

CHAPTER 1


Introduction

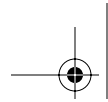
Most of the concepts in theoretical and empirical finance that have been developed over the last 50 years rest upon the assumption that the return or price distribution for financial assets follows a normal distribution. Yet, with rare exception, studies that have investigated the validity of this assumption since the 1960s fail to find support for the normal distribution—or *Gaussian* distribution as it is also called. Moreover, there is ample empirical evidence that many, if not most, financial return series are heavy-tailed and possibly skewed.

The “tails” of the distribution are where the extreme values occur. Empirical distributions for stock prices and returns have found that the extreme values are more likely than would be predicted by the normal distribution. This means that, between periods where the market exhibits relatively modest changes in prices and returns, there will be periods where there are changes that, are much higher (i.e., crashes and booms) than predicted by the normal distribution. This is not only of concern to financial theorists, but also to practitioners who are, in view of the frequency of sharp market down turns in the equity markets, troubled by the “... compelling evidence that something is rotten in the foundation of the statistical edifice ...” used, for example, to produce probability estimates for financial risk assessment.¹ Heavy or fat tails can help explain larger price fluctuations for stocks over short time periods than can be explained by changes in fundamental economic variables as observed by Robert Shiller (1981).

Mathematical models of the stock market developed through the joint efforts of economists and physicists have provided support for price and return distributions with heavy tails. This has been done by

¹ Hope 1999, p. 16.





modeling the interaction of market agents.² While these mathematical models by their nature are a gross simplification of real-world financial markets, they provide sufficient structure to analyze return distributions. Computer simulations of these models have been found to generate fat tails and other statistical characteristics that have been observed in real-world financial markets.

The first fundamental attack on the assumption that price or return distribution are not normally distributed was in the 1960s by Benoit Mandelbrot (1963). He strongly rejected normality as a distributional model for asset returns. Examining various time series on commodity returns and interest rates, Mandelbrot conjectured that financial returns are more appropriately described by a nonnormal stable distribution. To distinguish between Gaussian and non-Gaussian stable distributions, the latter are often referred to as “stable Paretian” distributions³ or “Lévy stable” distributions.⁴ His early investigations on asset returns were carried further by Eugene Fama (1965a, 1965b), among others, and led to a consolidation of the hypothesis that asset returns can be better described as a stable Paretian distribution.


There was obviously considerable concern in the finance profession by the findings of Mandelbrot and Fama. Shortly after the publication of the Mandelbrot paper, Paul Cootner (1964) expressed his concern regarding the implications of those findings for the statistical tests that had been published in prominent scholarly journals in economics and finance. He warned that:

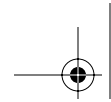
Almost without exception, past econometric work is meaningless. Surely, before consigning centuries of work to the ash pile, we should like to have some assurance that all our work is truly useless. If we have permitted ourselves to be fooled for as long as this

² Probably the most well-known model is the Santa Fe Stock Market Model (see Arthur et al., 1997). There are others. Bak, Paczuski, and Shubik (1996) and Lux (1999) analyze the interaction between two categories of market agents: “rational investors” and “noise traders.” Rational agents act on fundamental information in order to analyze risk-return opportunities and then act to optimize their utility function. “Noise” traders are market agents whose behavior is governed only by their analysis of market dynamics. Their choice at which to transact (buy or sell) may imitate the choice of other market agents. Cont and Bouchaud (2000) develop a model based on herding or crowd behavior that has been observed in financial markets.

³ The reason for this name is to emphasize the fact that the tails of the non-Gaussian stable distribution have Pareto power-type decay.

⁴ This name honors Paul Lévy for his seminal work introducing and characterizing the class of non-Gaussian stable distributions.





into believing that the Gaussian assumption is a workable one, is it not possible that the Paretian revolution is similarly illusory? (Cootner 1964, 337)



While that evidence has been supplied in numerous studies, the “normality” assumption remains the cornerstone of many leading theories used in finance.

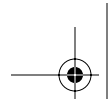
There was also concern at a theoretical level because one feature of the stable Paretian distribution is that there is an *infinite* variance. The variance of a return distribution for a highly diversified portfolio was just beginning to be accepted as the appropriate measure of risk in finance. It was one of the only two parameters needed in the theory of portfolio selection that had been developed by Harry Markowitz (1952, 1959). Moreover, the key feature of the framework developed by Markowitz, commonly referred to as “mean-variance analysis,” is the notion of how to diversification benefits investors. The underlying principle is that the risk (as measured by the variance) of a portfolio of returns consisting of stocks whose returns did not have perfect positive correlation would be less than the weighted average of the risk of the individual stocks comprising the portfolio. This quantification of diversification became known as “Markowitz diversification.”

