

## Summary

Volatility refers to the degree to which financial prices fluctuate. Financial markets pass through calm and hectic periods: volatility is not constant, and high volatility (as well as low volatility) tends to cluster. Modelling and forecasting financial market volatility has been an active research area in financial economics since the 1980s. Models for time varying volatility play an important role in financial risk management, portfolio allocation, option pricing, and testing economic theories.

A classical way of modelling volatility is by using GARCH models. Research in the 1990s showed that discrete time GARCH models give surprisingly good forecasts of volatility when applied to daily financial returns. Over the past decade intraday high-frequency data became widely available. These data, recording thousands of price ticks per day, provide a useful source of information in addition to daily close-to-close returns.

This thesis introduces a continuous time asset price model, the *scaling model*, that incorporates the intraday price movements into the discrete time GARCH model. The scaling model is quite versatile and poses no constraints on the intraday price process (in formal terms, it is nonparametric). The model greatly enhances classical GARCH models based on daily returns. We develop GARCH parameter estimators that use high-frequency data. These estimators are shown to be up to twenty times more efficient than classical estimators in practical applications. The thesis also develops a GARCH recursion for incorporating high-frequency data into volatility forecasts. These forecasts outperform forecasts based on classical GARCH models.

The scaling model allows us to introduce *volatility proxies*. Volatility itself is unobservable. This is why we work with proxies, i.e. variables that summarize per day the fluctuations in the intraday price process. A good proxy is a good stand-in for the day's volatility. We define a large class of proxies. This class includes commonly applied statistics such as the absolute close-to-close return, the intraday high-low range, and the realized volatility. The thesis presents a theory for optimizing volatility proxies. Proxies can be combined in a simple way into a single, more efficient proxy (without having to make model assumptions for the GARCH volatilities, nor for the intraday price process). The successful use of volatility proxies forms the basis for the mentioned improvements in GARCH parameter estimation and volatility forecasting.

It is a stylized fact that decreasing equity prices tend to be followed by an increase in volatility. Our empirical applications to the S&P 500 index suggest that the downward price movements are the main driver of the daily volatility process. One obtains good proxies by summing absolute returns and high-low ranges over 10-minute intervals rather than summing squared values. The scaling model gives a good description for the S&P 500 index over the years 1988–2008. The scaling model describes financial processes in the words: “In probabilistic terms every day is the same day, up to a random scale factor.

This scale factor is fixed within the day and intensifies (or diminishes) all price movements during the day.”

Our way of using high-frequency data for modelling daily volatility differs from approaches in the current literature. Nowadays it is common practice to assume a semimartingale model for financial processes and to interpret volatility in terms of quadratic variation, which may be estimated up to a certain degree of accuracy by summing squared returns over five or two-minute intervals, the so-called realized quadratic variation (or realized variance). The theory presented in this thesis does not use the theory of semimartingales. The scaling model may, but need not, be a semimartingale. The scaling model is a practical and useful model that may be seen as an alternative to semimartingale models. It gives good measurements and forecasts of GARCH volatility, but does not strive to conform to economic principles such as no-arbitrage. This approach has its advantages. We do not have to introduce financial market microstructure effects to explain why prices do not satisfy the martingale property over short time intervals, nor do we have to worry how traders make a living in a world where arbitrage is impossible. Our approach allows one to easily, effectively, and realistically incorporate intraday price movements into discrete time GARCH models.