Essays on nonparametric econometrics of stochastic volatility

Zu, Y.

Chapter 1

Introduction

In financial economics, it is general practice to model the time series of a financial instrument prices as a stochastic process. According to Back (1991), under the assumption of no arbitrage opportunities and a finite expected return, within a frictionless setting, the efficient logarithmic price process must be a special martingale, which is a martingale that admits a unique canonical decomposition (Protter (2005)).

In this thesis, I focus on a subclass of special semimartingales, namely the one-dimensional Brownian semimartingale, as the model of efficient logarithmic price process. Let $\{X_t\}_{t\geqslant 0}$ be the efficient logarithmic price process, as a Brownian semimartingale it has the following representation:

$$X_t = X_0 + \int_0^t \mu_s \mathrm{d}s + \int_0^t \sigma_s \mathrm{d}W_s, \tag{1.1}$$

where μ and σ are predictable processes and W is a standard Brownian motion.

The special semimartingale is a very large class of stochastic processes and in particular it allows discontinuities in the sample paths of the process and it includes multivariate processes. However, the Brownian semimartingale model I consider here is a univariate process with continuous sample paths. Models for multivariate price process is an interesting and important topic by itself, but it is not considered in this thesis. Although jumps occur in realistic financial price data, they need a separate methodology of modeling, and the model we consider here can be viewed as being applicable to the data with jumps filtered out.

Definition of return and volatility

Financial return is fundamental to any financial analysis. Although there are various notions of financial returns (see Campbell, Lo, and MacKinlay (1997), Tsay (2005)), the one we use here is the so-called logarithmic return. Let $r(t, h)$ denote the logarithmic return at time t over the period h, then

$$r(t, h) := X_t - X_{t-h} = \int_{t-h}^t \mu_s \mathrm{d}s + \int_{t-h}^t \sigma_s \mathrm{d}W_s.$$

Notice the definition needs the period h to be meaningful, besides the time instance t.

Volatility is a concept that describes the variation of financial returns, it has important applications in a number of financial decision-making processes, such as risk management, asset allocation and option pricing. There are different notions of volatility, see Andersen, Bollerslev, and Diebold (2010) for an overview. One volatility notion that in particular is relevant to this thesis is the *integrated volatility*. Let the integrated volatility at time t over the period h be denoted $\mathrm{IV}(t, h)$, then

$$\mathrm{IV}(t, h) := \int_{t-h}^t \sigma_s^2 \mathrm{d}s.$$

Integrated volatility is also meaningful for a specific time period. Another notion of volatility is the *spot volatility*. The spot volatility at time t is just σ_t^2, it is related to the integrated volatility with the relation

$$\sigma_t^2 = \lim_{h \to 0} \frac{1}{h} \int_{t-h}^t \sigma_s^2 \mathrm{d}s.$$

The spot volatility is defined for a time point t and it does not need a specific time period to be meaningful. σ_t^2 represents the volatility per unit of time, so if t is measured in years, then σ_t^2 represents volatility on an annual basis.

Notice that the volatility notions in this thesis are related to the variance instead of the standard deviation, and we follow this convention throughout.

Stochastic volatility models

The subject of this thesis is stochastic volatility models. Although the definition of stochastic volatility is not universal, the concept of stochastic volatility in this thesis is the same as that of the *genuine stochastic volatility* as defined in Andersen and Benzoni (2009), "where stochastic volatility refers to a modeling scheme for the volatility process

in which the return variation dynamics include an unobservable shock which cannot be predicted using current available information."

The efficient logarithmic price process (1.1) we consider in the beginning already has a generic *stochastic* volatility form, because of the randomness of the σ process. However, (1.1) is not the actual starting point of the stochastic volatility modeling literature from a historical perspective, although it is the unified framework we use nowadays. Shephard and Andersen (2009) discuss in detail the origins of the modern stochastic volatility modeling literature, which is summarized here.

