CHAPTER 9

The Role of Mathematics in the Financial Markets

WALTER SCHACHERMAYER

The financial markets have not only experienced a stormy development in past years, but the methods used for evaluating the quality and the risk of an investment have changed too. While even 30 years ago the tools available to a successful investor were only his "flair", besides legal and economic knowledge, today there is a multitude of quantitative methods available. The concepts of "arbitrage" and the "Black-Scholes formula" now play a central role in the valuation and hedging of options. The significance of this formula was recognized in 1997 by the award of the Nobel Prize for Economics to Robert C. Merton and Myron Scholes: Fisher Black, who had died in 1995, was also honored, but Nobel Prizes are not awarded posthumously.

In this short article I will attempt to present a generally understandable survey of stochastic mathematical finance, the theory that stands behind these methods. In particular, I would like to indicate the strengths, but also the weaknesses, of modeling financial markets by stochastic (i.e., chance dependent) processes.

Let us begin our discussion with the older sister of mathematical finance, classical insurance mathematics. Ever since Sir Edmond Halley, the pupil and friend of Isaac Newton, known above all for his eponymous comet, published a "mortality table" in the year 1693, actuaries have employed the same method for setting insurance premiums: the "equivalence principle".

We illustrate this by an extremely simple example.

Suppose that a 40-year-old woman takes out a one-year term insurance: In the case of her death in the course of the following year, her heirs will be paid the insured sum S, for example $S =$ € 100,000, at the end of the year; otherwise nothing is paid out. How should an insurance company calculate the premium for such a contract?

Here is where probability theory enters the story, in a very simple way. The death, or the survival, of the woman is modeled as a *random variable*, just like throwing a coin. Now, however, the chances are not $50:50$: instead they must be derived from the fact that 40-year-old women have a known probability, which we denote by q_{40}, of dying in the course of the following year. A mortality table (awful word!) is simply a listing of the values q_y and q_x, where y (resp. x) runs through the possible ages (e.g., $0, 1, 2, \ldots, 110$) for women (resp. men).

The premium for the contract is determined as the *expected value* of the benefit to the insurance company. Thus, in our example,

$$\text{premium} = q_{40} \cdot \text{insured sum}\,. \tag{1}$$

If we further assume that q_{40} can be taken to be 0.0012 (in accordance with a modern mortality table) we find

$$\text{premium} = 0.0012 \cdot € \, 100,000 = € \, 120. \tag{2}$$

A trifle that we have left out: although the premium is paid at the beginning of the year, the benefit of the insurance will not be paid before the end of the year, and the interest effect must be taken into account. This is done by determining an actuarial interest rate (for example $i = 4\,\%$) and discounting the premium appropriately: here

$$\text{premium} = € \, \frac{120}{1.04} \approx € \, 115.38. \tag{3}$$

Simple as this process may seem, this modus operandi is, *in nuce*, exactly what actuaries have done for centuries: they calculate the premium as the discounted expected value of the insurance benefit.

Perhaps you would now like to object that the insurance company also has costs (sales, administration, etc.), and that these must be considered too. This is of course correct, and these costs are incorporated as an addition to the premium.

But when we discount the cost, what really is the mathematical justification for using the *expected value* in assessing the premium? The reason lies in the *Law of Large Numbers*, which states the following: So long as the assumed probability q_{40} does in reality model the mortality of 40-year-old women correctly, the insurance company will on average neither win nor lose if it concludes "many" independent contracts of this kind. How to interpret "many" can be quantified in a mathematically precise way from the Limit Theorems.

Financial mathematics, more precisely stochastic mathematical finance, is—at least at the first glance—fundamentally different from insurance mathematics. The reasoning based on the Law of Large Numbers is replaced by the concept of *arbitrage*.

In order to motivate this concept we again consider a very simple example: If in Frankfurt the dollar is exchanged at \$ 1.05 per euro, then it will be exchanged at (almost) the same rate in New York: For if the rate there were, for example, \$ 1.0499 per euro, then *arbitrageurs* would immediately simultaneously exchange euros against dollars in New York and dollars against euros in Frankfurt, thereby realizing a risk free profit. If, instead, the rate in New York were \$ 1.0501 per euro they would transact in the opposite direction, again making a risk free profit. With an exchanged volume of € 10 M, for example, a relatively moderate sum in global foreign exchange, where transactions can be made within a second, the arbitrage profit would be about € 950.