Fama (1965c) revisited the notion of diversification if stock returns followed a stable Paretian distribution rather than a normal distribution. As we will see in Chapter 7, there are four parameters that describe a stable Paretian distribution. One of those parameters, the characteristic exponent of the distribution (also called the “tail index”), is critical in analyzing the benefits of diversification—reducing the dispersion of stock returns as the number of holdings increases. Fama derived the boundaries for the parameter so that an increase in the number of holdings in a portfolio provides the benefit of diversification. However, if the parameter was not within the narrow range he derived, increasing holdings could result in a greater dispersion of stock returns!⁵

As one would expect in the development of ideas in any field, defendants of the prevailing theories went on the offensive. One attack on the stable Paretian distribution was that there is no closed-form solution to obtain the necessary information about the distribution—probability density, distribution functions, and quantile, concepts of a probability distribution that we will describe in Chapter 3. While this may have been a valid criticism at one time, advances in computational finance make it fairly straightforward to fit observed returns to determine the parameters of a stable Paretian distribution. Thus, this criticism is no longer valid.

⁵ For a more general derivation, see Chapter 9 in Rachev and Mittnik (2000).





The major attack in the 1970s and 1980s centered around the claim that while the empirical evidence does not support the normal distribution, it is also not consistent with the stable Paretian distribution. For example, it was observed that asset return distributions are not as heavy tailed as the stable Paretian distribution would predict and furthermore they did not remain constant under temporal aggregation.⁶ That is, while there was no disagreement that the distribution of returns for assets were found to have heavier tails relative to the normal distribution, it was thinner than a stable Paretian distribution. Studies that came to such conclusions were typically based on a statistical test of the tail of the empirical distributions.⁷ However, the test that has been typically used to estimate the tails is highly unreliable for testing whether a distribution follows a stable Paretian distribution because sample sizes in excess of 100,000 are required to obtain reasonably accurate estimates. In other words, even if we were generating a small sample of a true stable Paretian distribution and then estimate the tail thickness with the estimator that has been used, we would most probably generate tail thickness estimates which contradict the stable assumption. There are other technical problems with these studies.⁸

Partly in response to these empirical “inconsistencies,” various alternatives to the stable Paretian distribution that had a finite variance were proposed in the literature. One alternative is the Student- t distribution, a distribution that under certain conditions not only has a finite variance⁹ but also allows for tails with more observations than the normal distribution. Battberg and Gonedes (1974) presented evidence supporting the Student t -distribution. Yet another distribution that has been proposed is a finite mixture of normal distributions. Kon (1984) found that this alternative explains daily returns for stocks better than the Student- t distribution.



A major drawback of all these alternative models is their lack of stability. As has been stressed by Mandelbrot and argued by Rachev and Mitnik (2000), among others, the stability property is highly desirable for asset returns. This is particularly evident in the context of portfolio analysis and risk management. Only for stable (which includes the Gaussian as a special case) distributed returns of independent assets does one obtain the property that the linear combination of the returns (portfolio returns) follow again a stable distribution. The independence assumption

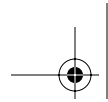
⁶ See Officer (1972), Akgiray and Booth (1988), and Akgiray and Lamoureux (1989).

⁷ The procedure for estimating the tail that has been used is the Hill estimator.

⁸ See Rachev and Mitnik (2000).

⁹ Specifically, the degrees of freedom must be greater than 2. With 2 degrees of freedom or less, a Student- t distribution has infinite variance as well.





of the returns can be replaced by assuming that the returns are jointly stable distributed (similar to the multivariate Gaussian case). Then again any portfolio return has a stable distribution, while the returns of the assets in the portfolios are jointly dependent with multivariate stable distribution. While the Gaussian distribution shares this feature, it is only one particular member of a large and flexible class of distributions, which also allows for skewness and heavy-tailedness.


This stability feature of the stable Paretian distribution that is not shared by other non-Gaussian distributions allows the generalization of the cornerstone normal-based financial theories and, thus, to build a coherent and more general framework for financial modeling. The generalizations are only possible because of specific probabilistic properties that are unique to stable (Gaussian and non-Gaussian) distributions, namely, the stability property, the Central Limit Theorem, and the Invariance Principle that we will describe in Chapter 7.