The first origin of stochastic volatility models is the empirical modeling literature of financial volatility. The popularity of empirical volatility models start from the ARCH model introduced in Engle (1982) and the GARCH model introduced in Bollerslev (1986), which successfully model the stylized facts present in financial returns, namely a fat tailed return distribution and volatility clustering, while at the same time retaining simplicity in estimation and inference. From an empirical, or statistical, perspective, (some) stochastic volatility models have been developed as extensions to these ARCH type models. Unlike the ARCH type model, where the volatility is deterministic conditional on the return history, the stochastic volatility model introduce a separate stochastic term for the evolution of volatility. This stochastic specification improves the flexibility of the model to capture the fat-tailedness of the return distribution; furthermore, the newly appeared correlation coefficient between this stochastic volatility shock and the return shock makes the so-called leverage effect easier to model. Discrete-time stochastic volatility models are usually derived from this origin, for example, the Stochastic Autoregressive Volatility class of model as described in Andersen et al. (2010). Some continuous-time models are found by deriving the continuous-time limit of discrete-time volatility models; one example is the the GARCH diffusion process by Nelson (1990).

Another origin of stochastic volatility models is the option pricing literature. In financial economics, stochastic volatility models are developed mainly for the purpose of option pricing. The well known Black-Scholes model uses a Geometric Brownian motion process to model the underlying price process, in which the volatility parameter for the logarithmic price process is a constant. However, a realistic logarithmic price process does not seem to have a constant diffusion parameter, and furthermore, the volatilities implied by observed option prices through the Black-Scholes formula exhibit the well-known phenomenon of a volatility smile. A natural approach to cope with these mismatches is to let the volatility parameter be a stochastic process itself — leading to continuous-time stochastic volatility models. A well-known example of this type is the Heston (1993)

model. Although empirical soundness is an aspect to consider, analytical tractability of option pricing formula is also a very important aspect for modeling, as reflected in the popularity of the Heston model.

Estimation of stochastic volatility models is known to be difficult, because of the unobservability of the volatility process and the intractability of the likelihood function. However, significant progress has been made in the past decades, for example, the Method of Moments by Taylor (1986), the Generalized Method of Moments by Melino and Turnbull (1990), the Simulated Method of Moments by Duffie and Singleton (1993), the Gaussian Quasi-Maximum Likelihood Estimator by Harvey, Ruiz, and Shephard (1994), the Bayesian Markov Chain Monte Carlo by Jacquier, Polson, and Rossi (2002) and Kim, Shephard, and Chib (1998), the Efficient Method of Moments by Gallant, Hsieh, and Tauchen (1997), the Simulated Maximum Likelihood Estimator by Danielsson (1994) and the Monte Carlo Maximum Likelihood Estimator by Sandmann and Koopman (1998).

Although the field of stochastic volatility modeling is dominated by parametric approaches, nonparametric methods also exist. Nonparametric estimation of spot volatility has been considered by Nelson (1992), Nelson and Foster (1995), Foster and Nelson (1996), Mykland and Zhang (2008), Fan and Wang (2008) and Kristensen (2010). Other nonparametric approaches include Franke, Härdle, and Kreiß (2003), Van Es, Spreij, and Van Zanten (2003) and Van Es, Spreij, and Van Zanten (2005), who nonparametrically estimate the stationary density of the volatility process; Reno (2006), Reno (2008), Comte, Genon-Catalot, and Rozenholc (2010) and Comte, Lacour, and Rozenholc (2010), who nonparametrically estimate the drift and diffusion functions of the volatility process.

High-frequency data and estimating intergrated volatility

Stochastic volatility models were developed at a time when only low frequency data were available. The advent of commonly available high-frequency data (such as that introduced in Dacorogna, Gençay, Muller, Olsen, and Pictet (2001)), such as minute-by-minute data, have made a regime shift in the field of volatility measurement. Instead of estimating stochastic volatility models using daily (or lower frequency) data and infering volatility from the model, researchers have started to be interested in measuring daily integrated volatility using a *realized volatility* measure based on high frequency data; two seminal papers on this topic are Andersen, Bollerslev, Diebold, and Labys (2001) and Barndorff-Nielsen and Shephard (2002).

However, when using ultra high frequency data (e.g. second-by-second data or even

higher), the so-called market microstructure noise in such data will prevent the realized volatility estimator to deliver a reasonable estimate of integrated volatility. Several estimators have been proposed to correct the effects of noise. These include the first unbiased (but inconsistent) estimator by Zhou (1996); the subsampling based Two Scale Realized Variance (TSRV) estimator by Zhang, Mykland, and Aït-Sahalia (2005) and the Multi-scale Realized Variance (MSRV) estimator by Zhang (2006); the Realized Kernel (RK) estimator by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2010); the pre-averaging method by Podolskij and Vetter (2009), and Jacod, Li, Mykland, Podolskij, and Vetter (2009); and the recently proposed Quasi-Maximum Likelihood Estimation (QMLE) estimator by Xiu (2010), which is based on the likelihood function from the constant variance model in Aït-Sahalia, Mykland, and Zhang (2005).

