# FUTURE TRENDS FROM PAST CYCLES

## IDENTIFYING SHARE PRICE TRENDS AND TURNING POINTS THROUGH CYCLE, CHANNEL AND PROBABILITY ANALYSIS

BRIAN J. MILLARD

## About the Author

Brian J. Millard's background was as a scientist and until 1980 he was a senior lecturer at the University of London, publishing over 70 scientific papers.

He later became interested in the work of J. M. Hurst on cycles and channels in the stock market and, as this interest grew, spent time carrying out research in this field. Following his landmark book *Stocks and Shares Simplified*, published in 1980, Brian wrote a further five books on the application of scientific methods to the stock market.

His books on channel analysis are now universally recognised as taking forward the work of J. M. Hurst to a higher level by analysing price movement and especially the occurrence of predictable cycles in market data. Brian also published software to enable traders to apply his methods.

## Other Books by Brian Millard

*Quantitative Mass Spectrometry*

*Stocks and Shares Simplified: A Guide for the Smaller Investor*

*Traded Options Simplified*

*Millard on Channel Analysis: The Key to Share Price Prediction*

*Profitable Charting Techniques*

*Winning on the Stock Market*

*Channels & Cycles: A Tribute to J. M. Hurst*

## Brian Millard (1937-2009)

Brian worked tirelessly for the last two years, building on his earlier research and publications, to create the theories and techniques contained within this book. He died very suddenly in July 2009 before he could see the publication of this work.

I don't have a record of the many who helped in validating and testing the work presented – but I thank you for your contribution.

I would especially like to thank Martin Smart, who gave up his time to proof this work for my father, and Louise Hinchen at Harriman House, who has given me such support in bringing this book to you.

A great intellect and passion has been taken from the world; it is a lesser place for the loss.

Sandy has lost a wonderful husband, Alastair and I the best dad anyone could have, and his grandchildren – Rebecca, Eleanor, Madeline and Edward – a grandpa who was so proud of them.

I hope that this book brings you much success with your investments.

Simon Millard

Northampton, October 2009

# 1. Introduction

*"The further back you can look, the farther forward you can see."*

Winston S. Churchill

It is now just about ten years since I wrote my last book, *Channels & Cycles: A Tribute to J.M. Hurst*. I have been gratified by the response I have had to that book and the many kind comments I have received about my approach to trading on the stock market.

I still acknowledge that Hurst set me on the road along which I have been travelling since that time in 1979 (is it really 30 years ago?) when I picked up his book *The Profit Magic of Stock Transaction Timing*. This book was reprinted by Traders Press in 2000, and I urge readers to take advantage of its restored availability. For those around when it was first published in 1971 it was a breath of common sense in showing what is possible when approaching the markets with a measured, logical technique based on firm mathematical and scientific logic. New readers will see it in a different light, because now there are many authors and many software packages that use these important principles. To these new readers it might not now appear as revolutionary as it did when first published, but they will still enjoy Hurst's writing style and the book's logical approach to the improved timing of buying and selling decisions.

When considering a title for this book, I started by thinking about what goes through a trader's mind when contemplating a trade. It is usually: 'I think the price of this security will rise/fall within *n* hours/days/weeks/months.' The thought process of a much smaller number of traders will be: 'There is a high probability that the price of this security will rise/fall within *n* hours/days/weeks/months.'

Provided the second type of trader has carried out his analysis carefully, he will in the long run be more successful than the first type of trader.

You can now see my dilemma in choosing the title for this book. Should I stress the novelty aspect of using cycles in a way that hasn't been attempted before, or should I stress the importance of making sure that probability is always on your side? In the end I opted for novelty, being aware that novelty is always appealing.

Readers of my books on channel analysis will be aware that the most difficult aspect of that technique is in deciding when a channel has changed direction. How the channel has behaved in the gap between its last calculated value (half of the span used to calculate it back in time from the present) is open to interpretation. It is this which has been the subject of my research over the last ten years and which lies behind my writings in this book.

In *Channels & Cycles* I made the point that channel analysis could be carried out with a paper and pencil and that a computer was not strictly essential. However, the study of cycles and their relationship to channels has now moved on so much that a computer is absolutely essential. Some of the scans of cycles that I will describe take tens of millions of calculations to perform and quite clearly a fast computer is therefore essential.

Except for the cycle scans described in this book that are unique to the packages Channalyze and CCS Visions, the isolation of single cycles can be performed by any software package that allows you to specify your own calculations. These types of calculation can also be carried out in a spreadsheet application such as Microsoft Excel. Of course, in all of these cases a good amount of accurate historical data is required so that long-term cycles can be studied.

Channels can be drawn by Channalyze, but for other software packages a paper and pencil can still be useful for drawing constant depth channels based on a centred moving average which has been calculated by the computer software. These constant depth channels should not be confused with Bollinger bands, which virtually all software packages can produce. Bollinger bands are not of constant depth and, unlike channels, bear no obvious relationship to cycles.

I chose the title *Future Trends from Past Cycles* because it clearly describes what the book is all about. Although no mention is made of channels in this title, channels are an essential part of taking an investment decision. The channel is based on the future trend, but knowledge of the future trend does not in itself give the trader the optimum time to place a trade. I took a rather different approach in *Channels & Cycles*, placing more weight on a discussion of channels and less on cycles. The balance in this book is reversed, since quite clearly all channels are derived from cycles, or rather sums of cycles.

## Definition of a Trend

Different readers will have different perceptions of what constitutes a *trend*. There are many uses of this word in technical analysis. There is, however, one property that is often overlooked, but which is essential, and that is a timescale. Each trend should have a timescale associated with it. Once this is accepted, then it is perfectly reasonable to separate trends into short-term, medium-term and long-term.

## Definition of a Cycle

A perfect cycle in the market is a sine wave. Associated with the sine wave is a *wavelength*, measured in minutes, hours, days, weeks, years, etc and an *amplitude*. The amplitude is measured in whatever units the security is quoted, such as points for an index, a ratio for foreign exchange and currency units such as pounds, dollars and euros for stocks. A third parameter is the phase, which is how far along the cycle is from some arbitrary starting point. We will see later that sine waves are described by an equation, which allows us to know its future path provided we know the three parameters which are unique to that particular sine wave.

Unfortunately, market cycles are not stable throughout their lifetime, and their amplitude will change after a period of time, rendering predictions unreliable. The way around this dilemma is to use only those cycles which have been stable for a short period of time and should remain so in the near future.

## Determining Trends

The title of this book explains what it is all about. It is to determine as far as possible the future direction of a trend of interest, whatever its timescale. The title also explains that this can be done by an analysis of cycles present in market data. While stable-cycle analysis is the prime method, it is only one of the three analyses which must be carried out before a decision is taken about the future direction. These three analyses are:

1. Cycle analysis: determining those cycles which are currently stable and adding the appropriate cycles in order to estimate the future path of the trend. The current stability of the cycle sum is checked by means of a comparator, a centred nine-day moving average.
2. Channel analysis: confirming that the trend estimated by the cycle sum is supported by channel analysis.
3. Probability analysis: confirming that the trend is also supported by an analysis of the predicted probable price range a few days into the future.

Following this sequence of analyses, the ideal situation will be when a predicted change in direction of a trend is imminent and is likely to happen a few days into the future. The progress of the estimated trend can then be followed each day as more price information arrives. If the new information indicates that the trend estimation is beginning to vary from the initial prediction, then move to another security.

Before taking a position in the security, it is essential that the change in direction is confirmed. Because of the fact that market cycles are constantly varying, changes in direction must never be anticipated.

### TREND POSITION

Naturally, an estimation of a future trend will be liable to an uncertainty in its position over the course of time. This is illustrated in Figure 1.1.

