CHAPTER 15

VOLATILITY SMILES

How close are the market prices of options to those predicted by Black–Scholes? Do traders really use Black–Scholes when determining a price for an option? Are the probability distributions of asset prices really lognormal? What research has been carried out to test the validity of the Black–Scholes formulas? In this chapter we answer these questions. We explain that traders do use the Black–Scholes model—but not in exactly the way that Black and Scholes originally intended. This is because they allow the volatility used to price an option to depend on its strike price and time to maturity.

A plot of the implied volatility of an option as a function of its strike price is known as a *volatility smile*. In this chapter we describe the volatility smiles that traders use in equity and foreign currency markets. We explain the relationship between a volatility smile and the risk-neutral probability distribution being assumed for the future asset price. We also discuss how option traders allow volatility to be a function of option maturity and how they use volatility matrices as pricing tools. The final part of the chapter summarizes some of the work researchers have carried out to test Black–Scholes.

15.1 PUT–CALL PARITY REVISITED

Put–call parity, which we explained in Chapter 8, provides a good starting point for understanding volatility smiles. It is an important relationship between the price, c, of a European call and the price, p, of a European put:

$$p + S_0e^{-qT} = c + Ke^{-rT} \tag{15.1}$$

The call and the put have the same strike price, K, and time to maturity, T. The variable S_0 is the price of the underlying asset today, r is the risk-free interest rate for maturity T, and q is the yield on the asset.

A key feature of the put–call parity relationship is that it is based on a relatively simple no-arbitrage argument. It does not require any assumption about the future probability distribution of the asset price. It is true both when the asset price distribution is lognormal and when it is not lognormal.

Suppose that, for a particular value of the volatility, p_{BS} and c_{BS} are the values of European put and call options calculated using the Black–Scholes model. Suppose further that p_{mkt} and c_{mkt} are the market values of these options. Because put–call parity holds for the Black–Scholes model, we must have

$$p_{\text{BS}} + S_0e^{-qT} = c_{\text{BS}} + Ke^{-rT}$$

Because it also holds for the market prices, we have

$$p_{\text{mkt}} + S_0e^{-qT} = c_{\text{mkt}} + Ke^{-rT}$$

Subtracting these two equations gives

$$p_{\text{BS}} - p_{\text{mkt}} = c_{\text{BS}} - c_{\text{mkt}} \tag{15.2}$$

This shows that the dollar pricing error when the Black–Scholes model is used to price a European put option should be exactly the same as the dollar pricing error when it is used to price a European call option with the same strike price and time to maturity.

Suppose that the implied volatility of the put option is 22%. This means that $p_{\text{BS}} = p_{\text{mkt}}$ when a volatility of 22% is used in the Black–Scholes model. From equation (15.2), it follows that $c_{\text{BS}} = c_{\text{mkt}}$ when this volatility is used. The implied volatility of the call is therefore also 22%. This argument shows that the implied volatility of a European call option is always the same as the implied volatility of a European put option when the two have the same strike price and maturity date. To put this another way, for a given strike price and maturity, the correct volatility to use in conjunction with the Black–Scholes model to price a European call should always be the same as that used to price a European put.

This is also approximately true for American options. It follows that when traders refer to the relationship between implied volatility and strike price, or to the relationship between implied volatility and maturity, they do not need to state whether they are talking about calls or puts. The relationship is the same for both.

Example 15.1 The value of the Australian dollar is \$0.60. The risk-free interest rate is 5% per annum in the United States and 10% per annum in Australia. The market price of a European call option on the Australian dollar with a maturity of one year and a strike price of \$0.59 is 0.0236. DerivaGem shows that the implied volatility of the call is 14.5%. For there to be no arbitrage, the put–call parity relationship in equation (15.1) must apply with q equal to the foreign risk-free rate. The price, p, of a European put option with a strike price of \$0.59 and maturity of one year therefore satisfies

$$p + 0.60e^{-0.10\times 1} = 0.0236 + 0.59e^{-0.05\times 1}$$

so that $p = 0.0419$. DerivaGem shows that, when the put has this price, its implied volatility is also 14.5%. This is what we expect from the analysis just given.

15.2 FOREIGN CURRENCY OPTIONS

The volatility smile used by traders to price foreign currency options has the general form shown in Figure 15.1. The volatility is relatively low for at-the-money options. It becomes progressively higher as an option moves either in the money or out of the money.

In Appendix 15A we show how to determine the risk-neutral probability distribution for an asset price at a future time from the volatility smile given by options maturing at that time. We refer to this as the *implied distribution*. The volatility smile in Figure 15.1 corresponds to the probability distribution shown by the solid line in Figure 15.2. A lognormal distribution with the

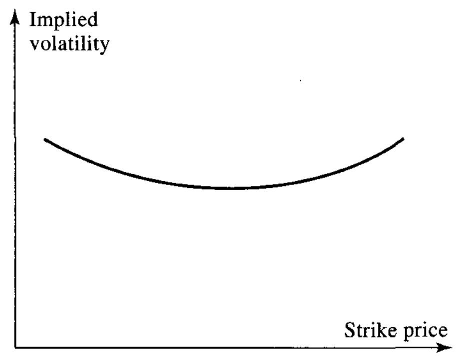

Figure 15.1 Volatility smile for foreign currency options

same mean and standard deviation as the implied distribution is shown by the dashed line in Figure 15.2. It can be seen that the implied distribution has heavier tails than the lognormal distribution.[1]

To see that Figure 15.1 and 15.2 are consistent with each other, consider first a deep-out-of-the-money call option with a high strike price of K_2. This option pays off only if the exchange rate proves to be above K_2. Figure 15.2 shows that the probability of this is higher for the implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price for the option. A relatively high price leads to a relatively