Whether one finds this smooth functioning of the international financial market good or bad, is a different question, which we shall not analyze: If the suggestion of imposing a global turnover tax (of the order of a fraction of a tenth of a percent) on financial transactions were to be adopted, as proposed by J. Tobin (Economics Nobel Prize 1981), this would introduce "friction", and the situation would change quickly.

But back to the concept of arbitrage: You may perhaps argue that the activity of a shoe shop owner who buys a pair of shoes for € 30 and later sells them for € 60 is akin to arbitrage; the difference, however, is that the shoe seller's input is in finding customers, contacting his suppliers, storing the shoes, etc. In contrast, the exchanges are organized so that the prices are transparent to all market participants and large volumes can be transacted at little cost. In our foreign exchange example the full € 950 will not remain as an arbitrage profit to the arbitrageurs because of the transaction costs, but for the "big players" the relative significance of the transaction cost is extremely low.

We can now define an essential pillar of the theory, as employed by Black, Scholes and Merton. They took the *no-arbitrage principle* as fundamental in their mathematical modeling of the financial markets, ignoring the transaction costs (as a first approximation): There should be no arbitrage opportunities in the mathematical model of a financial market. The plausible argument behind this is: As soon as there are arbitrage opportunities, no matter how small, then, as in the preceding example, arbitrageurs will reduce them quickly to zero, precisely by exploiting those arbitrage opportunities. In liquid financial markets, e.g., foreign exchange markets but also in large share and commodity markets, reality comes very close to this mathematical postulate.

Next let us elucidate the no-arbitrage principle with a somewhat less simple example than the location arbitrage sketched above, namely with the *forward rate* of a currency. I can buy a *forward contract* giving me *the right and the obligation* to exchange a certain amount, e.g., € 10,000 into dollars at a rate agreed today at a fixed time point, e.g., in a year.

For such a contract to be possible there must be other market participants ready to make such a contract in the other direction, i.e., to buy the right and the obligation to exchange the same amount from dollars to euros at the agreed rate at the same time.

The *forward rate* for the dollar is the exchange rate at which the players in the financial markets are prepared to make such a contract today.

Can one say something intelligent about the level of the forward price, other than the lapidary assertion that this price will swing according to offer and demand on the foreign exchange futures markets? The answer is yes, and it is amazingly simple.

Suppose for simplicity that today the interest level for a one year "risk free" term deposit (that is—to a first approximation—government bonds with one year remaining to run) is equally high in euros and dollars. I claim that then the forward rate for the euros in dollars must coincide with today's ("spot" or "cash") rate of the euros in dollars. Suppose, for example, that the forward rate for the euros is higher than the cash rate, e.g., € 1.06 versus $ 1.05. In this case an arbitrageur will today borrow dollars for a year, exchange them into euros, deposit these euros for a year, and simultaneously make a *forward contract* to exchange the (compounded) sum back from euros into dollars after the year. Our assumption that the interest levels are the same in dollars and euros implies that the result of this combination of transactions must cancel out, so long as the cash rate is equal to the forward rate. If the forward rate, though, were higher than the cash rate the difference would remain as a profit to the arbitrageur! What is remarkable about this arbitrage transaction is that this profit is achieved without any net investment of capital and is completely risk free: the profit is completely independent of whether in the subsequent year the rate of the euro against the dollar rises, falls, or stays the same. It is the nature of an arbitrage profit that it arises from a combination of transactions, each of which is individually risky—maybe very risky—but that the countervailing risks mutually cancel.

The alert reader may remark that we have used the same level of interest both for a loan (in dollars) and for a deposit (in euros); on the other hand we all know that one has to pay higher interest on a loan than one receives for a risk free deposit (in the same currency), since otherwise one could make an obvious arbitrage. With this argument as with the transaction cost: for small investors this difference is very significant, but the "big players", however, can in the same conditions go "long" or "short", i.e., in the context of our example, can deposit or lend money at the same interest rate.