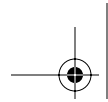
ORGANIZATION OF THE BOOK

The book is divided into five parts. Part One of the book includes six chapters that provide an introduction to the essential elements of probability theory and statistics for understanding the analysis of financial times series, risk management, and option pricing. In Chapters 2, 3, and 4 we explain how a probability distribution is used to describe the potential outcomes of a random variable, the general properties of probability distributions (including statistical moments), and the different types of probability distributions. In Chapter 3 we look at the normal probability distribution and its appeal.

In Chapter 5 we move from the probability distribution of a single random variable to that of multiple random variables, introducing the concept of a joint probability distribution, marginal probability distribution, and correlation and covariance that is commonly used to measure how random variables move together. We also discuss the multivariate normal distribution and a special class of distributions, the elliptical distribution. The limitations of correlation as a measure of the dependence between two random variables and how that limitation can be overcome by using copulas is provided in Chapter 6.

The stable distribution is the focus of Chapter 7. We explain the properties of stable distributions and considerations in the application of the distribution. We conclude the chapter with a brief introduction to smoothly truncated stable distributions that have been suggested for various applications in finance. In Chapter 8, we explain methodologies







for testing whether the probability distribution for a time series of returns for a particular asset follows a specific distribution and then present methodologies to fit a stable distribution from an empirical distribution.

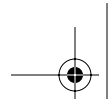
The two chapters in Part Two of the book cover stochastic processes. The theory of stochastic processes in *discrete time* is an important tool when examining the characteristics of financial time series data. An introduction to this tool is provided in Chapter 9. One of the simplest time series models is provided by the linear models of the *autoregressive moving average* (ARMA). However, when the focus is on modeling financial return data, it is sometimes necessary to incorporate time-varying conditional volatility. Statistical models for doing so are the *autoregressive conditional heteroskedasticity* (ARCH) model and *generalized ARCH* (GARCH) model and they are described in the chapter. Our approach in Chapter 9 is to motivate the reader as to why an understanding of this theory is important, rather than set forth a rigorous analytical presentation.

In some applications it might be more useful if we had the opportunity to model stochastic phenomena on a continuous-time scale rather than in discrete time. An example is the valuation of options (Black-Scholes option pricing model). The tool used in this case is continuous-time stochastic processes. In Chapter 10, we describe this tool and the most prominent representatives of the class of continuous-time stochastic processes: Brownian motion, Geometric Brownian motion, and the Poisson process.

Part Three provides the first of the three applications to finance. Application to portfolio selection is covered in this part, beginning in Chapter 11 with a description of recent empirical evidence on the return distribution for common stock and bonds that supports the stable Paretian hypothesis and clearly refutes the normal (Gaussian) hypothesis. Some desirable features of investment risk measures, the limitations of using the most popular measure of risk in finance and risk management, the variance, and a discussion of alternative risk measures are covered in Chapter 12. Also in that chapter, we describe two disjointed categories of risk measures, dispersion measures and safety risk measures, and review some of the most well known of each along with their properties.

There are two basic approaches to the problem of portfolio selection under uncertainty—one based on utility theory and the other based on reward-risk analysis. The former approach offers a mathematically rigorous treatment of the portfolio selection problem but appears sometimes detached from the world because it requires that asset managers specify their utility function and choose a distributional assumption for the returns. The latter approach is one that is more practical to implement. According to this approach, portfolio choice is made with respect






to two criteria—the expected portfolio return and portfolio risk—with the preferred portfolio being the one that has higher expected return and lower risk. The most popular reward-risk measure is the Sharpe ratio (see Sharpe 1966). In Chapter 13, we describe some new reward-risk measures that take into account the observed phenomena that assets returns distributions are fat tailed and skewed.


Applications to the management of market, credit, and operational risk are the subject of Part 4. Chapter 14 which covers market risk management begins with a review of the adoption of Value at Risk (VaR) by bank regulators for determining risk-based capital requirements and various methodologies for measuring VaR. We then discuss the stable VaR approach and present empirical evidence comparing VaR modeling based on the normal distribution with that of the stable Paretian distribution. We conclude with an explanation of an alternative market risk measure to VaR, Expected Tail Loss (or Conditional VaR) and the advantage of using this risk measure in portfolio optimization.