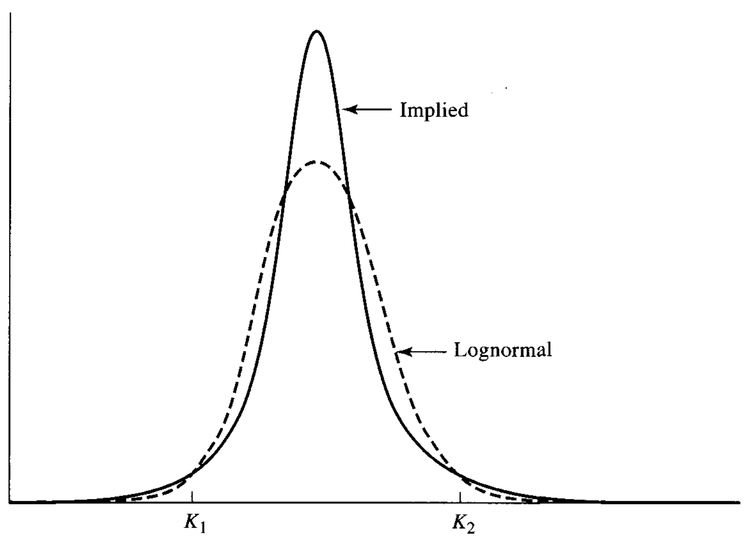

Figure 15.2 Implied distribution and lognormal distribution for foreign currency options

[1] This is known as kurtosis. Note that, in addition to having a heavier tail, the implied distribution is more "peaked". Both small and large movements in the exchange rate are more likely than with the lognormal distribution. Intermediate movements are less likely.

high implied volatility—and this is exactly what we observe in Figure 15.1 for the option. The two figures are therefore consistent with each other for high strike prices. Consider next a deep-out-of-the-money put option with a low strike price of K_1. This option pays off only if the exchange rate proves to be below K_1. Figure 15.2 shows that the probability of this is also higher for implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price, and a relatively high implied volatility, for this option as well. Again, this is exactly what we observe in Figure 15.1.

Reason for the Smile in Foreign Currency Options

We have just shown that the smile used by traders for foreign currency options implies that they consider that the lognormal distribution understates the probability of extreme movements in exchange rates. To test whether they are right, we examined the daily movements in 12 different exchange rates over a 10-year period. As a first step we calculated the standard deviation of daily percentage change in each exchange rate. We then noted how often the actual percentage change exceeded one standard deviation, two standard deviations, and so on. Finally, we calculated how often this would have happened if the percentage changes had been normally distributed. (The lognormal model implies that percentage changes are almost exactly normally distributed over a one-day time period.) The results are shown in Table 15.1.[2]

Daily changes exceed three standard deviations on 1.34% of days. The lognormal model predicts that this should happen on only 0.27% of days. Daily changes exceed four, five, and six standard deviations on 0.29%, 0.08%, and 0.03% of days, respectively. The lognormal model predicts that we should hardly ever observe this happening. The table therefore provides evidence to support the existence of heavy tails and the volatility smile used by traders.

Why are exchange rates not lognormally distributed? Two of the conditions for an asset price to have a lognormal distribution are:

1. The volatility of the asset is constant.
2. The price of the asset changes smoothly with no jumps.

In practice, neither of these conditions is satisfied for an exchange rate. The volatility of an

Table 15.1 Percent of days when daily exchange rate moves are greater than one, two,..., six standard deviations (S.D. = standard deviation of daily change)

	Real world	*Lognormal model*
>1 S.D.	25.04	31.73
>2 S.D.	5.27	4.55
>3 S.D.	1.34	0.27
>4 S.D.	0.29	0.01
>5 S.D.	0.08	0.00
>6 S.D.	0.03	0.00

[2] This table is taken from J. C. Hull and A. White, "Value at Risk When Daily Changes in Market Variables Are Not Normally Distributed," *Journal of Derivatives*, 5, no. 3 (Spring 1998), 9–19.

exchange rate is far from constant, and exchange rates frequently exhibit jumps.[3] It turns out that the effect of both a nonconstant volatility and jumps is that extreme outcomes become more likely.

The impact of jumps and nonconstant volatility depends on the option maturity. The percentage impact of a nonconstant volatility on prices becomes more pronounced as the maturity of the option is increased, but the volatility smile created by the nonconstant volatility usually becomes less pronounced. The percentage impact of jumps on both prices and the volatility smile becomes less pronounced as the maturity of the option is increased. When we look at sufficiently long dated options, jumps tend to get "averaged out", so that the stock price distribution when there are jumps is almost indistinguishable from the one obtained when the stock price changes smoothly.

15.3 EQUITY OPTIONS

The volatility smile used by traders to price equity options (both those on individual stocks and those on stock indices) has the general form shown in Figure 15.3. This is sometimes referred to as a *volatility skew*. The volatility decreases as the strike price increases. The volatility used to price a low-strike-price option (i.e., a deep-out-of-the-money put or a deep-in-the-money call) is significantly higher than that used to price a high-strike-price option (i.e., a deep-in-the-money put or a deep-out-of-the-money call).

The volatility smile for equity options corresponds to the implied probability distribution given by the solid line in Figure 15.4. A lognormal distribution with the same mean and standard

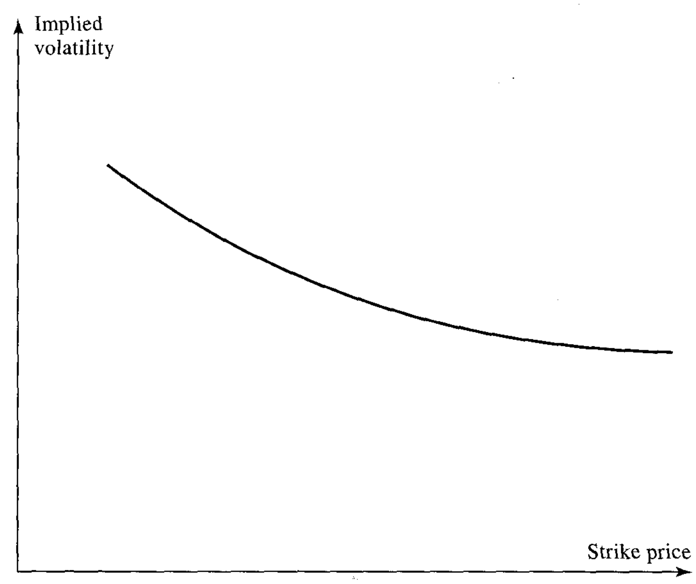

Figure 15.3 Volatility smile for equities

[3] Often the jumps are in response to the actions of central banks.