***Figure 1.1 – There will be an uncertainty in the location of the central trend line.***

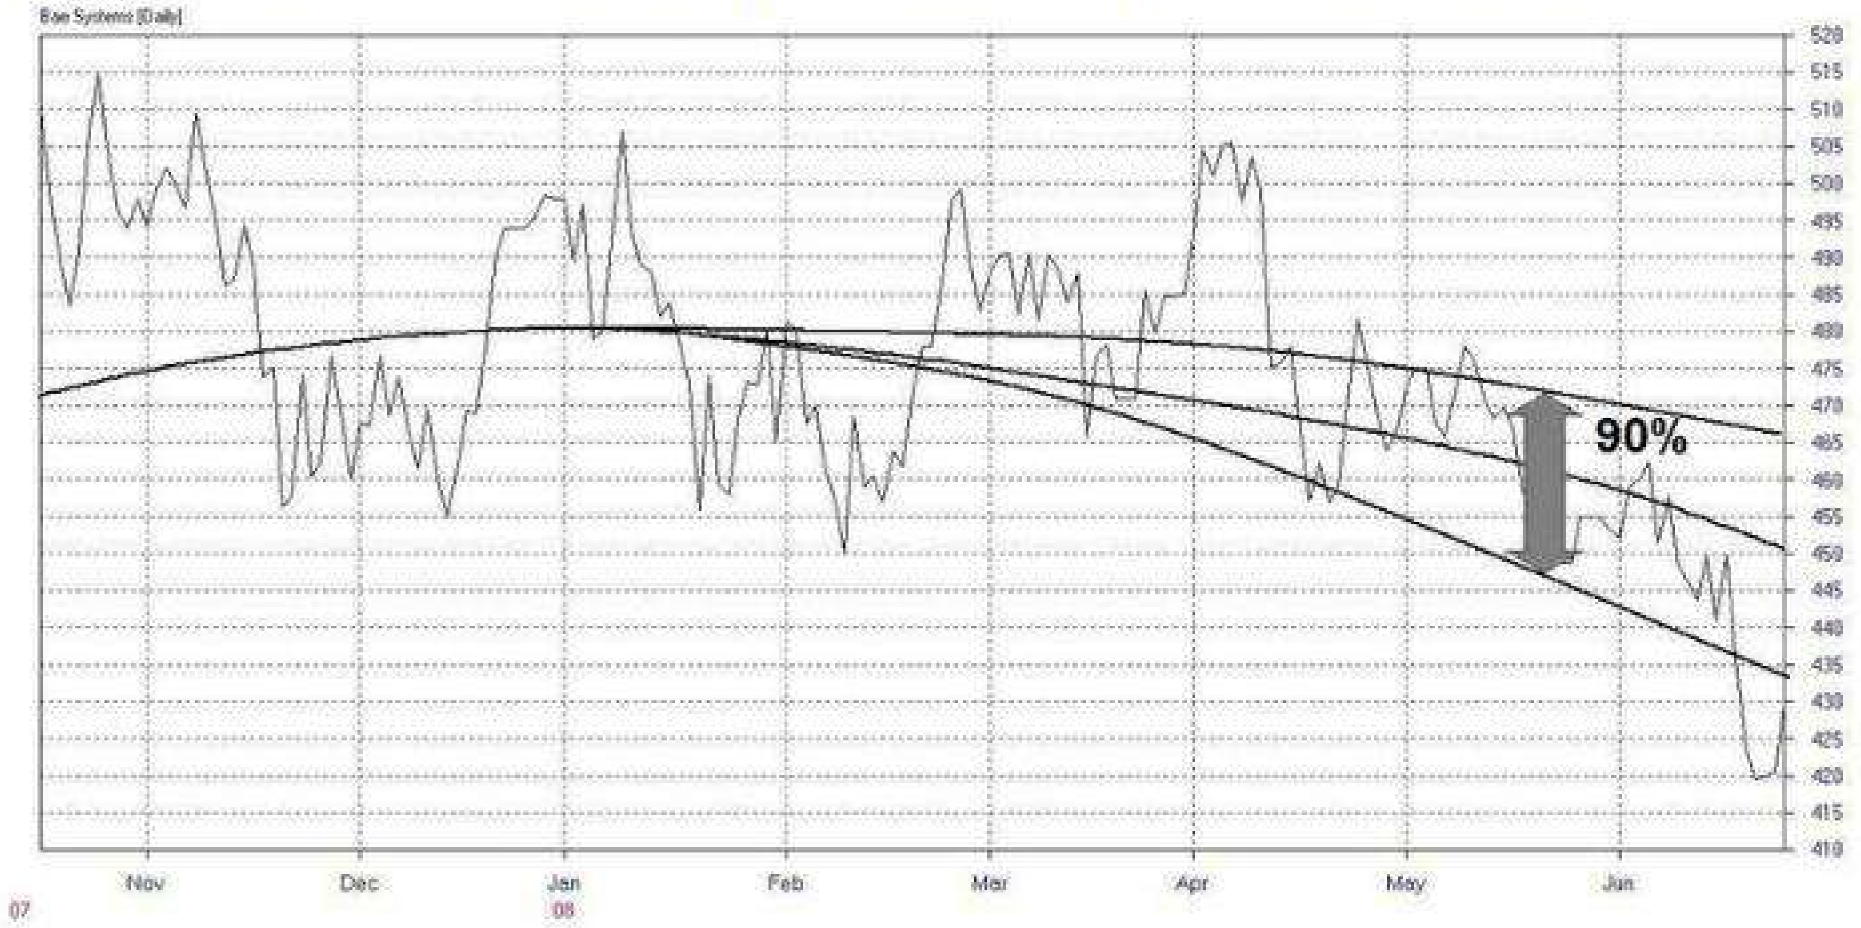


**In the example shown in Figure 1.1 there is a 90% probability that the trend line lies between the two positions indicated by the vertical arrow.**

### PRICE POSITION

Once we have determined the likely position of the trend, our attention will turn to the relationship between the value of the security, stock, index or foreign exchange ratio, and the estimated position of the trend. In my previous work on channel analysis I showed that the values of securities oscillate around a central trend and that there is a limit to this oscillation. This limit depends upon the timescale of the trend and differs from one security to another.

We will be able to draw probability boundaries that will give the probability of a price being at a particular place *relative* to a predicted central trend. This concept is illustrated in Figure 1.2.

***Figure 1.2 – Probability boundaries if we assume that the central trend is in the correct position. In this case there is a 95% probability that the price lies between the two boundaries.***

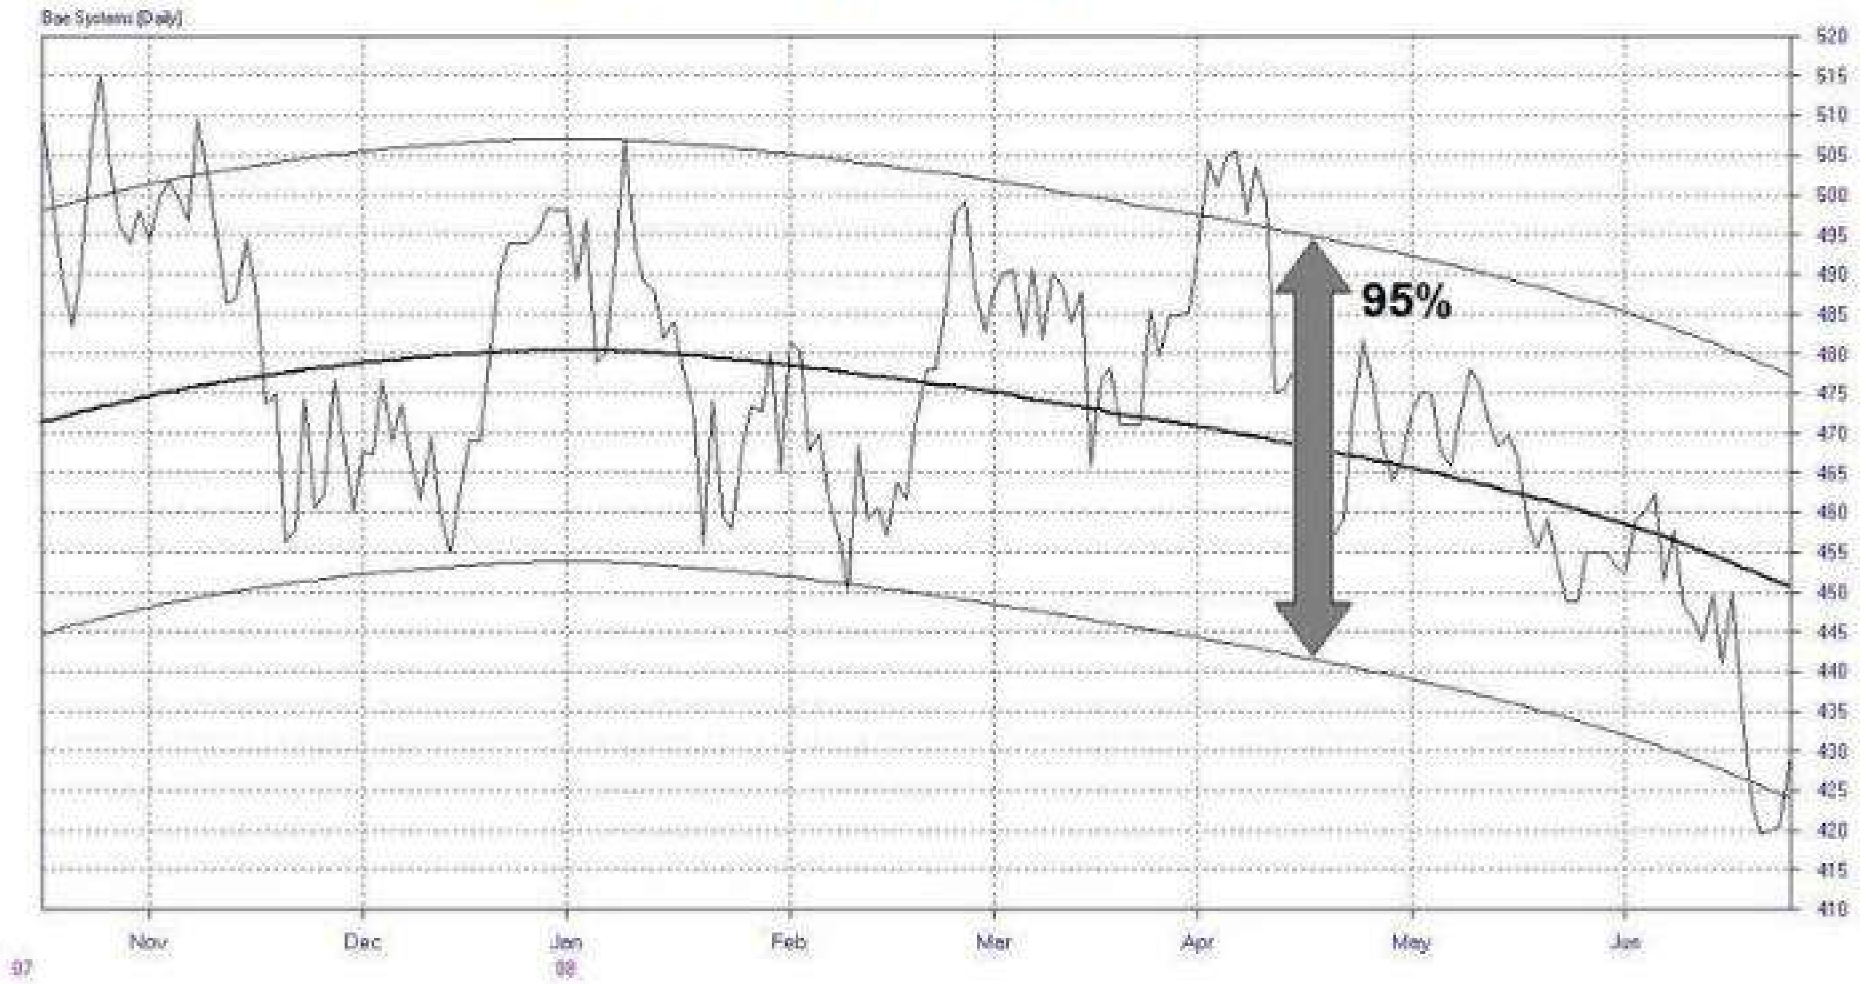


In the example shown in Figure 1.2 the boundaries are placed at a point where there is a 95% probability that the price lies between these two. In fact we will see later that the probability of the price being at an exact position relative to the boundary changes as we move between these two boundaries can be determined at least approximately. When we are at, say, the lower boundary there is a high probability (but *not* 100%) that the price will rebound upwards. This means that there is a low probability that the price will fall further. Conversely, when we are at the upper boundary there is a high probability (but *not* 100%) that the price will fall back. This means that there is a low probability that the price will rise further. These boundaries are at a constant vertical distance apart throughout the whole set of data. Thus they form a constant depth channel, which is the basis of *channel analysis*. The probabilities are calculated from the distances between the value of a centred moving average at each point and the price at that point. A more exact explanation of the change in probability across the channel will be given in a later chapter.

## COMBINING THESE TWO PROBABILITIES

It can now be seen that the prediction of the future movement of a price depends upon knowing the probable position of the future trend and the position of the price value within the constant depth channel constructed around this trend. If we use p(price) to be the probability of a future value for the stock price, p(trend) to be the probability associated with the trend position and p(relative) to be the probability of the price position relative to the channel boundaries, then it follows that:

p(price) = p(relative) x p(trend)

Based on a great deal of research into channel analysis, a reasonably accurate value can be obtained for the probability relative to the central trend (p(relative)). This is provided of course that an extensive amount of historical daily closing data is available for the particular stock, currency or commodity. The greater the amount of data, the more accurate is the value obtained. At least 1000 points must be considered to be the minimum required.

We will see later that at the present state of development of cycle prediction, we can arrive at only an approximate probability for (p(trend)), i.e. the future position of the central trend.

This is because, unlike the p(relative) calculation, which is based on real data and real values for the moving average which have been used, the future cycles are themselves estimates. They are based on the assumption that a stable cycle will remain stable for a reasonable time into the future. However, as we have mentioned, cycles will change both in amplitude and wavelength. This change will add uncertainty to the estimated future position of each cycle and therefore to the estimated sum of cycles.

## MORE ABOUT PROBABILITY

Probability is rarely discussed in books on technical analysis, most likely because it appears to be an academic topic which is difficult for the lay person to understand. However, this depends upon how the topic is addressed, and whether it is absolutely essential to the understanding of the rest of the book. I have taken the view that it can be addressed in a simplified way that will make it easy for the lay person to grasp the important fundamentals and that it is not necessary to concentrate on terrifying equations with integral signs and exponential terms to achieve this.

