–52. 37
- 38 38
- 39 39
- 40 40
- 41 41
- 42 42
- 43 43

# Proof of Raw Subject Index

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43

---

**Page: 780**

volatility  
Wiener  
continuous-time  
diffusion  
implied volatility  
forecasting  
return volatility  
uncertainty  
price process

---

**Page: 781**

volatility forecasting  
volatility modeling

---

**Page: 782**

ARCH  
GARCH  
conditional expectation  
information set  
volatility process  
filtering problem  
stochastic volatility models  
model specification  
forecast evaluation tests

---

**Page: 783**

noise  
realized volatility  
stochastic differential equation  
drift  
Brownian motion  
diffusive coefficient  
spot volatility

---

**Page: 784**

integrated variance  
integrated volatility  
squared returns

---

**Page: 785**

forecast quality  
simulated series

---

**Page: 786**

fully efficient forecasts  
higher frequency “intraday” data  
discrete-time  
univariate volatility forecasting  
stochastic volatility  
realized volatility  
multivariate  
forecasting conditional covariances and  
correlations  
forecast evaluation

---

**Page: 787**

forecast loss function  
point forecast

---

**Page: 788**

conditional volatility  
optimal forecast  
interval forecasting  
probability forecasting  
sign forecasting

---

**Page: 789**

volatility dynamics  
density forecasting

---

**Page: 790**

risk management  
Value-at-Risk  
Expected Shortfall

---

**Page: 792**

covariance risk  
time-varying betas  
conditional Sharpe ratios  
stochastic discount factor

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43

1			1
2	<b>Page: 793</b>	<b>Page: 805</b>	2
3	covariance forecasting	long memory	3
4	asset allocation	Fractionally Integrated FIGARCH	4
5	time-varying covariance		5
6	CAPM	<b>Page: 806</b>	6
7	<b>Page: 794</b>	Fractional Integrated EGARCH	7
8	option valuation	FIEGARCH	8
9	dynamic volatility	Component GARCH	9
10	Black–Scholes–Merton (BSM) option pricing	volatility impulse response function	10
11	model		11
12	stochastic volatility	<b>Page: 807</b>	12
13	<b>Page: 795</b>	Maximum Likelihood Estimation	13
14	integrated volatility		14
15	integrated variance	<b>Page: 808</b>	15
16	realized volatility	Bayesian inference	16
17		corresponding Quasi-Maximum Likelihood	17
18	<b>Page: 796</b>	Estimates	18
19	volatility forecasting	QMLE	19
20			20
21	<b>Page: 798</b>	<b>Page: 809</b>	21
22	GARCH volatility	density forecasting	22
23	rolling regressions	forecast distributions	23
24	RiskMetrics		24
25	rolling sample windows	<b>Page: 810</b>	25
26		quantile predictions	26
27	<b>Page: 799</b>	Extreme Value Theory	27
28	filter	EVT	28
29	data generating process		29
30		<b>Page: 811</b>	30
31	<b>Page: 800</b>	Monte Carlo simulations	31
32	Integrated GARCH		32
33	IGARCH	<b>Page: 813</b>	33
34		Generalized Error Distribution	34
35	<b>Page: 801</b>		35
36	term-structure-of-variance	<b>Page: 814</b>	36
37		Stochastic Volatility	37
38	<b>Page: 803</b>	SV	38
39	“leverage” effect @l	latent variables	39
40	“volatility feedback” effect @v	filtering	40
41	Threshold GARCH	smoothing	41
42	TGARCH	simulated moments, Markov Chain Monte Carlo	42
43	Asymmetric GARCH	(MCMC)	43
	AGARCH		
	Exponential GARCH	<b>Page: 815</b>	
	EGARCH	model specification	
		Mixture of Distributions Hypothesis	
		MDH	



1			1
2	<b>Page: 816</b>	<b>Page: 830</b>	2
3	mixture model	realized volatility	3
4		high-frequency data	4
5	<b>Page: 817</b>		5
6	Stochastic Autoregressive Volatility	<b>Page: 832</b>	6
7	SARV	semi-martingale	7
8			8
9	<b>Page: 818</b>	<b>Page: 834</b>	9
10	jump	realized volatility modeling	10
11	return variability		11
12	return variance		12
13		<b>Page: 835</b>	13
14	<b>Page: 819</b>	diffusive volatility	14
15	conditional variance	jump component	15
16		realized volatility forecasting	16
17	<b>Page: 820</b>		17
18	latent volatility process	<b>Page: 836</b>	18
19		sampling frequency	19
20	<b>Page: 821</b>		20
21	moment matching	<b>Page: 838</b>	21
22	Kalman filter	implied volatility	22
23	state space		23
24		<b>Page: 839</b>	24
25	<b>Page: 822</b>	multivariate volatility	25
26	transition equation	portfolio weights	26
27			27
28	<b>Page: 823</b>	<b>Page: 840</b>	28
29	efficient method of simulated moments	exponential smoothing	29
30	Efficient Method of Moments	RiskMetrics	30
31	EMM	time-varying covariance matrices	31
32	Method of Simulated Moments (MSM)		32
33		<b>Page: 841</b>	33
34	<b>Page: 825</b>	multivariate GARCH	34
35	simulation error		35
36		<b>Page: 842</b>	36
37	<b>Page: 826</b>	conditional covariance matrix	37
38	Markov Chain Monte Carlo (MCMC)		38
39	inference	<b>Page: 843</b>	39
40	forecasting	forecast horizon	40
41	Gibbs sampler		41
42		<b>Page: 844</b>	42
43	<b>Page: 827</b>	vech model	43
	particle filter		
	out-of-sample forecasting	<b>Page: 845</b>	
		dynamic conditional correlations	
		Constant Conditional Correlation	

1	<hr/>	volatility forecast evaluation	1
2	<b>Page: 846</b>	<hr/>	2
3	likelihood function	<b>Page: 856</b>	3
4	<hr/>	Mincer–Zarnowitz volatility regression	4
5	<b>Page: 847</b>	daily volatility	5
6	multivariate stochastic volatility	<hr/>	6
7	factor models	<b>Page: 857</b>	7
8	Arbitrage Pricing Theory (APT)	volatility proxy	8
9	<hr/>	bid–ask bounces effect	9
10	<b>Page: 849</b>	<hr/>	10
11	realized covariances and correlations	<b>Page: 858</b>	11
12	realized covariation	volatility signature plot	12
13	covariation	<hr/>	13
14	realized volatility errors	<b>Page: 859</b>	14
15	<hr/>	interval forecast	15
16	<b>Page: 850</b>	Value-at-Risk evaluation	16
17	realized covariation	<hr/>	17
18	realized correlation	<b>Page: 860</b>	18
19	<hr/>	interval forecast	19
20	<b>Page: 851</b>	probability forecast evaluation	20
21	regime switching type models	market timing tests	21
22	<hr/>	<hr/>	22
23	<b>Page: 853</b>	<b>Page: 861</b>	23
24	market microstructure frictions	density forecast evaluation	24
25	<hr/>	Probability Integral Transform (PIT)	25
26	<b>Page: 854</b>	<hr/>	26
27	forecasting volatility	<b>Page: 863</b>	27
28	Value-at-Risk	Ljung–Box statistics	28
29	density forecasting	serial correlation	29
30	forecast loss function	<hr/>	30
31	optimal forecast	<b>Page: 864</b>	31
32	<hr/>	forecast combination	32
33	<b>Page: 855</b>		33
34	forecast evaluation regression		34
35			35
36			36
37			37
38			38
39			39
40			40
41			41
42			42
43			43