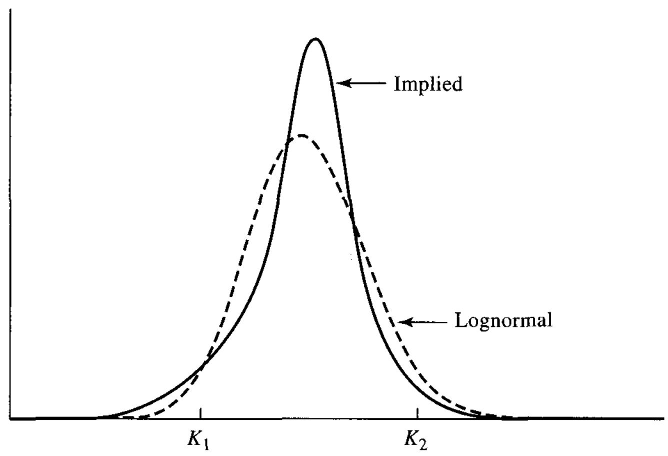

Figure 15.4 Implied distribution and lognormal distribution for equity options

deviation as the implied distribution is shown by the dotted line. It can be seen that the implied distribution has a heavier left tail and a less heavy right tail than the lognormal distribution.

To see that Figures 15.3 and 15.4 are consistent with each other, we proceed as for Figures 15.1 and 15.2 and consider options that are deep out of the money. From Figure 15.4 a deep-out-of-the-money call with a strike price of K_2 has a lower price when the implied distribution is used than when the lognormal distribution is used. This is because the option pays off only if the stock price proves to be above K_2, and the probability of this is lower for the implied probability distribution than for the lognormal distribution. Therefore we expect the implied distribution to give a relatively low price for the option. A relatively low price leads to a relatively low implied volatility—and this is exactly what we observe in Figure 15.3 for the option. Consider next a deep-out-of-the-money put option with a strike price of K_1. This option pays off only if the stock price proves to be below K_1. Figure 15.3 shows that the probability of this is higher for implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price, and a relatively high implied volatility, for this option. Again, this is exactly what we observe in Figure 15.3.

The Reason for the Smile in Equity Options

One possible explanation for the smile in equity options concerns leverage. As a company's equity declines in value, the company's leverage increases. As a result the volatility of its equity increases, making even lower stock prices more likely. As a company's equity increases in value, leverage decreases. As a result the volatility of its equity declines, making higher stock prices less likely. This argument shows that we can expect the volatility of equity to be a decreasing function of price and is consistent with Figures 15.3 and 15.4.

It is an interesting observation that the pattern in Figure 15.3 for equities has existed only since the stock market crash of October 1987. Prior to October 1987 implied volatilities were much less dependent on strike price. This has led Mark Rubinstein to suggest that one reason for the pattern in Figure 15.3 may be "crashophobia". Traders are concerned about the possibility of another crash similar to October 1987, and they price options accordingly. There is some empirical support for crashophobia. Whenever the market declines (increases), there is a tendency for the skew in Figure 15.3 to become more (less) pronounced.

15.4 THE VOLATILITY TERM STRUCTURE AND VOLATILITY SURFACES

In addition to a volatility smile, traders use a volatility term structure when pricing options. This means that the volatility used to price an at-the-money option depends on the maturity of the option. Volatility tends to be an increasing function of maturity when short-dated volatilities are historically low. This is because there is then an expectation that volatilities will increase. Similarly, volatility tends to be an decreasing function of maturity when short-dated volatilities are historically high. This is because there is then an expectation that volatilities will decrease.

Volatility surfaces combine volatility smiles with the volatility term structure to tabulate the volatilities appropriate for pricing an option with any strike price and any maturity. An example of a volatility surfaces that might be used for foreign currency options is shown in Table 15.2.

One dimension of Table 15.2 is strike price; the other is time to maturity. The main body of the table shows implied volatilities calculated from the Black–Scholes model. At any given time, some of the entries in the table are likely to correspond to options for which reliable market data are available. The implied volatilities for these options are calculated directly from their market prices and entered into the table. The rest of the table is determined using linear interpolation.

When a new option has to be valued, financial engineers look up the appropriate volatility in the table. For example, when valuing a nine-month option with a strike price of 1.05, a financial engineer would interpolate between 13.4 and 14.0 to obtain a volatility of 13.7%. This is the volatility that would be used in the Black–Scholes formula (or in the binomial tree model, which we will discuss further in Chapter 18).

The shape of the volatility smile depends on the option maturity. As illustrated in Table 15.2, the smile tends to become less pronounced as the option maturity increases. Define T as the time to maturity and F_0 as the forward price of the asset. Some financial engineers choose to define the volatility smile as the relationship between implied volatility and

$$\frac{1}{\sqrt{T}} \ln \frac{K}{F_0}$$

Table 15.2 Volatility surface

	Strike price				
	0.90	***0.95***	***1.00***	***1.05***	***1.10***
1 month	14.2	13.0	12.0	13.1	14.5
3 month	14.0	13.0	12.0	13.1	14.2
6 month	14.1	13.3	12.5	13.4	14.3
1 year	14.7	14.0	13.5	14.0	14.8
2 year	15.0	14.4	14.0	14.5	15.1
5 year	14.8	14.6	14.4	14.7	15.0

rather than as the relationship between the implied volatility and K. The smile is then usually much less dependent on the time to maturity.[4]

The Role of the Model

How important is the pricing model if traders are prepared to use a different volatility for every option? It can be argued that the Black–Scholes model is no more than a sophisticated interpolation tool used by traders for ensuring that an option is priced consistently with the market prices of other actively traded options. If traders stopped using Black–Scholes and switched to another plausible model, the volatility surface would change and the shape of the smile would change. But arguably, the dollar prices quoted in the market would not change appreciably.

15.5 GREEK LETTERS

The volatility smile complicates the calculation of Greek letters. Derman describes a number of volatility regimes or rules of thumb than are sometimes assumed by traders.[5] The simplest of these is known as the *sticky strike rule*. This assumes that the implied volatility of an option remains constant from one day to the next. It means that Greek letters calculated using the Black–Scholes assumptions are correct provided that the volatility used for an option is its current implied volatility.