It should be stressed here that it is not essential that the reader should grasp clearly the subjects of statistics and probability in order to profit from this book. These topics are addressed purely so that the reader can understand that there is a mathematical underpinning of the issues that are discussed later. There are many texts available and many snippets on the internet that have the objective of making statistics and probability more easily comprehensible to the layperson. Thus readers may, if they wish, skip over those chapters which discuss probability. Of course, the calculations require data, and sources of data in the correct format for using in spreadsheets and at no cost to the user are widely available online. For readers who would prefer to access these probability calculations by a quicker route, they are also available in the Channalyze program.

If the trader does not remain disciplined and forgets about making sure that probability is on his or her side, then failure is the almost certain outcome. Just like the casinos in Las Vegas, in the long term you will win if the odds are on your side. It only takes a small shift in the odds in your favour to improve your performance out of all recognition. Conversely, a small shift the wrong way can lead you to despair. It is essential to monitor the trend prediction constantly to make sure that a cycle that will cause an adverse shift in these trends is not rapidly gaining in amplitude.

One final message – you will not be 100% correct in your decision making. There is no such thing as a cast iron, guaranteed profit in stock, commodity or currency markets.

## 2. Risk and the Markets

In my previous book (*Channels & Cycles*) I spent some time in the first chapter discussing money management. This is an important subject for traders and many have come unstuck owing to a failure to divide their capital into a number of parts which are traded in different stocks. I have always suggested dividing into eight, but certainly into not fewer than six. This means a total loss in one trade, while painful, can be accommodated. The discipline to stay with this partition of capital must be maintained at all times. Never put more than this one part into any one position, however positive the outcome might look at the time the trade or investment is being considered. Quite obviously the existence of risk means that there is never a trade which has 100% probability of success. Even those with a 90% probability of profit still have a 10% probability of failure and there will be many occasions in your life as a trader where the outcome that is only 10% probable is the one which occurs. You are bound to make a loss in some trades, so make sure you are not wiped out by it through lack of discipline.

The minimum strategy might be said to preserve capital at all times. This seems to be a negative approach since the aim of a trader or investor should be to increase capital over a period of time. A much more positive approach can be achieved by the quite simple strategy of applying probability to maximise profit and reduce losses.

## Sources of Risk

The risk to the performance of a stock comes from several sources, and the trader must be aware of these so that a rounded view can be taken at all times. The net performance of the stock is the result of all these influences. As far as the individual stock is concerned, there are a number of influences that can affect the stock price.

### INDIVIDUAL STOCK

There is the risk to the price of an individual stock due to a deteriorating performance of the company that it represents. This performance may involve a cut in the dividend, to which investors do not take too kindly. It might take the form of an announcement about future prospects with the usual comment about difficult trading conditions over the coming year. It might involve an obvious drop-off in sales, a failed takeover bid for another company or for the company itself, paying too high a price for another company, or the purchase of another company which is likely to under-perform in the immediate future. There will be other risks, but the above will serve to give a good overall view.

### MARKET RISK

The old Chinese saying 'a falling tide will maroon all boats' is very apposite when applied to the other risks: those of the sector, the whole market, worldwide markets and major political events.

The trader must keep more than a weather eye on these other influences. If the sector of which the stock is a member is becoming unfashionable, then it is unlikely, although not impossible, that any individual stock will outperform the sector. However, the balance of probability is that the odds are against any large increase in the stock price unless there are special circumstances such as a takeover that can be identified. Thus it will be useful to maintain, and indeed analyse, a chart of the sector index, just as one would analyse the individual stock.

It goes without saying that the major market index should be constantly monitored. If the market is headed strongly downwards, then it is unlikely that an individual stock will continue to move in the opposite direction for any appreciable length of time. This is easily checked by looking at the advances and declines in stocks in the financial pages the next morning after an above average fall in the market.

The other side of this coin is when a stock in which you are invested does not rise with the rest of the market. This is particularly galling, but can be mostly avoided if the principles shown in this book are followed.

#### *Link between individual daily stock prices*

It is interesting to examine the behaviour of individual stocks from different market sectors to see if changes in one stock are reflected by changes in another stock. In the London market the three stocks AstraZeneca, Jardine Lloyd Thompson and Tesco were examined. These are from very diverse sectors – pharmaceuticals, holiday travel and food retail. Thus in the absence of a stock market, it can be assumed that these stocks would be rising or falling quite independently of each other. If they are not totally independent of each other, then this must be due to the influence of the market itself upon overall investor sentiment, which manifests itself in a degree of synchronisation between stocks.

The closing prices of these three stocks were studied for the period from the beginning of 2003 to February 2009, a total of 1597 individual daily values. The daily changes of each of these were then calculated, giving 1596 such values for each stock. The assumption was then made that the daily changes within each stock were independent of each other.

The relationship between the daily change in one stock with each of the other two was examined. In other words, if the price of AstraZeneca fell on a particular day, did the prices of Jardine Lloyd Thompson and Tesco also fall? If the price of one rose on a particular day, did the prices of the others also rise? Logic would dictate that this is most unlikely, so the test boils down to: what proportion of the 1596 changes were in the same direction for two or all three of these stocks? The results are shown in Table 2.1.

***Table 2.1 – Daily changes in AstraZeneca, Jardine Lloyd Thompson and Tesco from 2003 to 2009.***

***Table 2.2 – Synchronised daily changes for AstraZeneca (AZN), Jardine Lloyd Thompson (JLT) and Tesco (TSCO).***

If the behaviour of these three stocks is totally independent, then we would expect the prices to move in the same direction 50% of the time. If only a small number of changes were taken, then the result could differ from the value of 50% due to the effect of probability. However, with a large number such as 1596 the actual percentage of synchronised movements should be close to 50%. The values of 40, 58, 59 and 63 are sufficiently removed from 50% that it can be considered that these movements are not totally independent, but show a small measure of dependence upon each other. Where does this dependence come from? It comes of course from the overall effect of the market.

While the effect of the market over the long term (a period of about six years was used for the study shown in Table 2.2) is not very large, over the shorter term it will be much more important. Large changes in the market will cause the majority of stocks to move in the same direction. The figures in Table 2.2 reflect the fact that over the long run, there are only a few such large changes in the market, the rest being of such low magnitude that they have only a minimal effect.

The effect of various changes in the market, as measured by an index, on the behaviour of individual stocks is illustrated in Figure 2.1. Here the change in the market index over various three-month periods is plotted against the change in the individual stocks over the same period.

***Figure 2.1 – The relationship between the performance of individual stocks and the market itself over three-month periods. The vertical axis is the percentage of rising stocks. This shows that once the percentage rise or fall in the index exceeds 10%, the majority of stocks follow suit.***

The relationship shown in Figure 2.1 is interesting – it is of course not a straight line. When the index has more or less stood still over a three-month period, then around half of the stocks rise and half of them fall over this same period. For changes of 10% or more in the index in a three-month period, then the majority of stocks follow suit in rising or falling. It is this relationship which makes it difficult to buck the market during periods of high volatility, such as that which we experienced during the 2007-2009 crisis.

**WORLD MARKETS RISK**

Investors in the major markets around the world should always take note of what the other major markets are doing. Thus although London might be doing very well up until 2:30pm UK time, a major fall on Wall Street at the opening will see London respond in a negative fashion. This can be very frustrating for UK investors who might find that the price of a stock that has no connection whatsoever with the United States suddenly goes into reverse simply because of the behaviour of Wall Street. The opposite is of course also true in that a large fall in London up until lunchtime will frequently cause a fall at the opening of Wall Street.

The behaviour of major world markets over the period from the end of October 2008 to the end of January 2009 is shown in Table 2.3.

***Table 2.3 – Change in some world markets over three months from the end of October 2008 to the end of January 2009.***

Whilst there is a wide variation in the percentage fall in the markets shown in Table 2.3, the fact is that they all fell over this period. Taking the general relationship from Table 2.3, it can be seen that the approximate percentage of rising stocks in each of these markets would be as shown in Table 2.4. The best of these markets was China, where only 68% of stocks fell during the period. The worst was the United States, where 90% of stocks fell during the period. In terms of probability, there is no doubt that the odds were stacked against the trader during this time. Naturally, within this period there will have been short-term rises in most of these stocks, with some rises being insignificant, while others would have been quite useful. A short-term trader could therefore have made profitable trades over a period when the overall change was unprofitable.

***Table 2.4 – Approximate percentage of rising stocks over the period from the end of October 2008 to the end of January 2009.***

The message from this is to tailor your trading horizon to match the timescale of the trend that has been identified. However, there is no doubt that if the medium-term trend is downward, then the profit from short-term trades will be far, far less than if the medium-term trend is upwards. This issue of the timescales of various trends is of course addressed in a later chapter.

Political events such as an interference with oil supplies can have a major effect on all markets. If such events happen, the trader who can quickly see the implication is the one who will benefit the most, provided he or she takes the appropriate action.

It is difficult to put a numerical value of probability on those risks which have just been highlighted as being outside of those due simply to the stock itself. However, in monitoring the stock price over a period of time we will be able to determine at least approximate probabilities of the range of future movement.

***London and Wall Street daily closes***

The same exercise as was performed on AstraZeneca, Jardine Lloyd Thompson and Tesco can be performed on the closing values of London (FTSE 100) and Wall Street (DJIA).

This is shown in Table 2.5.

***Table 2.5 – Daily changes in closing values for FTSE 100 and DJIA. The 'Next day' column is for synchronisation of the FTSE 100 Index with the previous day's close of the DJIA.***

In this exercise the value in points of the rise or fall has been ignored and only whether both markets have risen or both markets have fallen is noted. Thus out of 2296 daily changes, 1127 were synchronised. If pure chance was operating, then the expected number of synchronisations would be half of the total, i.e. 1148. The actual value of 1127 is close enough, given the large number of changes being taken into account, to arrive at the conclusion that there is virtually no synchronisation between daily rises and falls if the level of the rise or fall is ignored. Since Wall Street closes well after London, the final column notes the behaviour of London the following day, when the previous day's performance of Wall Street is known to London traders. The level of synchronisation is still that which would be expected from purely random behaviour.

However, just as large changes in a market will affect the majority of stocks in that market (as we saw in Figure 2.1), then we can understand how sometimes one market can influence another if enough traders are taking note of other markets. This can be evaluated by studying the effect of unusual rises or falls in Wall Street on the behaviour of the London market the next day. The situation is clouded rather by the overlap of the trading sessions. Thus if Wall Street falls one day by a large amount, then London will usually follow suit the next morning. However, if Wall Street then rises at the start of the next session, then London may recover some of the ground lost in the morning session.

Show in Table 2.6 are the rises and falls in London which follow the direction of Wall Street the previous day, versus the percentage rise or fall in Wall Street. This enables us to determine at what level a change in Wall Street is very likely to affect the London market the next day.

Until we get to a level of around 2.5% of movement in the DJIA, then the effect on the next day's movement in London is minimal. However, once we see movements of 3% and above we see an increasing effect on London stocks, with around three-quarters of stocks moving in the same direction.