Credit risk management is the subject of Chapter 15, where we provide a description of credit risk (which consists of credit default risk, credit spread risk, and downgrade risk), an overview of the credit risk framework for banks as set forth in the Basel Accord II, credit risk models (structural models and reduced form models), and commercially available credit risk tools. We also present a framework for integrating market and credit risk management.

Our coverage of operational risk in Chapter 16 starts with a discussion of the distributions suggested by the Basel Committee for Regulatory Supervision for measuring exposure to operational risk. We then present evidence against the measure suggested by regulators for measuring exposure to operational risk by showing that the stable Paretian distribution may provide the best fit to the frequency and severity data.

Option pricing models depend on the assumption regarding the distribution of returns. In the three chapters in Part 5 we look at the most popular model for pricing options, the Black-Scholes model and how it can be extended. Chapter 17 covers the basic features of options, how options can be valued using the binomial model, and how one can obtain a continuous-time option pricing model by iteratively refining the binomial model. In Chapter 18 we then introduce the most popular continuous-time model for option valuation, the Black-Scholes model, looking at the assumptions and their importance. In Chapter 19 we look at several topics related to option pricing: the smile effect, continuous-time generalizations of the geometric Brownian motion for option pricing (stochastic volatility models and so-called “local volatility models”), models with jumps, models with heavy-tailed returns, and generalization of the discrete time model for pricing options.





REFERENCES

Akgiray, A., and G. G. Booth. 1988. "The Stable-Law Model of Stock Returns." *Journal of Business and Economic Statistics* 6: 51–57.

Akgiray, V. and C. G. Lamoureux. 1989. "Estimation of Stable-Law Parameters: A Comparative Study." *Journal of Business and Economic Statistics* 7: 85–93

Arthur, W. B., J. H. Hollan, B. LeBaron, R. Palmer, and P. Tayler. 1997. "Asset Pricing under Endogeneous Expectations in an Artificial Stock Market." In W.B. Arthur, S. Durlauf, and D. Lane (eds.), *The Economy as an Evolving Complex System, II, Vol. XXVII* of SFI Studies in the Sciences of Complexit. Redwood, CA: Addison-Wesley.

Bak, P., M. Paczuski, and M. Shubik. 1996. "Price Variations in a Stock Market with Many Agents." No 1132, Cowles Foundation Discussion Papers, Cowles Foundation, Yale University.

Blattberg, R. C. and N. J. Gonedes. 1974. "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices." *Journal of Business* 47: 244–280.

Cont, R. and J.P. Bouchaud. 2000. "Herd Behavior and Aggregate Fluctuations in Financial Markets." *Macroeconomic Dynamics* 4: 170–196.

Cootner, P.H. 1964. *The Random Character of Stock Market Prices*. Cambridge, MA: The M.I.T. Press.

Fama, E. 1963a. *The Distribution of Daily Differences of Stock Prices: A Test of Mandelbrot's Stable Paretian Hypothesis*. Doctoral Dissertation, Graduate School of Business, University of Chicago.


Fama, E. F. 1963b. "Mandelbrot and the Stable Paretian Hypothesis." *Journal of Business* 36: 420–429.

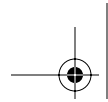
Fama, E. F. 1965c. "Portfolio Analysis in a Stable Paretian Market." *Management Science* 11: 404–419.

Hoppe, R. 1999. "It's Time We Buried Value-at-Risk." *Risk Professional*, Issue 1/5, July/August: 16.

Kon, S., 1984. "Models of Stock Returns—A Comparison." *Journal of Finance* 39: 147–165.

Lux, T. 1998. "The Socio-Economic Dynamics of Speculative Markets." *Journal of Economic Behavior and Organization* 33: 143–165.





Mandelbrot, B. B. 1963. "The Variation of Certain Speculative Prices." *Journal of Business* 36: 394–419.

Markowitz, H. M. 1952. "Portfolio Selection." *Journal of Finance* 7: 77–91.

Markowitz, H. M. 1959. *Portfolio Selection: Efficient Diversification of Investment*. New York: John Wiley.

Officer, R. R. 1972. "The Distribution of Stock Returns." *Journal of the American Statistical Association* 76: 807–812.

Rachev, S. T., and S. Mittnik. 2000. *Stable Paretian Models in Finance*. John Wiley & Sons, Chichester.

Sharpe, W. F. 1966. "Mutual Funds Performance." *Journal of Business*, January: 119–138.

Shiller, R. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends." *American Economic Review* 71: 421–436.