A more complicated rule is known as the *sticky delta* rule. This assumes that the relationship we observe between an option price and S/K today will apply tomorrow. As the price of the underlying asset changes the implied volatility of the option is assumed to change to reflect the option's moneyness (i.e., the extent to which it is in or out of the money). If we use the sticky delta rule, the formulas for Greek letters given in the Chapter 14 are no longer correct. For example, delta of a call option is given by

$$\frac{\partial c_{\mathrm{BS}}}{\partial S} + \frac{\partial c_{\mathrm{BS}}}{\partial \sigma_{\mathrm{imp}}} \frac{\partial \sigma_{\mathrm{imp}}}{\partial S}$$

where c_{BS} is the Black–Scholes price of the option expressed as a function of the asset price S and the implied volatility σ_{imp}. Consider the impact of this formula on the delta of an equity call option. From Figure 15.3, volatility is a decreasing function of the strike price K. Alternatively it can be regarded as an increasing function of S/K. Under the sticky delta model, therefore, the volatility increases as the asset price increases, so that

$$\frac{\partial \sigma_{\mathrm{imp}}}{\partial S} > 0$$

As a result, delta is higher than that given by the Black–Scholes assumptions.

It turns out that the sticky strike and sticky delta rules do not correspond to internally consistent models (except when the volatility smile is flat for all maturities). A model that can be made exactly consistent with the smiles is known as the *implied volatility function* model or the *implied tree model*. We will explain this model in Chapter 20.

[4] For a discussion of this approach, see S. Natenberg, *Option Pricing and Volatility: Advanced Trading Strategies and Techniques*, 2nd edn., McGraw-Hill, New York, 1994; R. Tompkins, *Options Analysis: A State of the Art Guide to Options Pricing*, Irwin, Burr Ridge, IL,1994.

[5] See E. Derman, "Regimes of Volatility," *RISK*, April 1999, pp. 54–59.

In practice many banks try to ensure that their exposure to the changes in the volatility surface that are most commonly observed is reasonably small. One technique for identifying these changes is principal components analysis, which we discuss in Chapter 16.

15.6 WHEN A SINGLE LARGE JUMP IS ANTICIPATED

Suppose that a stock price is currently \$50 and an important news announcement in a few days is expected to either increase the stock price by \$8 or reduce it by \$8. (This announcement could concern the outcome of a takeover attempt or the verdict in an important lawsuit.) The probability distribution of the stock price in, say, three months might then consist of a mixture of two lognormal distributions, the first corresponding to favorable news, and the second to unfavorable news. The situation is illustrated in Figure 15.5. The solid line shows the mixtures-of-lognormals distribution for the stock price in three months; the dashed line shows a lognormal distribution with the same mean and standard deviation as this distribution. Assume that favorable news and unfavorable news are equally likely.[6] Assume also that after the news (favorable or unfavorable) the volatility will be constant at 20% for three months.

Consider a three-month European call option on the stock with a strike price of \$50. We assume that the risk-free interest rate is 5% per annum. Because the news announcement is expected very soon, the value of the option assuming favorable news can be calculated from the Black–Scholes formula with $S_0 = 58$, $K = 50$, $r = 5\%$, $\sigma = 20\%$, and $T = 0.25$. It is 8.743. Similarly, the value of the option assuming unfavorable news can be calculated from the Black–Scholes formula with $S_0 = 42$, $K = 50$, $r = 5\%$, $\sigma = 20\%$, and $T = 0.25$. It is 0.101. The value of the call option today should therefore be

$$0.5 \times 8.743 + 0.5 \times 0.101 = 4.422$$

The implied volatility calculated from this option price is 41.48%.

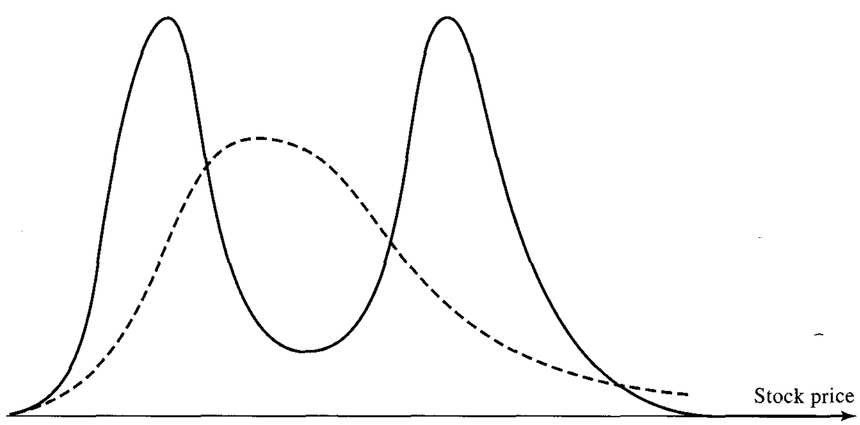

Figure 15.5 Effect of a single large jump: the solid line is the true distribution; the dashed line is the lognormal distribution

[6] Strictly speaking, we are assuming that the probabilities are equally likely in a risk-neutral world.

Table 15.3 Implied volatilities in a situation where an important announcement is imminent

Strike price ($)	*Call price if good news ($)*	*Call price if bad news ($)*	*Call price today ($)*	*Implied volatility (%)*
35	23.435	7.471	15.453	30.95
40	18.497	3.169	10.833	35.46
45	13.565	0.771	7.168	39.94
50	8.743	0.101	4.422	41.48
55	4.546	0.008	2.277	39.27
60	1.764	0.000	0.882	35.66
65	0.494	0.000	0.247	32.50

A similar calculation can be made for other strike prices and a volatility smile constructed. The results of doing this are shown in Table 15.3, and the volatility smile is shown in Figure 15.6.[7] It turns out that we are in the opposite situation to that of Figure 15.1. At-the-money options have higher volatilities than either out-of-the-money or in-the-money options.