The question then arises as to whether Wall Street is affected by a large change in London on the same day? This is shown in Table 2.7.

***Table 2.6 – Daily changes in the DJIA greater than the specified percentage and whether London moved in the same direction the next day. The data was obtained over 2300 days.***

***Table 2.7 – Daily changes in the FTSE 100 Index greater than the specified percentage and if Wall Street moved in the same direction the same day. The data was obtained over 2300 days.***

Tables 2.6 and 2.7 are interesting in that the number of days (out of 2300) that these various percentage changes occurred were very similar, e.g. for greater than 2% there were 216 days in the Dow and 222 days in the FTSE 100 Index. It also has to be acknowledged that traders in the US take less notice of what happens in London than traders in London do about New York. It is not until London movement gets to 4% and above that New York stocks start to see an effect, causing a larger percentage of them to move in the same direction.

In conclusion, being aware of the behaviour of other markets in the short term will help the trader to avoid placing buying and selling orders on a day when these external forces will increase the probability of an adverse move in the particular stock of interest.

# 3. How Prices Move (I)

In this chapter we will need to understand the concept of probability at a very simple level. This simple introduction will be very helpful in understanding the issues in Chapter 4.

Probability can be expressed either as a percentage or as a fraction. The difference is simply that a fractional probability is multiplied by 100 in order to arrive at the equivalent percentage probability.

In mathematics, a probability of an event X is represented by a number in the range from 0 to 1 and written as p(X). A probability of 0 means that the event is impossible, while a probability of 1 means that the event is certain. Thus if X is the event where a zero is thrown on a dice, then p(X) = 0; the event is impossible because there is no zero on a six-sided dice.

If we take a six-sided dice as an example, there is only one face out of six with a particular number such as a three. Thus the probability of throwing a three, p(3), is 1/6, which is 0.166 or 16.6%.

If we wish to link two events together, such as rolling two dice, the probability of a certain outcome is simply the probabilities of each single outcome multiplied together.

Thus, as we have stated, the probability of throwing a three on a dice is 1/6. The probability of throwing a three on the second dice is also 1/6. Thus the probability of throwing two threes is 1/6 x 1/6 = 1/36 or 0.0277 (2.77%). Quite low!

Mathematically, if we call the second event Y, then its probability is p(Y). Thus the probability of both events X and Y happening is given by:

p(X) x p(Y)

In general, therefore, if we wish to see the outcome of a certain number of events, then the probability of all of these happening is obtained by multiplying together all of the individual probabilities. Those readers who wish to exercise their minds further could check out the probability of drawing a flush in five-card poker (all five cards of the same suit). This is very low!

One important point – the probability of an event not happening (call it p(not X)) is (1 minus the probability of it happening). Thus, using our p(X) for the probability of the event happening, the probability of it not happening is given by:

p(not X) = 1 – p(X)

From this, we can see that the probability of not throwing two sixes will be 100 - 2.77 = 97.23%

Various probabilities of everyday occurrences are shown in Figure 3.1.

**FIGURE 3.1 – THE PROBABILITIES ASSOCIATED WITH THE OCCURRENCE OF VARIOUS OUTCOMES.**

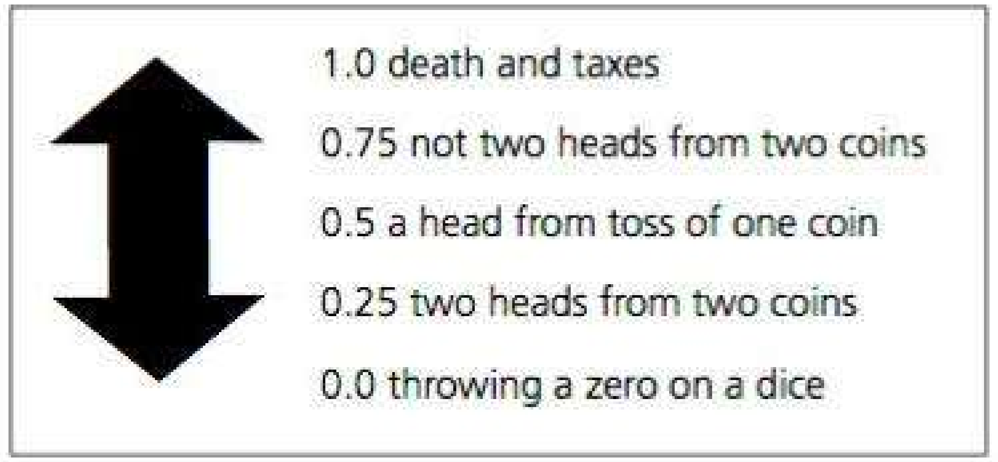

## Coin Tossing

Most people are confused by the probability of an event such as the next toss of a coin resulting in a head if there has already been a sequence of heads. In the toss of a coin each toss is independent of the previous toss, so that even if a sequence of ten successive heads has been achieved, the chance of the next throw being a head is still 50%. The coin, being inanimate, has no knowledge of its previous history.

### SEQUENCE OF TOSSES

The probability of the next toss achieving a head or a tail is not the same as determining the probability of a sequence of, say, heads or tails before the coin tossing experiment begins. The probabilities of runs of successive heads can be determined by multiplying together the probabilities of each event.

Thus, probability of tossing:

- 1 head = 0.5 (50%)
- 2 heads = 0.5 x 0.5 = 0.25 (25%)
- 3 heads = 0.5 x 0.5 x 0.5 = 0.125 (12.5%)
- 4 heads = 0.5 x 0.5 x 0.5 x 0.5 = 0.0625 (6.25%)

And so on.

Note that the probability of achieving, say, five heads is exactly the same as the probability of achieving four heads and then a tail. If the required outcome is three successive heads, then the expected proportion is 12.5%. As the number of tosses increases, then the proportion that is actually achieved settles towards the expected value. Since 1000 tosses is a reasonably large number, then the sequence of three heads should occur fairly close to 125 times in 1000 tosses. In other words the most likely value is 125 times, but since by their very nature tests based on probabilities do not give exact, repeatable outcomes then a value close to 125, rather than exactly 125, might be obtained.

It is possible to arrive at a probable range for the outcomes of such experiments using standard deviation. This is a measure of the scatter of a set of measurements around a mean value. The two words '*accuracy*' and '*precision*' can be used in this context. Accuracy refers to how close the mean of a set of results is to the true value, while precision refers to the scatter about the mean value. It is better to be roughly accurate than precisely wrong. A set of measurements with high precision will have a small standard deviation. However, it cannot be deduced from this that the mean of the results is closer to the true value than the mean of another experiment in which the standard deviation is higher.

In coin-tossing experiments, the larger the number of tosses, the closer will the end result be to the true result. This can be seen from the formula for standard deviation in coin-tossing experiments:

standard deviation = 0.5 * square root (N)

where N is the number of tosses. Thus for the number of heads in 1000 tosses, the standard deviation is 15.8. From the properties of a normal distribution (see later) this means that 95% of the occurrences of heads will lie within a range of 16 (15.8 rounded up) either side of the expected value of 500. Note that the standard deviation from sets of numerical data is obtained by a more complex formula, but that it is not necessary to apply such a formula since the standard deviation is easily obtained by using a spreadsheet program (this function is usually named STDEV) or many technical analysis programs.

As far as stock markets, currency markets or commodity markets are concerned, in this book we are only concerned with the probability of one event occurring over a specified period of time and that is whether our security is going to rise or fall. We shall see in a later chapter how to obtain probabilities from moving averages. Even though we use these in quite a simple way, they will be the rock on which our trading strategy of predicting trends in these markets is built.

## Rising and Falling Trends

A rising trend can be considered to be a sequence of changes after which the end value is higher than the starting value, a falling trend one in which the end value is lower than the starting value, and a sideways trend one in which the end value is the same as the starting value. There has to be a timescale attached to these trends, the time being the number of changes in the sequence. In this book we are not concerned with intra-day data, so all our data is based on daily values.

There are two aspects of trends that can be considered. A rising trend will occur if we have a sequence of successive rises, irrespective of the amount of each rise. In other words, only the direction is important. A rising trend will even occur if there is not a sequence of successive rises, but if the sum of the individual rises is greater than the sum of the individual falls. The opposite applies to falling trends. We can study these two aspects using simple applications of probability.

### SEQUENCE OF RISES AND FALLS

If we take the data for AstraZeneca – which was used earlier to study the correlation between movements in AstraZeneca, Jardine Lloyd Thompson and Tesco – we will see that there were 784 rises, 749 falls and 63 no changes out of 1596 daily changes.

The question arises as to whether we can predict with any certainty whether the next day's close will be higher, lower or unchanged from the day before. Of course, we addressed in Chapter 2 the issue of large changes in Wall Street having an effect on London stocks the next day, so that in these special circumstances we saw that we could be correct around 75% of the time. Later in this book we will see how to predict whether a short-term trend is expected to change direction. However, it will become clear that we cannot pin this change of direction down to the exact day. We will only be able to confirm the change once it has happened and we certainly should not open a trade until we have this confirmation.

It is helpful in this discussion to examine more closely these rises and falls in AstraZeneca. Note that there is nothing special about AstraZeneca (apologies to any readers who are holders of that stock!) – the same conclusions will be reached from a study of any other stock, currency or commodity from any other market.

In simple terms, we can say that since there were 784 rises out of 1596 daily changes, the probability of a rise after the period covered by this set of data would be 784/1596 = 49%. To be more correct, we should say that the likelihood of a rise is 49%. Just as in the simple coin-tossing experiment we would only get approximately 50% heads out of a trial of 1000 coin tosses. This is because the issue of standard deviation comes into the picture. Thus the actual probability lies somewhere within a range of values around this value of 49%. However, for our purposes, taking just the likelihood of a rise is sufficient for us to be able to take the discussion forward.

The value of 49% means that there is a more or less equal chance that the price will rise or the price will fall.

We can calculate probabilities from the data as follows:

- Probability of a rise = 784/1596 = 0.494
- Probability of a '*not rise*' = 1 - 0.494 = 0.506
- Probability of a fall = 749/1596 = 0.469
- Probability of a '*not fall*' = 1 - 0.469 = 0.531
- Probability of no change = 63/1596 = 0.039

- Probability of change = 1 - 0.039 = 0.961

From these probabilities we can, just as in the case of the coin-tossing experiment, calculate the probabilities of seeing isolated successive rises, isolated successive falls or isolated successive '*no changes*' in the history of daily changes. By isolated we mean that a rising sequence is ended by a fall or '*no change*' movement, a falling sequence by a rise or '*no change*' movement and a '*no change*' sequence by a rise or fall. It is important to study isolated sequences; otherwise the issue would be confused, since a sequence of four successive rises would also contain two sequences of three successive rises.

Thus, the probability of four successive rises, started by a '*not rise*' and ended by a '*not rise*':

= 0.506 x 0.494 x 0.494 x 0.494 x 0.494 x 0.506 = 0.0152

Since we have 1596 daily changes, this should lead to 0.0152 x 1596, which gives 24 such sequences.

