1

Introduction

This book is about the modeling of asset returns when the normality assumption does not apply. It provides an up-to-date and step-by-step description of the tools that are useful for the modeling of non-Gaussian asset return distributions and for option pricing in the non-Gaussian context.

1.1 Financial markets and financial time series

For more than four decades, distributions of financial asset returns have been known to be non-Gaussian (see Mandelbrot, 1963, and Fama, 1965). The assumption of normality is stacked against two hard facts: First, the empirical distributions of asset returns have tails thicker than those from a normal distribution and appear to be negatively skewed. This means more extreme negative values, which has a very serious implication for risk management and portfolio selection. Second, returns are time dependent. Squared returns, absolute returns, and all measures and proxies of volatility exhibit strong serial correlation. This is now known as volatility clustering or conditional heteroskedasticity (Engle, 1982).

Financial modeling is all about capturing and exploiting patterns in the data including the two phenomena mentioned above. Chapter 2 discusses the unique statistical properties of financial market data and several so-called stylized facts. These stylized facts will be the basis for Part II where each chapter will tackle some specific features of financial market returns.

Chapter 3 describes the actual functioning and the microstructure of financial markets. Here, we present some theoretical models that may help explaining why asset returns are non-normal and time dependent. The foundation is built on Clark (1973) who postulates that non-normality and volatility clustering could be due to intermittent information arrivals.

1.2 Econometric modeling of asset returns

Part II is concerned with the time series aspects of asset returns. Chapters 4, 5, and 7 cover models for the second, third, and fourth moments and the tails of return distributions. These higher moments and tail measures are the hallmarks of non-Gaussian distributions. Chapter 6 deals with the dependence structure when the higher moments display significant departure from normality and returns appear to be time dependent. The dependence among the tail observations, described in Chapter 7, is very different from the dependence, described in Chapter 6, of the central and main part of the distributions, because of the differences in the underpinning statistical theories and the important fact that financial markets do behave very differently between normal and crisis periods.

Specifically, Chapter 4 covers models for volatility that include the better known GARCH (Generalized Autoregressive and Conditional Heteroskedasticity) class of models and some new extensions such as GARCH models with jumps and realized volatility models. With high-frequency data becoming more common these days, realized volatility is expected to remain an area of active research. This chapter also describes the lesser known or lesser discussed issues on GARCH aggregation and the relationship between stochastic volatility model in continuous time and the discrete time GARCH model.

Although time-varying volatility and volatility asymmetry may produce thick-tail and asymmetric distributions in asset return, volatility alone cannot explain away all the non-normality. To fully capture return distributions, we also need models for skewness and kurtosis. Chapter 5 does exactly that by fitting time-varying higher-moment conditional models to returns. It also describes tests for the adequacy of these conditional high moment models.

Chapters 4 and 5 are concerned with univariate time series characteristics. Chapter 6 shifts the focus to the relationships between and among the asset return series. This involves two main tasks. First, we have to extend the GARCH family to a (possibly large) number of assets in order to reproduce the joint dynamic of volatility. Second, we have to capture in this multivariate framework the non-normality of returns. Here, we move from a multivariate GARCH model with normal distribution to one with skewed Student t distribution that is designed to capture both fat-tailedness and asymmetry. We also examine an alternative approach that circumvents some difficulties in designing a multivariate distribution. The so-called copula approach is a tool that is able to join any type of marginal distribution. It has many theoretical appeals. But in many finance applications in practice, integration of the joint distribution is needed. This cannot be done analytically in the copula approach and will become more and more cumbersome as the number of assets or the dimension of the problem increases.

Chapter 7 presents an approach that is very different from all the other chapters in this part. It deals with models for only the tails of the distribution. We describe in this chapter various approaches for characterizing the behavior of extreme events. In particular, we explain how to model the distribution of maxima over subsamples and the distribution of exceedances above a high threshold. In a multivariate context, we also highlight the important difference between asymptotic dependence and asymptotic independence and describe some non-parametric statistical measures for both.

1.3 Applications of non-Gaussian econometrics

Part III presents some examples of applications of the models described in Part II. All these applications are important and of routine use in the finance industry. Readers should be convinced, after reading this part, that the non-Gaussian models are not only indispensable in financial modeling, but they can also be very rewarding. Specifically, Chapter 8 deals with risk management and the Value-at-Risk (VaR) measure introduced by the Basel Accords. The industry benchmark, the RiskMetrics model, is based on normal distributions. We describe alternative techniques that are more appropriate for non-Gaussian distributions.

Chapter 9 is concerned with portfolio construction and asset allocation. Markowitz's mean-variance analysis is appropriate for Gaussian distributions or quadratic utility function only.[1] In the context of non-normal returns, this approach may not hold anymore. The main idea is that the investor's expected utility may be approximated as a function of mean, variance, but also of higher moments of the portfolio return. A rational investor would be averse to negative skewness and high kurtosis and in favor of positive skewness. We will also show in this chapter how downside risk constraint affects asset allocation decisions.

1.4 Option pricing with non-Gaussian distributions

Part IV deals with derivative assets and considers option pricing when the underlying asset return has a non-Gaussian distribution. The seminal contributions by Black and Scholes (1973) and Merton (1973) laid the foundation of pricing by no-arbitrage and, later, pricing by equivalent martingale measure. This model, which essentially assumes normality and time independence of price changes, has been shown for a long time to be unable to reproduce some well-known stylized facts such as the volatility smile or the term structure of volatilities. As for the modeling of asset returns, option pricing models have to incorporate the volatility clustering and the non-normality of the conditional distribution. Among these models, the most well-known is the stochastic volatility model and the models with jumps. Due to the mathematical

[1] Only the mean and variance terms are relevant when asset pricing is based on a quadratic utility function.

content of these option pricing models, we have provided five support chapters in the appendices in Part V, so that Part IV is not overly cluttered by mathematical abstraction. We have made every effort in ensuring that these two parts are accessible to non-mathematician readers.

In Chapter 10, we go through the fundamental building blocks of the Black-Scholes-Merton (BSM) model and use it as an example to introduce the key mathematical concepts such as Brownian motion and stochastic calculus (Chapter 13) and martingale and changing measure (Chapter 14). The BSM model is based on the underlying asset return having a normal distribution. Almost as soon as the importance of this model was recognized, the implied volatility smile was reported, indicating that the normality assumption is inconsistent with option price data. Nevertheless and despite the BSM pricing irregularities, the popularity of the BSM model survives even today.

From the BSM model it emerges that options can be priced using risk neutral densities (RND). Breeden and Litzenberger (1978) were the first to realize that the RND can be recovered from the option prices. Not only that RND can be used for pricing other, typically the less liquid and more exotic, derivatives written on the same underlying, but also researchers have found RND to be more informative and more responsive to news than the actual densities obtained from the prices of the underlying asset. But, the most important fact is that these empirically obtained RNDs are almost exclusively non-Gaussian. Chapter 11 covers a whole range of parametric and non-parametric methods for extracting RNDs. These techniques do not assume a specific model for the underlying asset. This is the reason why we called this chapter the "non-structural" approach to option pricing.

In the last chapter, Chapter 12, we put extensions of the BSM model into what we call "structural" option pricing models: structural in the sense that we now have a specific dynamic for the underlying asset price and sometimes a specific dynamic for the volatility also. This chapter is the most mathematically demanding and would require the support of Chapters 15, 16, and 17 for the mathematically less inclined readers. But this chapter also truly reflects the non-Gaussian nature of the underlying asset distributions in that jumps of "all shapes and sizes" are permitted at both return and volatility levels. At the time of writing, this is the cutting edge of option pricing as we know it!