15.7 EMPIRICAL RESEARCH

A number of problems arise in carrying out empirical research to test the Black–Scholes and other option pricing models.[8] The first problem is that any statistical hypothesis about how options are priced has to be a joint hypothesis to the effect that (1) the option pricing formula is

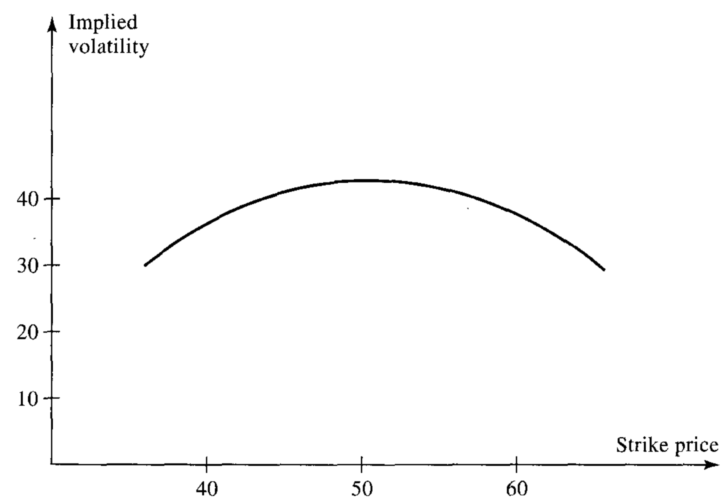

Figure 15.6 Volatility smile for situation in Table 15.3

[7] In this case, the smile is a frown!

[8] See the end-of-chapter references for citations to the studies reviewed in this section.

correct and (2) markets are efficient. If the hypothesis is rejected, it may be the case that (1) is untrue, (2) is untrue, or both (1) and (2) are untrue. A second problem is that the stock price volatility is an unobservable variable. One approach is to estimate the volatility from historical stock price data. Alternatively, implied volatilities can be used in some way. A third problem for the researcher is to make sure that data on the stock price and option price are synchronous. For example, if the option is thinly traded, it is not likely to be acceptable to compare closing option prices with closing stock prices. The closing option price might correspond to a trade at 1:00 p.m., whereas the closing stock price corresponds to a trade at 4:00 p.m.

Black and Scholes (1972) and Galai (1977) have tested whether it is possible to make excess returns above the risk-free rate of interest by buying options that are undervalued by the market (relative to the theoretical price) and selling options that are overvalued by the market (relative to the theoretical price). Black and Scholes used data from the over-the-counter options market where options are dividend protected. Galai used data from the Chicago Board Options Exchange (CBOE) where options are not protected against the effects of cash dividends. Galai used Black's approximation as described in Section 12.13 to incorporate the effect of anticipated dividends into the option price. Both studies showed that, in the absence of transactions costs, significant excess returns over the risk-free rate can be obtained by buying undervalued options and selling overvalued options. However, it is possible that these excess returns are available only to market makers and that, when transactions costs are considered, they vanish.

A number of researchers have chosen to make no assumptions about the behavior of stock prices and have tested whether arbitrage strategies can be used to make a riskless profit in options markets. Garman (1976) provides a computational procedure for finding any arbitrage possibilities that exist in a given situation. One frequently cited study by Klemkosky and Resnick (1979) tests whether the relationship in equation (8.8) is ever violated. It concludes that some small arbitrage profits are possible from using the relationship. These are due mainly to the overpricing of American calls.

Chiras and Manaster (1978) carried out a study using CBOE data to compare a weighted implied volatility from options on a stock at a point in time with the volatility calculated from historical data. They found that the former provide a much better forecast of the volatility of the stock price during the life of the option. We can conclude that option traders are using more than just historical data when determining future volatilities. Chiras and Manaster also tested to see whether it was possible to make above-average returns by buying options with low implied volatilities and selling options with high implied volatilities. The strategy showed a profit of 10% per month. The Chiras and Manaster study can be interpreted as providing good support for the Black–Scholes model and showing that the CBOE was inefficient in some respects.

MacBeth and Merville (1979) tested the Black–Scholes model using a different approach. They looked at different call options on the same stock at the same time and compared the volatilities implied by the option prices. The stocks chosen were AT&T, Avon, Kodak, Exxon, IBM, and Xerox, and the time period considered was the year 1976. They found that implied volatilities tended to be relatively high for in-the-money options and relatively low for out-of-the-money options. A relatively high implied volatility is indicative of a relatively high option price, and a relatively low implied volatility is indicative of a relatively low option price. Therefore, if it is assumed that Black–Scholes prices at-the-money options correctly, it can be concluded that out-of-the-money (high strike price) call options are overpriced by Black–Scholes and in-the-money (low strike price) call options are underpriced by Black–Scholes. These effects become more pronounced as the time to maturity increases and the degree to which the option is in or out of the money increases. MacBeth and Merville's results are consistent with Figure 15.3. The results were confirmed by Lauterbach and Schultz (1990) in a later study concerned with the pricing of warrants.

Rubinstein has done a great deal of research similar to that of MacBeth and Merville. No clear-cut pattern emerged from his early research, but the research in his 1994 paper and joint 1996 paper with Jackwerth gives results consistent with Figure 15.3. Options with low strike prices have much higher volatilities than those with high strike prices. As mentioned previously in the chapter, leverage and the resultant negative correlation between volatility and stock price may partially account for the finding. It is also possible that investors fear a repeat of the crash of 1987.

A number of authors have researched the pricing of options on assets other than stocks. For example, Shastri and Tandon (June 1986) and Bodurtha and Courtadon (1987) have examined the market prices of currency options; in another paper, Shastri and Tandon (December 1986) have examined the market prices of futures options; and Chance (1986) has examined the market prices of index options.

In most cases, the mispricing by Black–Scholes is not sufficient to present profitable opportunities to investors when transactions costs and bid–offer spreads are taken into account. When profitable opportunities are sought, it is important to bear in mind that, even for a market maker, some time must elapse between a profitable opportunity being identified and action being taken. This delay, even if it is only to the next trade, can be sufficient to eliminate the profitable opportunity.

SUMMARY

The Black–Scholes model and its extensions assume that the probability distribution of the underlying asset at any given future time is lognormal. This assumption is not the one made by traders. They assume the probability distribution of an equity price has a heavier left tail and a less heavy right tail than the lognormal distribution. They also assume that the probability of an exchange rate has a heavier right tail and a heavier left tail than the lognormal distribution.

Traders use volatility smiles to allow for nonlognormality. The volatility smile defines the relationship between the implied volatility of an option and its strike price. For equity options, the volatility smile tends to be downward sloping. This means that out-of-the-money puts and in-the-money calls tend to have high implied volatilities, whereas out-of-the-money calls and in-the-money puts tend to have low implied volatilities. For foreign currency options, the volatility smile is U-shaped. Deep-out-of-the-money and deep-in-the-money options have higher implied volatilities than at-the-money options.