As a next step we will discard the simplifying assumption that the interest rates (for a one year risk free deposit) are the same in dollars and euros: suppose, for example, that the corresponding dollar-interest is 4 % while the euro-interest is only 3 %. If one again thinks through the argument developed above one sees immediately that it can be applied in this situation too: the only difference is that the ratio of the forward rate of the euro into dollars to the cash rate is no longer 1 : 1 but now must be 1.04 : 1.03.

We here recommend the skeptical reader to take up the financial part of the daily paper and check empirically that these considerations are not just arid theorizing. The reader will be able to convince herself that, as discussed, the level of the forward rates between two currencies really depends on the ratio of the interest levels of the currencies—and only on this. And this is not because of supervision by a regulatory agency or something similar (as was not the case in the simple example of location arbitrage), but because worldwide market participants exploit any arbitrage opportunities immediately, and so bring them to vanish (or, more precisely, reduce to such a small amount that even arbitrageurs with very small relative transaction costs can no longer profit from them). This certainly goes for currencies for which the cash and futures markets show a high liquidity (i.e., high transaction volume and small transaction costs). The euro versus the dollar is of course a prime example.

Up to now our thoughts on arbitrage have required only a very elementary mathematical discussion. This changes abruptly on turning to other contracts handled on the futures exchanges: an option (more precisely, a European call option) certifies the *right, but not the obligation*, to buy a certain quantity of an *underlying stock*, e.g., foreign currency, shares, etc., at a certain *expiration time* and at a certain *strike price*. To illustrate the economic sense of such contracts: In the preceding foreign exchange example there could be good reasons for investors to ensure themselves the right to exchange euros into dollars in a year's time, to protect themselves against an increasing dollar rate, but not to bind themselves to this transaction if it would lead to a loss in the case of a falling dollar rate.

Naturally one can no longer—as with a forward contract—purchase such an option at zero cost: the buyer has to pay a price to purchase the option.

Again we have the question of whether we can say something intelligent about this price or whether we must simply refer to market forces. And again the answer is yes; though the situation is no longer as simple as for the forward rate discussed above.

Above it was enough to consider "buy-and-hold strategies": When we review our argument on the determination of the unique arbitrage free forward rate, we see that only four transactions (a loan, a deposit, a currency exchange at the cash rate, and a forward contract) were necessary. In order to cream off the arbitrage profit (if the forward rate is not at the level postulated by the theory), the arbitrageur can close these transactions today, simply wait a year, and then chalk up a risk free profit according to the contractual right and obligation (buy-and-hold).

These elementary strategies are not adequate for ferreting out arbitrage opportunities in the context of options. It is relatively easy to see (and also to prove mathematically) that one cannot deduce any nontrivial statements on the price of an option from arbitrage arguments concerning buy-and-hold strategies.

Now the market permits not only buy-and-hold, that is, static trading strategies, but it also allows one to trade dynamically. In mathematical modeling one says "in continuous time"; we call a trading strategy *dynamic* when it permits buying and selling at any time, but of course using only the information available at the time of trading (if I had access to the exchange reports of tomorrow it would not be difficult to make an arbitrage profit). The mathematics has, with the theory of stochastic processes that were developed mainly for applications in the natural sciences, an outstanding instrument at hand for modeling the concept of a *dynamic trading strategy* precisely. (Keywords: filtrations, predictable processes, etc.)

The greater the possibilities of trading on the financial markets, the more the opportunities of compensating countervailing risks, drawing on the no-arbitrage argument for valuations.

In order to model the possible price developments of an option on the *underlying* (the "stock" or "asset"; e.g., shares in Company XYZ), we need to make assumptions about the price process $(S_t)_{0 \leq t \leq T}$. For each t in the interval $[0, T]$ we denote by S_t the price of the share at time t: The quantity T denotes the expiration time of the option (e.g., in a year), and we denote today's date by 0. We know today's price S_0, but, since we cannot see into the future, we model the variables S_t, for $0 \leq t \leq T$, as *random variables*. To specify the process $(S_t)_{0 \leq t \leq T}$, we need to impose further assumptions on the probability distributions of the random variable S_t.