With the success of estimating integrated volatility with high frequency data, subsequent research has tried to exploit the information in the estimated daily realized volatility measures. Shephard and Sheppard (2010) propose a so-called HEAVY model to jointly model daily returns and the estimated daily realized volatility, using a GARCH structure. Another strand of research tries to use the estimated realized volatility to estimate continuous time stochastic volatility models. These works exploit the moments implications of integrated volatility to construct Generalized Method of Moments type estimators of stochastic volatility models, which include Bollerslev and Zhou (2002), Barndorff-Nielsen and Shephard (2002), Corradi and Distaso (2006) and Todorov (2009). Both strands of research will bring new insights to our understanding of the evolution of the volatility process.

I have given a brief review on the stochastic volatility literature, with particular emphasis on those topics that are relevant to this thesis. This certainly cannot cover all important topics and papers in this vast field, such as literature on jump models, long memory models and multivariate high dimensional models.

Overview of the chapters

This thesis deals with applying nonparametric methods to various estimation and testing problems in stochastic volatility models. Nonparametric methods allow for much larger classes of models than classical parametric methods, thus provide more flexibility and less opportunity for misspecification. This is in particular useful in the econometrics of stochastic volatility, where no universally agreed good model for volatility exists up to now. The price to pay for the additional flexibility of nonparametric methods is lower

statistical efficiency and higher computational costs. However, the advance of information technology has led to easier access to large dataset in the finance industry, making the econometrics of stochastic volatility a perfect area to apply nonparametric methods. Furthermore, with the rapid improvement in computing power for academic use, the application of highly computational nonparametric methods is no longer infeasible nowadays.

Chapter 2 is based on Zu and Boswijk (2010), it is devoted to nonparametrically measuring spot volatility with high frequency financial data. Spot volatility, also called instantaneous volatility, measures the strength of variation in return at a certain time point. Based on the Two Scale Realized Variance estimator by Zhang et al. (2005), a spot volatility estimator for high frequency financial data is constructed. Consistency and the asymptotic distribution of the estimator are derived; the estimator converges to a mixed normal distribution. A two-step data-driven method is proposed to select the scale parameter and bandwidth parameter in the estimator. In carefully designed and practically meaningful Monte Carlo simulations, the usefulness of the estimator and the two-step method is demonstrated by comparison with a variety of existing methods.

Stochastic volatility model validation is another subject of this thesis, Chapter 3 and 4 are devoted to this topic. Although the estimation of stochastic volatility models has attracted much interest, specification tests for such models is still an underdeveloped area — this may due to the fact that stochastic volatility models are so difficult to estimate that specification testing has been considered of secondary importance. However, with the progress made in the estimation of stochastic volatility models in the past decade, specification test procedures for such models become more important. These two chapters try to fill this need, and test the stochastic volatility model specification by comparing certain aspects of the model with the corresponding nonparametrically estimated alternatives. Since there is no uniformly most powerful test in a nonparametric testing problem, these tests are useful complements to each other.

Chapter 3 is based on Zu and Boswijk (2009), it develops nonparametric specification tests for stochastic volatility models by comparing the nonparametically estimated return density and distribution functions with their parametric counterparts, respectively. We apply our tests to an empirical example and find some of our tests to show a remarkable ability to reject models.

Chapter 4 is based on Zu (2010), it develops a specification test for stochastic volatility models by comparing the nonparametric *kernel deconvolution density estimator* of the stationary integrated volatility density with its parametric counterpart, where the L_2

distance is used to measure the discrepancy. The asymptotic null distribution of the test is established, and the consistency and the local power properties of the test are derived. The parametric integrated volatility density of a stochastic volatility model usually does not have a closed-form expression, for which two approximation methods are developed. The performance of the approximation methods together with the finite sample properties of the test are studied in Monte Carlo experiments, and the test is applied to an empirical example.