Often traders also use a volatility term structure. The implied volatility of an option then depends on its life. When volatility smiles and volatility term structures are combined, they produce a volatility surface. This defines implied volatility as a function of both the strike price and the time to maturity.

SUGGESTIONS FOR FURTHER READING

Bakshi, G., C. Cao, and Z. Chen, "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, no. 5 (December 1997), 2004–49.

Black, F., "How to Use the Holes in Black–Scholes," *RISK*, March 1988.

Black, F., and M. Scholes, "The Valuation of Option Contracts and a Test of Market Efficiency," *Journal of Finance*, 27 (May 1972), 399–418.

Bodurtha, J. N., and G. R. Courtadon, "Tests of an American Option Pricing Model on the Foreign Currency Options Market," *Journal of Financial and Quantitative Analysis*, 22 (June 1987), 153–68.

Breeden, M., and R. Litzenberger, "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business*, 51 (1978), 621–51.

Chance, D. M., "Empirical Tests of the Pricing of Index Call Options," *Advances in Futures and Options Research*, 1, pt. A (1986), 141–66.

Chiras, D., and S. Manaster, "The Information Content of Option Prices and a Test of Market Efficiency," *Journal of Financial Economics*, 6 (September 1978), 213–34.

Derman, E., "Regimes of Volatility," *RISK*, April 1999, pp. 55–59.

Dumas, B., J. Fleming, and R. E. Whaley, "Implied Volatility Functions: Empirical Tests," *Journal of Finance* 53, 6 (December 1998), 2059–2106.

Galai, D., "Tests of Market Efficiency and the Chicago Board Options Exchange." *Journal of Business*, 50 (April 1977), 167–97.

Garman, M., "An Algebra for Evaluating Hedge Portfolios," *Journal of Financial Economics*, 3 (October 1976), 403–27.

Harvey, C. R., and R. E. Whaley, "S&P 100 Index Option Volatility," *Journal of Finance*, 46 (1991), 1551–61.

Harvey, C. R., and R. E. Whaley, "Market Volatility Prediction and the Efficiency of the S&P 100 Index Option Market," *Journal of Financial Economics*, 31 (1992), 43–73.

Harvey, C. R., and R. E. Whaley, "Dividends and S&P 100 Index Option Valuations," *Journal of Futures Markets*, 12 (1992), 123–37.

Jackwerth, J. C., and M. Rubinstein, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51 (December 1996), 1611–31.

Klemkosky, R. C., and B. G. Resnick, "Put–Call Parity and Market Efficiency," *Journal of Finance*, 34 (December 1979), 1141–55.

Lauterbach, B., and P. Schultz, "Pricing Warrants: An Empirical Study of the Black–Scholes Model and Its Alternatives," *Journal of Finance*, 4, no. 4 (September 1990), 1181–1210.

MacBeth, J. D., and L. J. Merville, "An Empirical Examination of the Black–Scholes Call Option Pricing Model," *Journal of Finance*, 34 (December 1979), 1173–86.

Melick, W. R., and C. P. Thomas, "Recovering an Asset's Implied Probability Density Function from Option Prices: An Application to Crude Oil during the Gulf Crisis," *Journal of Financial and Quantitative Analysis*, 32, no. 1 (March 1997), 91–115.

Rubinstein, M., "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978," *Journal of Finance*, 40 (June 1985), 455–80.

Rubinstein, M., "Implied Binomial Trees," *Journal of Finance*, 49, 3 (July 1994), 771–818.

Shastri, K., and K. Tandon, "An Empirical Test of a Valuation Model for American Options on Futures Contracts," *Journal of Financial and Quantitative Analysis*, 21 (December 1986), 377–92.

Shastri, K., and K. Tandon, "Valuation of Foreign Currency Options: Some Empirical Tests," *Journal of Financial and Quantitative Analysis*, 21 (June 1986), 145–60.

Xu, X., and S. J. Taylor, "The Term Structure of Volatility Implied by Foreign Exchange Options," *Journal of Financial and Quantitative Analysis*, 29 (1994), 57–74.

QUESTIONS AND PROBLEMS (ANSWERS IN SOLUTIONS MANUAL)

15.1. What pattern of implied volatilities is likely to be observed when
  a. Both tails of the stock price distribution are less heavy than those of the lognormal distribution?
  b. The right tail is heavier, and the left tail is less heavy, than that of a lognormal distribution?

15.2. What pattern of implied volatilities is likely to be observed for six-month options when the volatility is uncertain and positively correlated to the stock price?

15.3. What pattern of implied volatilities is likely to be caused by jumps in the underlying asset price? Is the pattern likely to be more pronounced for a six-month option than for a three-month option?

15.4. A call and put option have the same strike price and time to maturity. Show that the difference between their prices should be the same for any option pricing model.

15.5. Explain carefully why Figure 15.4 is consistent with Figure 15.3.

15.6. The market price of a European call is \$3.00 and its Black–Scholes price is \$3.50. The Black–Scholes price of a European put option with the same strike price and time to maturity is \$1.00. What should the market price of this option be? Explain the reasons for your answer.

15.7. A stock price is currently \$20. Tomorrow, news is expected to be announced that will either increase the price by \$5 or decrease the price by \$5. What are the problems in using Black–Scholes to value one-month options on the stock?

15.8. What are the major problems in testing a stock option pricing model empirically?

15.9. Suppose that a central bank's policy is to allow an exchange rate to fluctuate between 0.97 and 1.03. What pattern of implied volatilities for options on the exchange rate would you expect to see?

15.10. Option traders sometimes refer to deep-out-of-the-money options as being options on volatility. Why do you think they do this?

15.11. A European call option on a certain stock has a strike price of \$30, a time to maturity of one year, and an implied volatility of 30%. A European put option on the same stock has a strike price of \$30, a time to maturity of one year, and an implied volatility of 33%. What is the arbitrage opportunity open to a trader? Does the arbitrage work only when the lognormal assumption underlying Black–Scholes holds? Explain carefully the reasons for your answer.