This theme is by no means new. Louis Bachelier, in his 1900 dissertation under the distinguished mathematician Henri Poincaré, had already proposed a model for the price-process $(S_t)_{0 \leq t \leq T}$ of a share, his motivation being to derive a *formula for the valuation of options*. He modeled the price of the share as a random (or *stochastic*) process: whether the price of our share will rise or fall tomorrow is to be described in a way similar to the throwing of a coin or the spin of a roulette wheel. Bachelier had an almost mystic belief that a *probability law* determined the events on the exchange:

> Si, à l'égard de plusieurs questions traitées dans cette étude, j'ai comparé les résultats de l'observation à ceux de la théorie, ce n'était pas pour vérifier des formules établies par les méthodes mathématiques, mais pour montrer seulement que le marché, à son insu, obéit à une loi qui le domine : la loi de la probabilité.[1]

[1] If, with respect to the various questions treated in this study, I have compared the results of observation with those of the theory, it has not been to verify the formulae established by mathematical methods, but only to demonstrate that the market, unaware, obeys the one law that dominates it: the law of probability.

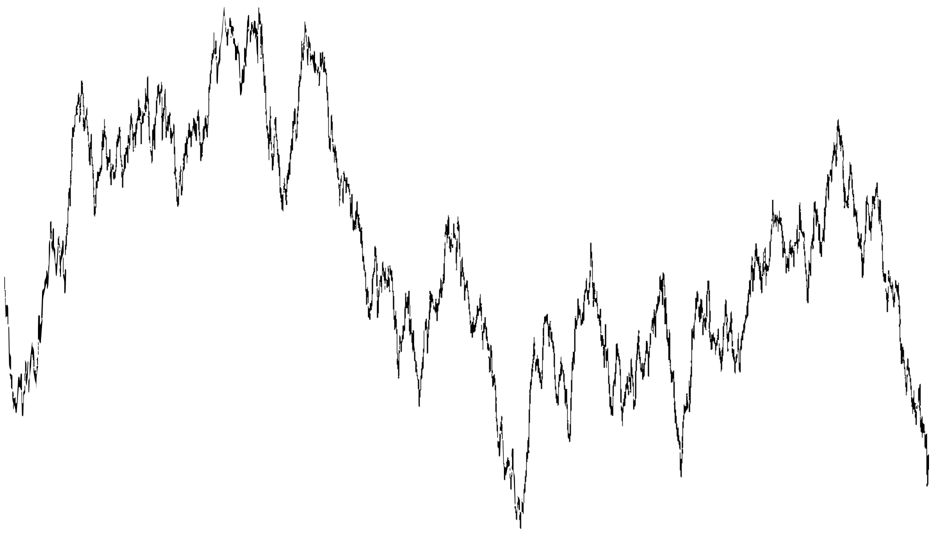

FIGURE 1. Path of a Brownian motion

As a concrete model he proposed the process $(S_t)_{0\leq t\leq T}$ known nowadays as *Brownian motion*: The increment $S_t - S_u$ of the price between two times $u < t$ is assumed to be normally distributed (that is, according to the famous *Gaussian bell curve*), with mean value 0 and variance proportional to the length of the interval $[u, t]$. Further, the increments over disjoint time intervals are assumed to be mutually independent.

For a fixed random element ω, i.e., for ω in the underlying probability space $(\Omega, \mathcal{F}, P)$, one has a path $(S_t(\omega))_{0\leq t\leq T}$; a typical simulated path is sketched in Figure 1.

It redounds to the pride of mathematical finance that Bachelier was thus the first to formulate the mathematical model of a Brownian motion. He was five years ahead of Einstein and Smoluchowski, who introduced this model into physics in 1905 to describe the behavior of gas molecules. The name Brownian motion derives from the fact that in 1826 the botanist Robert Brown detected a completely erratic behavior—similar to the simulated path in Figure 1—while observing particles in the microscope (though he did not attempt to model this behavior mathematically).