The probability of four successive falls:

= 0.469 x 0.469 x 0.469 x 0.469 x 0.531 = 0.0257

Since we have 1596 daily changes, this should lead to 0.0257 x 1596, which gives 41 such sequences.

The probability of four successive '*no changes*':

= 0.039 x 0.039 x 0.039 x 0.039 x 0.961= 0.000002

This gives a probability of effectively 0 for such sequences.

A careful analysis of successive rises, falls and no changes in AstraZeneca gives the results shown in Table 3.1. The sequence of one covers the case where there is an isolated rise, fall or no change. In a sense this would not be a sequence, but is still of interest and so has been included.

Table 3.1 should be compared with the expected values, calculated as shown above for a sequence of four rises, falls and no changes. These are shown in Table 3.2.

***Table 3.1 – Sequences of isolated rises, falls and no change in AstraZeneca over a period of 1596 daily changes.***

***Table 3.2 – Expected sequences of isolated rises, falls and no change in AstraZeneca over a period of 1596 daily changes.***

The actual sequences of two or more are so close to the expected values that it must be considered that the next day's move is independent of the previous day's move.

The implication is therefore that sequences of rising or falling daily changes are almost totally random. If these sequences are the only driving force for uptrends or downtrends then it would follow that trends are random and therefore traders may as well pick their positions with a pin. However, the other dimension to a trend is the extent of each upward or downward move, which will be investigated later in this book.

One last aspect of sequences of rises and falls to investigate is whether there is any relationship between the sequence of rises or falls and the total change in value over the length of time for which the sequence holds.

The relationship between the total rise and length of rising sequences is shown in Figure 3.2. Also drawn on the plot are the straight-line trends of these data. It can be seen that the rises are a good fit to this straight-line trend and therefore the price rise is essentially directly proportional to the length of the sequence. In other words, if we break down rising sequences into an average change per day, then long sequences do not give a different rise per day from that of short sequences.

*Figure 3.2 – The relationship between total price rise and sequence length.*

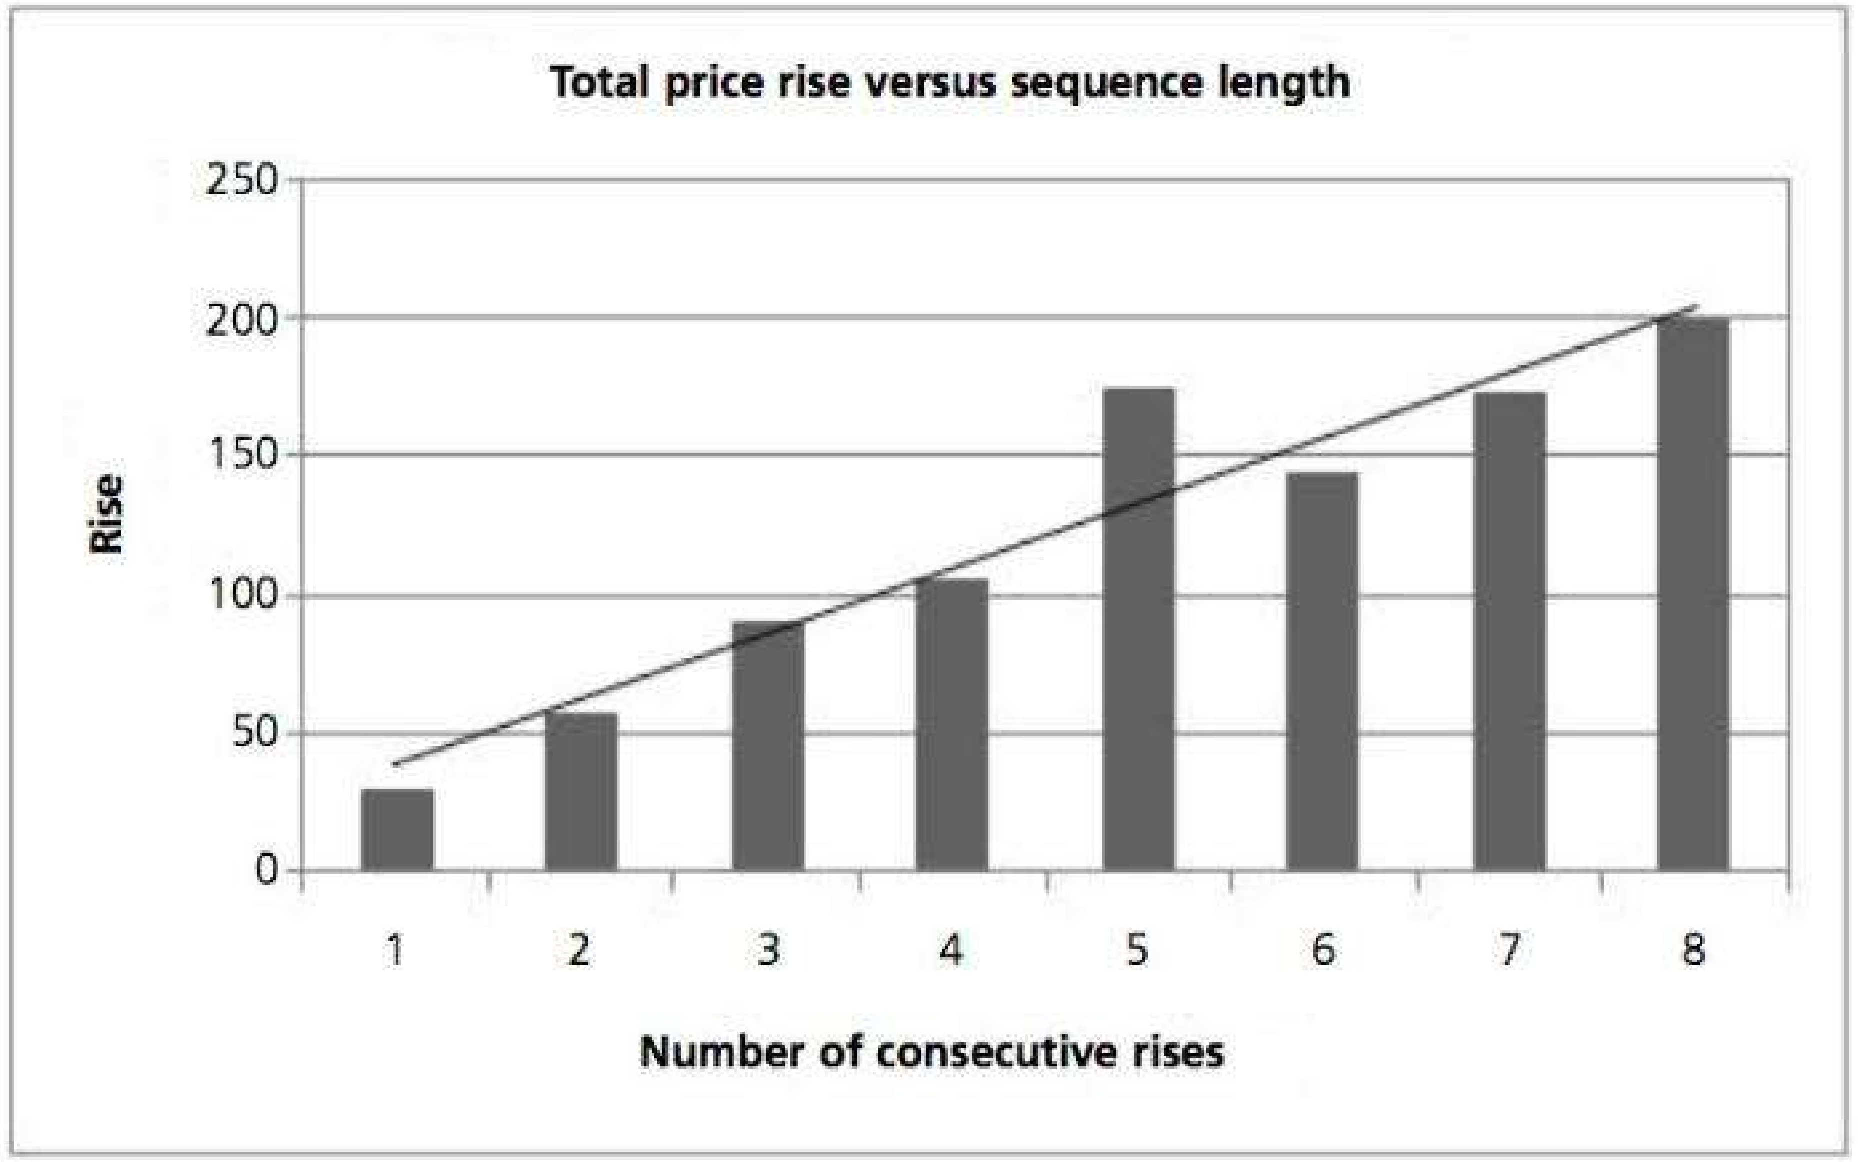


The relationship between the total fall and length of falling sequences is shown in Figure 3.3. Also drawn on the plot are the straight-line trends of these data. It can be seen that the falls are a good fit to this straight-line trend and therefore the price fall is essentially directly proportional to the length of the sequence.

In other words, if we break down falling sequences into an average change per day, then long sequences do not give a different fall per day from that of short sequences.