15.12. Suppose that the result of a major lawsuit affecting Microsoft is due to be announced tomorrow. Microsoft's stock price is currently \$60. If the ruling is favorable to Microsoft, the stock price is expected to jump to \$75. If it is unfavorable, the stock is expected to jump to \$50. What is the risk-neutral probability of a favorable ruling? Assume that the volatility of Microsoft's stock will be 25% for six months after the ruling if the ruling is favorable and 40% if it is unfavorable. Use DerivaGem to calculate the relationship between implied volatility and strike price for six-month European options on Microsoft today. Microsoft does not pay dividends. Assume that the six-month risk-free rate is 6%. Consider call options with strike prices of 30, 40, 50, 60, 70, and 80.

15.13. An exchange rate is currently 0.8000. The volatility of the exchange rate is quoted as 12% and interest rates in the two countries are the same. Using the lognormal assumption, estimate the probability that the exchange rate in three months will be (a) less than 0.7000, (b) between 0.7000 and 0.7500, (c) between 0.7500 and 0.8000, (d) between 0.8000 and 0.8500, (e) between 0.8500 and 0.9000, and (f) greater than 0.9000. Based on the volatility smile usually observed in the market for exchange rates, which of these estimates would you expect to be too low and which would you expect to be too high?

15.14. A stock price is \$40. A six-month European call option on the stock with a strike price of \$30 has an implied volatility of 35%. A six-month European call option on the stock with a strike price of \$50 has an implied volatility of 28%. The six-month risk-free rate is 5% and no dividends are expected. Explain why the two implied volatilities are different. Use DerivaGem to calculate the prices of the two options. Use put–call parity to calculate the prices of six-month European put options with strike prices of \$30 and \$50. Use DerivaGem to calculate the implied volatilities of these two put options.

15.15. "The Black–Scholes model is used by traders as an interpolation tool." Discuss this view.

ASSIGNMENT QUESTIONS

15.16. A company's stock is selling for \$4. The company has no outstanding debt. Analysts consider the liquidation value of the company to be at least \$300,000 and there are 100,000 shares outstanding. What volatility smile would you expect to see?

15.17. A company is currently awaiting the outcome of a major lawsuit. This is expected to be known within one month. The stock price is currently \$20. If the outcome is positive, the stock price is expected to be \$24 at the end of one month. If the outcome is negative, it is expected to be \$18 at this time. The one-month risk-free interest rate is 8% per annum.
 a. What is the risk-neutral probability of a positive outcome?
 b. What are the values of one-month call options with strike prices of \$19, \$20, \$21, \$22, and \$23?
 c. Use DerivaGem to calculate a volatility smile for one-month call options.
 d. Verify that the same volatility smile is obtained for one-month put options.

15.18. A futures price is currently \$40. The risk-free interest rate is 5%. Some news is expected tomorrow that will cause the volatility over the next three months to be either 10% or 30%. There is a 60% chance of the first outcome and a 40% chance of the second outcome. Use DerivaGem to calculate a volatility smile for three-month options.

15.19. Data for a number of foreign currencies are provided on the author's Web site:
 www.rotman.utoronto.ca/~hull
 Choose a currency and use the data to produce a table similar to Table 15.1.

15.20. Data for a number of stock indices are provided on the author's Web site:
 www.rotman.utoronto.ca/~hull
 Choose an index and test whether a three standard deviation down movement happens more often than a three standard deviation up movement.

15.21. Consider a European call and a European put with the same strike price and time to maturity. Show that they change in value by the same amount when the volatility increases from a level σ_1 to a new level σ_2 within a short period of time. (*Hint*: Use put–call parity.)

APPENDIX 15A

Determining Implied Risk-Neutral Distributions from Volatility Smiles

The price of a European call option on an asset with strike price K and maturity T is given by

$$c = e^{-rT}\int_{S_T=K}^{\infty}(S_T - K)g(S_T)\,dS_T$$

where r is the interest rate (assumed constant), S_T is the asset price at time T, and g is the risk-neutral probability density function of S_T. Differentiating once with respect to K, we obtain

$$\frac{\partial c}{\partial K} = -e^{-rT}\int_{S_T=K}^{\infty} g(S_T)\,dS_T$$

Differentiating again with respect to K gives

$$\frac{\partial^2 c}{\partial K^2} = e^{-rT}g(K)$$

This shows that the probability density function, g, is given by

$$g(K) = e^{rT}\frac{\partial^2 c}{\partial K^2}$$

This result, which is from Breeden and Litzenberger (1978), allows risk-neutral probability distributions to be estimated from volatility smiles. Suppose that c_1, c_2, and c_3 are the prices of European call options with maturity T and strike prices are $K - \delta$, K, and $K + \delta$, respectively. Assuming δ is small, an estimate of $g(K)$ is

$$e^{rT}\frac{c_1 + c_3 - 2c_2}{\delta^2}$$

CHAPTER 16

VALUE AT RISK

In Chapter 14 we examined measures such as delta, gamma, and vega for describing different aspects of the risk in a portfolio consisting of options and other financial assets. A financial institution usually calculates each of these measures each day for every market variable to which it is exposed. Often there are hundreds, or even thousands, of these market variables. A delta–gamma–vega analysis therefore leads to a huge number of different risk measures being produced each day. These risk measures provide valuable information for a trader who is responsible for managing the part of the financial institution's portfolio that is dependent on a particular market variable, but they are of limited use to senior management.

Value at risk (VaR) is an attempt to provide a single number summarizing the total risk in a portfolio of financial assets for senior management. It has become widely used by corporate treasurers and fund managers as well as by financial institutions. Central bank regulators also use VaR in determining the capital a bank is required to keep to reflect the market risks it is bearing.[1]

In this chapter we explain the VaR measure and describe the two main approaches for calculating it. These are the *historical simulation* approach and the *model-building* approach. Both are widely used by both financial and nonfinancial companies. There is no general agreement on which of the two is better.

16.1 THE VaR MEASURE

When using the value-at-risk measure, the manager in charge of a portfolio of financial instruments is interested in making a statement of the following form:

"We are X percent certain that we will not lose more than V dollars in the next N days."

The variable V is the VaR of the portfolio. It is a function of two parameters: N, the time horizon, and X, the confidence level. It is the loss level over N days that the manager is X% certain will not be exceeded.