After formalizing his model, Bachelier was able to approach the real aim of his work, namely the valuation of an option on a share whose price follows a Brownian motion $(S_t)_{0\leq t\leq T}$. If we fix the exercise time T and the strike price K, it is easy to specify the value C_T of the option *at the time* T: C_T is the larger of the two numbers 0 and $S_T - K$.

Indeed, if the price S_T of the underlying share is larger than K, then the value of the option is the difference $S_T - K$, since the owner of the option can buy a share at price K and then sell it immediately for S_T. But if S_T is smaller than K, then the option is quite worthless.

Thus we see that we can write the value C_T of the option at time T as a simple function of the random variable S_T. We do not know today (i.e., at time $t = 0$) the value that S_T will have at time $t = T$, but only the *probability distribution* of S_T. We illustrate the difference again with a simple example: When I roll a die I do not know the result of the throw in advance, but I assume—and with a fair die I have good reason to—that the probability distribution of the result is that each of the six possible digits will be thrown with the same chance (thus with probability $\frac{1}{6}$). Similarly, I do not know the value S_T of the share today, but I claim to know the probability distribution of S_T.

How can we compute the present value C_0 of the option? Bachelier does just what actuaries have done for ages: He takes the *expected value* of the value of C_T,

$$C_0 = E[C_T]. \tag{4}$$

This is an expression that can be evaluated easily, i.e., can be expressed in terms of a *formula*, since—according to our model assumption—we know the distribution S_T, namely a normal distribution with mean value S_0 and variance T.

One may object that Bachelier neglected the interest effect that we had to take into account when discussing life insurance. This objection is not a very strong one: Bachelier ignored the interest effect since he was interested in the valuation of options with a relatively short exercise time (T of the order of a few months), and interest rates were then low—at the time of the past *fin-de-siècle* too the interest rates were low and the share prices high! Of course, if one does want to consider the interest effect, there is no problem in building a discount factor into the formula:

$$C_0 = e^{-rT} E[C_T], \tag{5}$$

where r denotes the risk free interest. The decisive point is, however, a different one: The motivation for using the expected value is the Law of Large Numbers, which, from an economic point of view, is much less convincing than the no-arbitrage principle. We do not yet find the idea of a connection between these two approaches in Bachelier.

Bachelier's work sadly did not enjoy the attention it deserved. The economists, on the one hand, ignored it completely, and only 65 years later did the eminent economist Paul Samuelson (Economics Nobel Prize 1970) take this theme up again. But the mathematicians also took little notice of it. However, his work was not completely forgotten in the mathematical community; for example, it was cited in Kolmogorov's fundamental book on probability theory of 1932.

The essential breakthrough on the question of option valuation came first in 1973 with the work of Black and Scholes, and also of Merton. They took a slight variant of the model used by Bachelier for the share price: They assumed—as had Samuelson—that the *logarithm* $\ln(S_t)$ of the stock price process S_t would follow a Brownian motion with drift: i.e.,

$$\ln(S_t) = \ln(S_0) + \sigma W_t + \mu t, \tag{6}$$

where $\mu \in \mathbb{R}$, $\sigma > 0$ are appropriate normalization constants and W_t is a Brownian motion as defined by Bachelier.

Changing to the logarithm is a harmless step, and it corresponds to the difference between compound interest, when invested capital grows according to an exponential curve, and linear interest, when one neglects the effect of iterated interest. As is well known, the difference between these two approaches is not very significant over short time spans. It is similar for the difference between Bachelier's model of Brownian motion and the model (6), the *geometrical Brownian motion*, today often also referred to as the *Black-Scholes model*.

Black, Scholes and Merton were the next to break new ground, using a no-arbitrage argument for dynamical trading strategies. In essence this argument runs: *Suppose* that there really is a function $f(t, S)$ that determines the value of the option at each time point $0 \leq t \leq T$ in terms of the price S at time t of the *underlying*—this is the technical expression for the asset to whose value changes the option relates. Then one can differentiate this function $f(t, S)$ partially with respect to the variable S. Following practitioners' language, we call the quantity $\frac{\partial}{\partial S} f(t, S)$ the *Delta* of the option at time t at current stock-price S (for fixed t and S).