*Figure 3.3 – The relationship between total price fall and sequence length.*

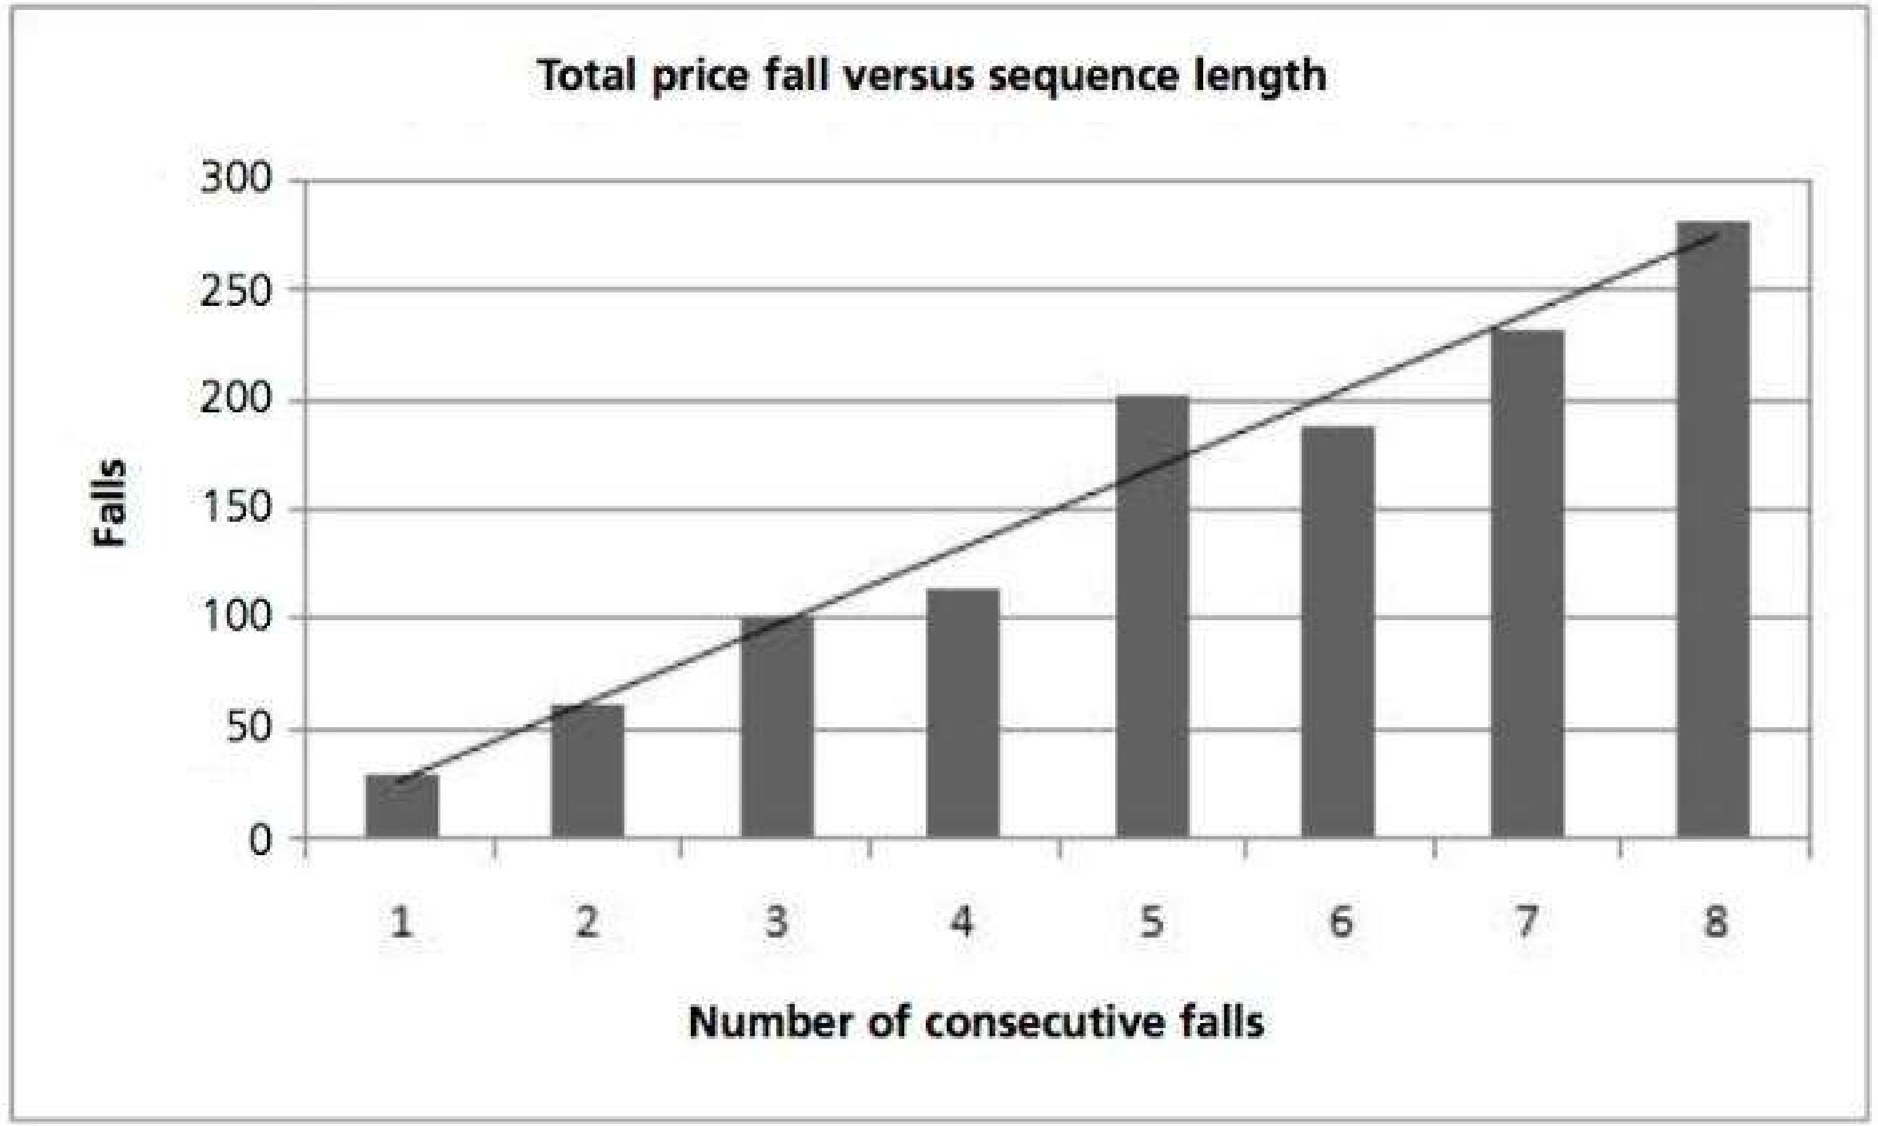


The results of the studies in this chapter would lead to the conclusion that price changes are random, since the number of successive rises or falls in a sequence are simply those that would be expected if there is an equal probability of a rise or fall each day. However, the next chapter will show that when it comes to putting values on daily changes rather than simply examining the direction of the change, we will see that they do not conform to a random distribution and therefore the prediction of future price levels is still possible.

# 4. How Prices Move (II)

The application of probability to sequences of rises or falls, even though of interest, was a rather trivial use of probability. When we come to the numerical values of these changes, we need to examine how probabilities can be arrived at by studying the distribution of price changes. By distribution we mean an analysis of the way in which the data is scattered about the mean value. It is necessary to introduce the statistical concept of the *normal distribution* in order to understand what can be deduced from a study of the actual value of the daily changes.

Note that it is not necessary to carry out laborious calculations – this can be done quite simply in a spreadsheet.

## The Normal Distribution

As an example, the average height of men in the United States is 69.1 inches (175.5cm) with a standard deviation of 2.9 inches (7.4cm). Quite obviously, there are men shorter and men taller than this, but a survey of the population would find that the vast majority of men would have heights close to this value. As we move away from this average value, the numbers of men with that particular height gets fewer and fewer.

The plot of numbers of men in a particular height category would appear as shown in Figure 4.1. The data falls into what is called a normal distribution around the mean of 69.1 inches. The *standard deviation* is a measure of the scatter of the data around the mean value. A fundamental property of the normal distribution is that:

1. 68% of the data falls within +-1 standard deviation of the mean
2. 95% of the data falls within +-2 standard deviations of the mean
3. 99.7% of the data falls within +-3 standard deviations of the mean.

Using the value of the standard deviation given in inches, this means that 68% of the men in the United States have heights in the range 66.2 to 72.0 inches and 95% of them have heights in the range 63.3 to 74.9 inches.

Since standard deviation is easily calculated in a spreadsheet such as Excel (using the STDDEV function), then the value of the normal distribution is that we can now state quite simply the limits between which 68% or 95% or 99.7% of the data will lie.