In calculating a bank's capital for market risk, regulators use $N = 10$ and $X = 99$. This means that they focus on the loss level over a 10-day period that is expected to be exceeded only 1% of the time. The capital they require the bank to keep is at least three times this VaR measure.[2]

[1] For a discussion of this, see P. Jackson, D. J. Maude, and W. Perraudin, "Bank Capital and Value at Risk," *Journal of Derivatives*, 4, no. 3 (Spring 1997), 73–90.

[2] To be more precise, the market risk capital required for a particular bank is k times the 10-day 99% VaR, where the multiplier k is chosen by the regulators and is at least 3.0.

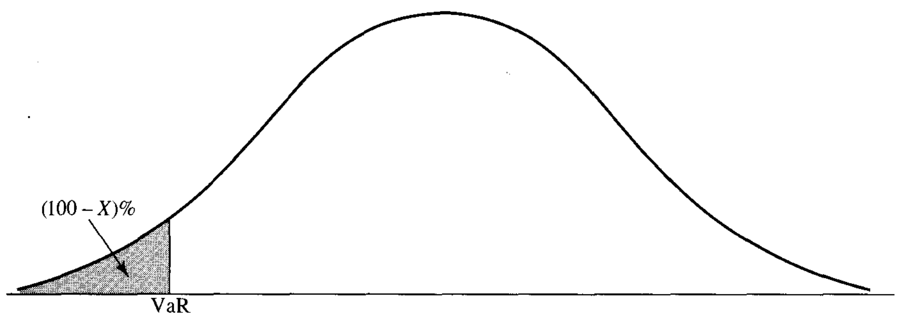

Figure 16.1 Calculation of VaR from the probability distribution of changes in the portfolio value; confidence level is $X\%$

In general, when N days is the time horizon and $X\%$ is the confidence level, VaR is the loss corresponding to the $(100 - X)$th percentile of the distribution of the change in the value of the portfolio over the next N days. For example, when $N = 5$ and $X = 97$, it is the third percentile of the distribution of changes in the value of the portfolio over the next five days. Figure 16.1 illustrates VaR for the situation where the change in the value of the portfolio is approximately normally distributed.

VaR is an attractive measure because it is easy to understand. In essence, it asks the simple question "How bad can things get?" This is the question all senior managers want answered. They are very comfortable with the idea of compressing all the Greek letters for all the market variables underlying the portfolio into a single number.

If we accept that it is useful to have a single number to describe the risk of a portfolio, an interesting question is whether VaR is the best alternative. Some researchers have argued that VaR may tempt traders to choose a portfolio with a return distribution similar to that in Figure 16.2. The portfolios in Figures 16.1 and 16.2 have the same VaR, but the portfolio in Figure 16.2 is much riskier because potential losses are much larger.

A measure that deals with the problem we have just mentioned is *Conditional VaR* (C-VaR).[3] Whereas VaR asks the question "How bad can things get?", C-VaR asks "If things do get bad, how much can we expect to lose?" C-VaR is the expected loss during an N-day period conditional

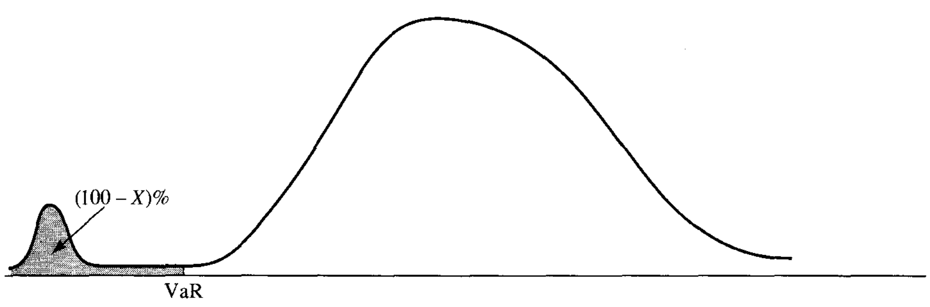

Figure 16.2 Alternative situation to Figure 16.1; VaR is the same, but the potential loss is larger

[3] This measure was suggested by P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent Measures of Risk," *Mathematical Finance*, 9 (1999), 203–28. These authors define certain properties that a good risk measure should have and show that the standard VaR measure does not have all of them.

that we are in the $(100 - X)\%$ left tail of the distribution. For example, with $X = 99$ and $N = 10$, C-VaR is the average amount we lose over a 10-day period assuming that a 1% worst-case event occurs.

In spite of its weaknesses, VaR (not C-VaR) is the most popular measure of risk among both regulators and senior management. We will therefore devote most of the rest of this chapter to how it can be measured.

The Time Horizon

In theory, VaR has two parameters. These are N, the time horizon measured in days, and X, the confidence interval. In practice, analysts almost invariably set $N = 1$ in the first instance. This is because there is not enough data to estimate directly the behavior of market variables over periods of time longer than one day. The usual assumption is

$$N\text{-day VaR} = 1\text{-day VaR} \times \sqrt{N}$$

This formula is exactly true when the changes in the value of the portfolio on successive days have independent identical normal distributions with mean zero. In other cases it is an approximation.

We mentioned earlier that regulators require a bank's capital to be at least three times the 10-day 99% VaR. Given the way a 10-day VaR is calculated, this capital level is, to all intents and purposes, $3 \times \sqrt{10} = 9.49$ times the 1-day 99% VaR.

16.2 HISTORICAL SIMULATION

Historical simulation is one popular way of estimating VaR. It involves using past data in a very direct way as a guide to what might happen in the future. Suppose that we wish to calculate VaR for a portfolio using a one-day time horizon, a 99% confidence level, and 500 days of data. The first step is to identify the market variables affecting the portfolio. These will typically be exchange rates, equity prices, interest rates, and so on. We then collect data on the movements in these market variables over the most recent 500 days. This provides us with 500 alternative scenarios for what can happen between today and tomorrow. Scenario 1 is where the percentage changes in the

Table 16.1 Data for VaR historical simulation calculation

Day	*Market variable 1*	*Market variable 2*	...	*Market variable N*
0	20.33	0.1132	...	65.37
1	20.78	0.1159	...	64.91
2	21.44	0.1162	...	65.02
3	20.97	0.1184	...	64.90
⋮	⋮	⋮	⋮	⋮
498	25.72	0.1312	...	62.22
499	25.75	0.1323	...	61.99
500	25.85	0.1343	...	62.10