For the purpose of illustration suppose that this Delta has the value $\frac{1}{2}$ for some fixed t and S. This means that when the value S of the underlying increases by €1 the value of the option increases by about 50 cents (t remaining fixed). The "about" is to be understood in the sense of differential calculus: this ratio 2 : 1, of the variation of the underlying to that of the option, fits better as the price variations become small, and is exact "in the limit".

This relation has an important economic consequence: If—still with fixed t and S—we create a portfolio, going "long" with one unit of the underlying (buying it) and simultaneously going "short" with two units of the option (selling it), then this portfolio is risk free against (small) price variations in the underlying: profits on the underlying will be compensated by losses on the option, and vice versa.

The portfolio is risk free only "locally", i.e., as long as t and S vary only a little, but the idea of a dynamical trading strategy allows one to adjust the balance of the portfolio by buying and selling so as to fit to the current Delta.

Now comes the no-arbitrage argument: A risk free portfolio formed in this way must yield the same interest as a risk free deposit. For, if not, one could find trading strategies, as already discussed, that would make arbitrage profits possible.

We have thus found an economic connection between the dynamics of the value of the portfolio and the risk free interest, and can express it mathematically in the form of an equation: if we adopt the model assumption (6), we are led to a *partial differential equation* that can be solved explicitly. The solution can be stated in the shape of a formula, namely the famous Black-Scholes formula:

$$f(t, S) = SN(d_1) - Ke^{-r(T-t)}N(d_2), \tag{7}$$

$$\text{where } d_1 = \frac{\ln(S/K) + (r + \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}}, \tag{8}$$

$$d_2 = \frac{\ln(S/K) + (r - \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}}. \tag{9}$$

Here N denotes the distribution function of the standard normal distribution, S is the value of the underlying at time t, while K and T are the strike price and exercise time of the call option, r is the risk free interest rate, and $\sigma > 0$ is the *volatility*, i.e., the parameter for the influence of the random Brownian motion W in the underlying model (6).

The concrete shape of the formula is not so important, and I have given it here only so that the reader can see how it can be used explicitly with concrete numerical values.

The following is more important. Today's value $f(0, S_0)$ of the call option is the *only possible arbitrage free price*. And more: The derivation of the formula also leads explicitly to dynamical trading strategies, with which arbitrage profits can be made if the market price of the option deviates from this theoretical value.

Finally, a total surprise: The price $f(0, S_0)$ also emerges from the approach corresponding to the Bachelier formula (5),

$$f(0, S_0) = e^{-rT} E_Q[C_T], \tag{10}$$

where the expected value is not with respect to the original probability measure P, but with respect to a modified *risk neutral* probability measure Q. The term risk neutral comes from the fact that, on the basis of this modified probability distribution, the appreciation of the stock is on average the same as for a risk free deposit.

A more detailed rationale for this fundamental connection between the no-arbitrage argument on the one hand and the spectacular revival (10) of the good old insurance mathematical equivalence principle on the other hand would take us beyond the confines of this article. It is the theme of the *Fundamental Theorem of Asset Pricing* that was developed around 1980 in the works of M. Harrison, D. Kreps, and S. Pliska, later extended by numerous other authors. An exact formulation of this fundamental theorem in a general mathematically precise framework was first given by Freddy Delbaen and the author in 1994.

We shall here develop only a very intuitive approach to what really happens in changing from the original "true" probability measure P to the modified risk neutral probability measure Q. Let us cut back to the very simple example of a one year term insurance for a 40-year-old woman. You probably reacted with surprise and scepticism on being told that the insurance company calculates the premium by using the expected value; because then the insurance company will on average not earn anything from these contracts. This scepticism is quite justified, for an insurance company has to be profit oriented.

The solution to this riddle lies in the fact that two different probability distributions are involved here too: on the one hand there is the true probability q_{40} that a 40-year-old woman will die in the course of the year; this true probability can be estimated very reliably from mortalities in the past. But another, carefully chosen, probability, denoted by q_{40}^{mod}, is applied in calculating the premium. The (mortality) profit on the policy results from the difference in these two values.