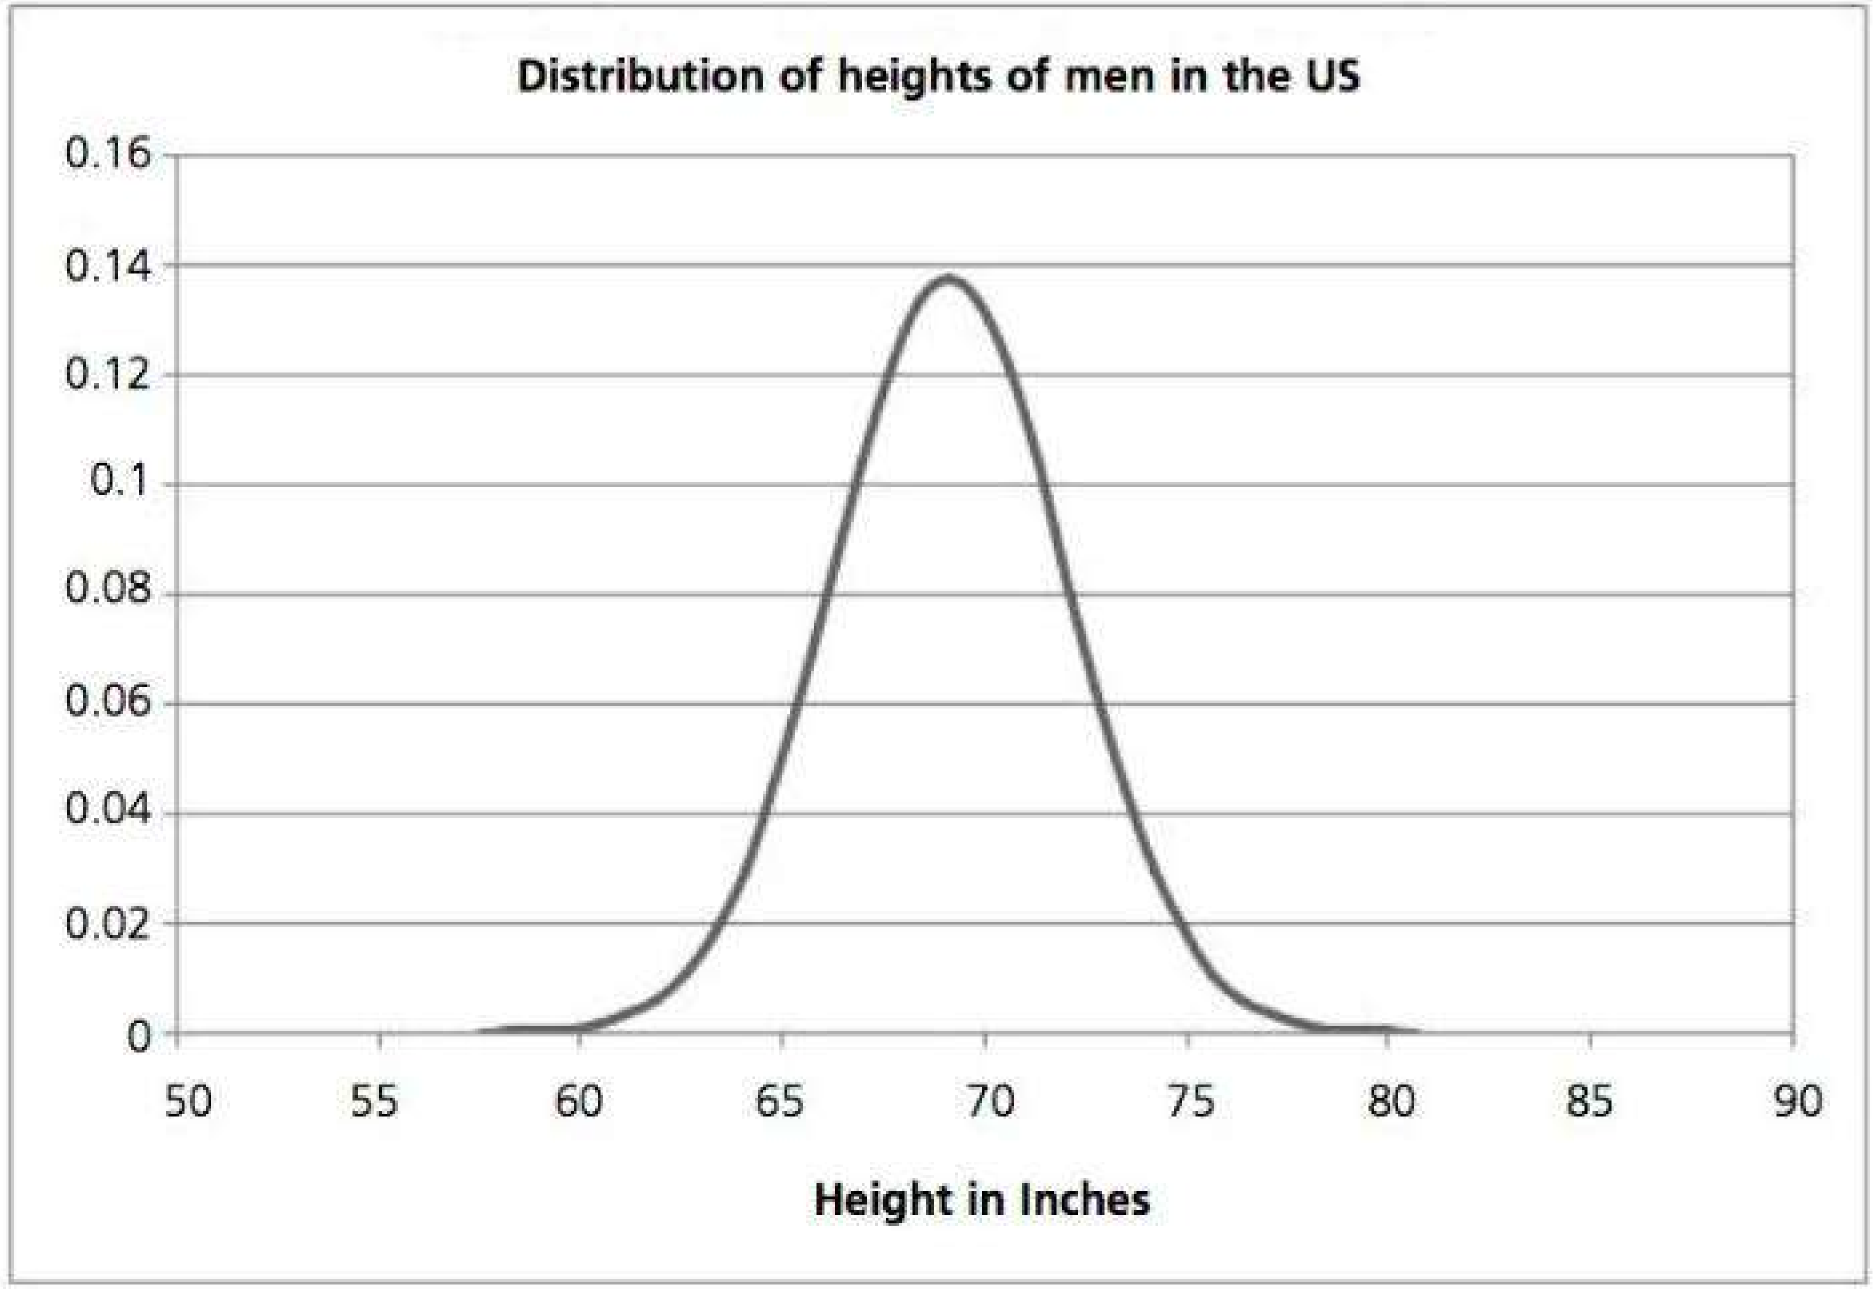


**Figure 4.1 – This shows a plot of heights of men in the United States. The average height is 69.1 inches and the standard deviation 2.9 inches.**

As we will see shortly, these limits are derived from areas under the distribution curve. They represent probabilities. Thus in this example the probability is that, in the United States, the next man you see in the airport or train or bus station will have a height of between 66.2 and 72 inches is 68%. The next time you are stuck in one of these public places, amuse yourself by taking your own mental and approximate survey.

Of course we have to be sure that the data we are examining are normally distributed. The simplest way is to plot the distribution and then see if it resembles a normal distribution such as that shown in Figure 4.1.

## THE STANDARD NORMAL DISTRIBUTION

What is called the standard normal distribution is a set of data with a mean of 0 and a standard deviation of 1. A plot of this data is shown in Figure 4.2.

**Figure 4.2 – The standard normal distribution curve has a mean value of 0 and a standard deviation of 1. In this plot the x-axis represents standard deviations from -5 to +5. The shaded area represents the 95% probability limits for a future value.**

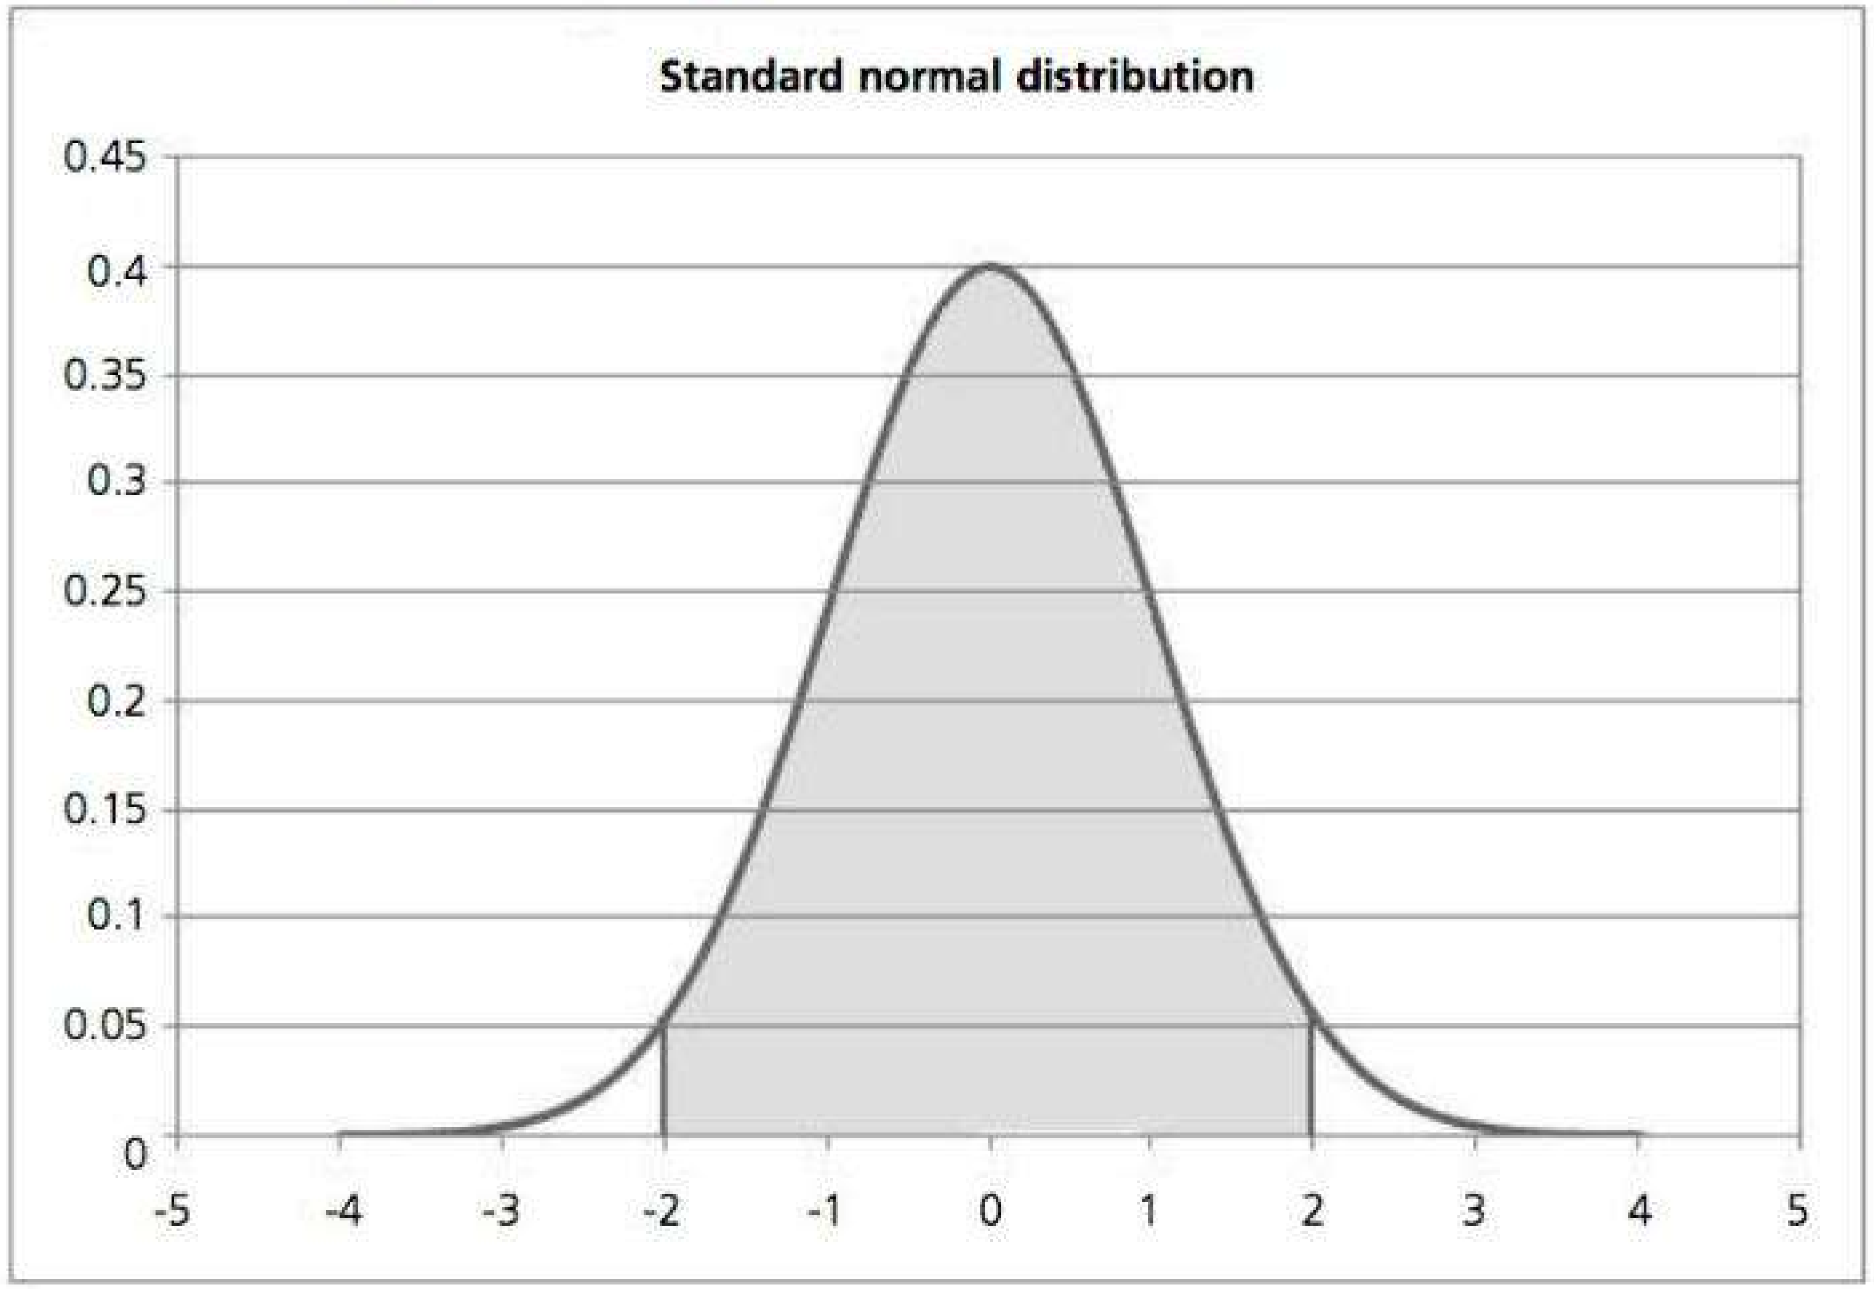


This is mentioned here because a value, x, from a normal distribution specified by a mean of μ and a standard deviation of σ can be converted to a corresponding value, z, in a standard normal distribution with the transformation $z = (x - \mu)/\sigma$. And, of course, in reverse, any value from a standard normal graph, say z, can be converted to a corresponding value on a normal distribution with a mean of μ and a standard deviation of σ by the formula $x = \mu + z*\sigma$.