The parallel to mathematical finance is now obvious, where it was equally essential to distinguish between the true measure P and the "modified" measure Q.

After this general discussion we pose the question of how well the Black-Scholes formula, and the hedging strategies derived from it, work in practice. This question essentially depends on whether the model of geometrical Brownian motion (6) describes reality correctly.

Let us cast a glance back to the real data from the financial market time series. Figure 2 presents the daily logarithmic returns, i.e., $\ln(S_{t+1}/S_t)$, of an Austrian share price index, where t runs through the trading days from April 1995 to

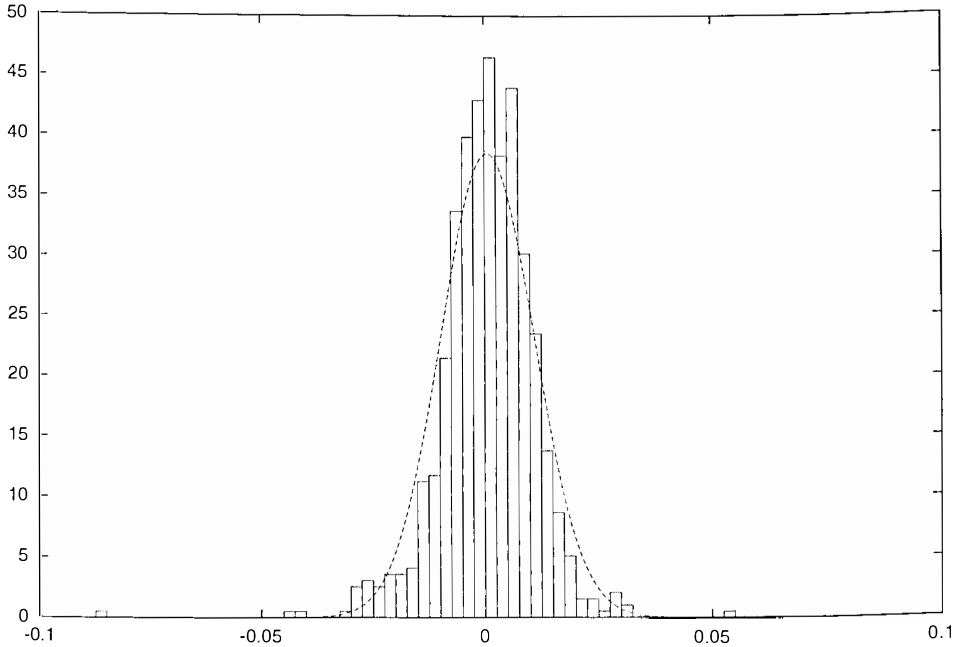

FIGURE 2. ATX Log-Returns (April 1995–June 1998): Comparison with the normal distribution

June 1998. If the assumptions of the Black-Scholes model are met, this random variable must be normally distributed, i.e., the empirical histogram must approximate the form of the normal distribution (shown dashed).

We see that the agreement is not so very good: the empirical histogram has too much mass near the mean value (in comparison with the theoretical normal distribution) and too little in the middle range. The most difficult problem for practical application is not so obvious to the eye, and concerns the "tails" of the distribution: the normal distribution underestimates the extreme events dramatically; and just these events, large fluctuations, are obviously of particular practical relevance.

These "stylized facts", as we observe in this example (i.e., in comparison to the normal distribution there is too much probability measure in the center and at the tails of the distribution, but too little in a middle range), reappear with remarkable persistence in such time series.

In place of an approximation by the normal distribution, Figure 3 shows the approximation of the same empirical histogram by one of a more general class of probability distributions, the hyperbolic distributions. The fit is considerably better and, although one cannot see it in this example with the naked eye, comprehensive empirical investigations confirm that the modeling of the extreme fluctuations by this more general class of distributions also agrees better with reality than does the normal distribution.

This provokes the question of why we do not replace the Black-Scholes model by a more general model that would describe reality better. Research is doing just this and, in increasing measure, practice too—and meanwhile progress is controversial. The situation immediately becomes essentially more complicated if

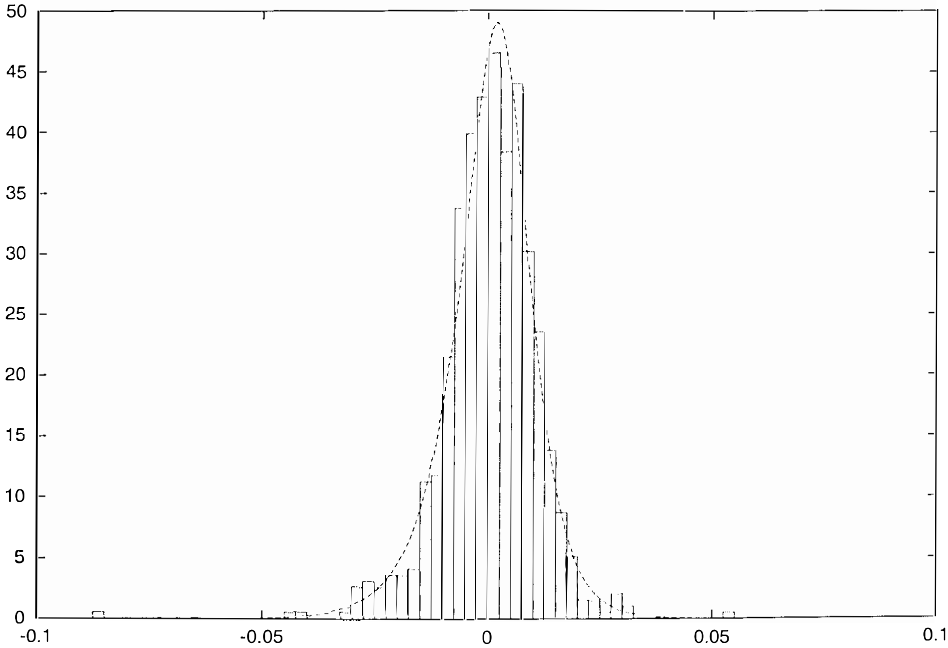

FIGURE 3. ATX Log-Returns (April 1995–June 1998): Comparison with the hyperbolic distribution

one steps out of the Black-Scholes model, since one can no longer derive unique prices and the corresponding trading strategies from pure no-arbitrage arguments. For this reason the Black-Scholes model continues to play a fundamental role for practitioners, although current research results are being adopted in practice with remarkable speed.

In this introductory presentation we cannot go into extensions of the Black-Scholes models, but only indicate a comprehensive list for further reading. I hope, however, to have conveyed the following message to the reader: For the practical application of the theory it is crucial to understand the chosen mathematical model and its assumptions thoroughly. This is particularly necessary for developing an understanding of in which aspects the model assumptions describe reality acceptably, and in which aspects this is not the case. This forms the basis for a critical awareness of the situations when the theory delivers valuable results, and in which situations great caution is called for.

References

[1] Bachelier, L. (1964) Théorie de la Spéculation. *Ann. Sci. Ecole Norm. Sup.* **17**, 21–86 (1900). English translation in: *The Random Character of Stock Market Prices* (P. Cootner, editor), MIT Press.

[2] Black, F., Scholes, M. (1973) The pricing of options and corporate liabilities. *Journal of Political Economy* **81**, 637–659.

[3] Delbaen, F., Schachermayer, W. (2006) The Mathematics of Arbitrage. *Springer Finance*, 371.

[4] Merton, R.C. (1973) The theory of rational option pricing. *Bell J. Econ. Manag. Sci.* **4**, 141–183.

[5] Samuelson, P.A. (1965) Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* **6**, 41–50.

[6] Schachermayer, W., Teichmann, J. (2005) How close are the Option Pricing Formulas of Bachelier and Black-Merton-Scholes? *Mathematical Finance* (1) **18**, 55–76.
[7] Shreve, S. (2004) Stochastic Calculus for Finance. *Springer Finance* (Vols. I and II).