Thus readers who wish to compare actual distributions to normal distributions can bring both the same scales. We will see the value of this when discussing the spread of data around a centred average in Chapter 7.

However, there is a highly valuable property of the normal distribution that allows us to predict the range in which a price should move on any particular day. We mentioned earlier that:

1. 68% of the data falls within +-1 standard deviation of the mean
2. 95% of the data falls within +-2 standard deviations of the mean
3. 99.7% of the data falls within +-3 standard deviations of the mean.

This can be converted to a probability that any particular day's change will be within those limits if it is assumed that the prices are normally distributed. The shaded area in Figure 4.2 represents the 95% probability that a future price will lie between these limits. All we need to calculate this is the mean value and the standard deviation computed from a reasonable amount of historical daily changes.

## CONVERTING DISTRIBUTIONS TO PROBABILITIES

The probabilities we have used so far have involved simply the probability levels of 68%, 95% and 99%. How do we find the probability of data lying between two values, or being greater or less than a specified value? The answer lies in the fact, which we have already mentioned, that it is the area under the distribution curve that gives us the probabilities. The probability of a data point lying between two values is obtained by taking the area under this curve as a proportion of the total area under the curve. Thus in Figure 4.2, the shaded portion under the curve lies between the values of +- 2 standard deviations. The area of this is exactly 95% of the total area under the curve. We therefore need a way of converting the normal distribution curves to a more usable probability curve. This is done by plotting the cumulative area as we move from the left to the right of the plotted data.

Calculating the area under the normal distribution curve is not a trivial process and involves a mathematical integration. However, a good enough approximation can be obtained by using the bin technique in a spreadsheet program or in Channalyze. This involves dividing the range along the x-axis into a number of equal segments (bins). The corresponding y value for each bin is then calculated for the normal distribution. All of these separate values can be summed to give a good approximation of the total area and the sum between two bin values will give the partial area of interest.

This process is illustrated in Figure 4.3. In this case 32 bins have been used to cover the range from -4 to +4 standard deviations. Quite obviously, the greater the number of bins into which we divide the x values, the closer will the end result approach the true profile of the normal distribution.

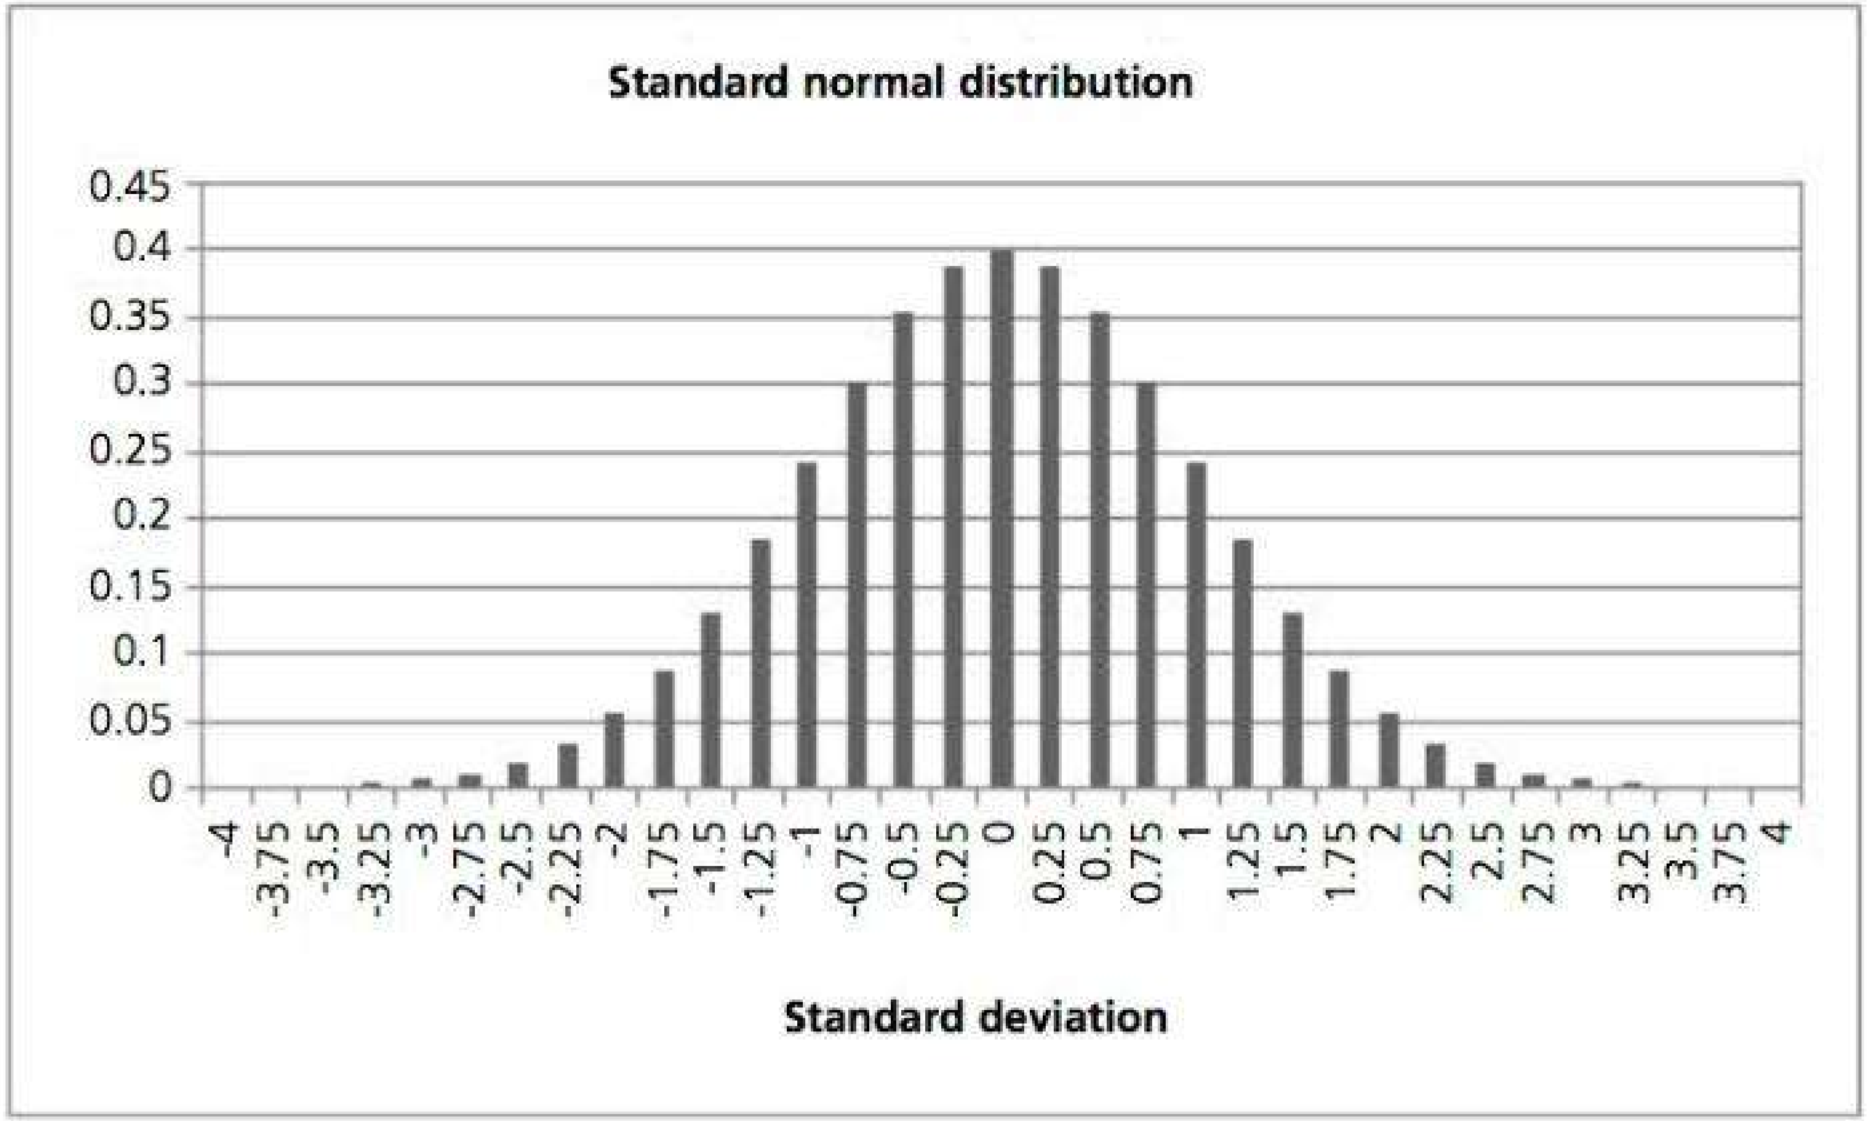


**Figure 4.3 – The area under the distribution curve can be approximated by adding up the values in each of the bins. The bins are of width 0.25 of a standard deviation.**

If we build a running total of the values in the bins as we move from left to right of this plot, and plot this total against the bin number, then we get a good approximation of the cumulative distribution. This will give us the plot shown in Figure 4.4.

**Figure 4.4 – The cumulative distribution derived from the normal distribution shown in Figure 4.3. The y-axis now depicts the probability, on a scale of 0 to 1, that a point from the distribution is greater than a value on the x-axis. At the value of 0 for standard deviation the probability is 0.5 (50%), i.e. there is an equal chance of the point being greater or less than this value.**

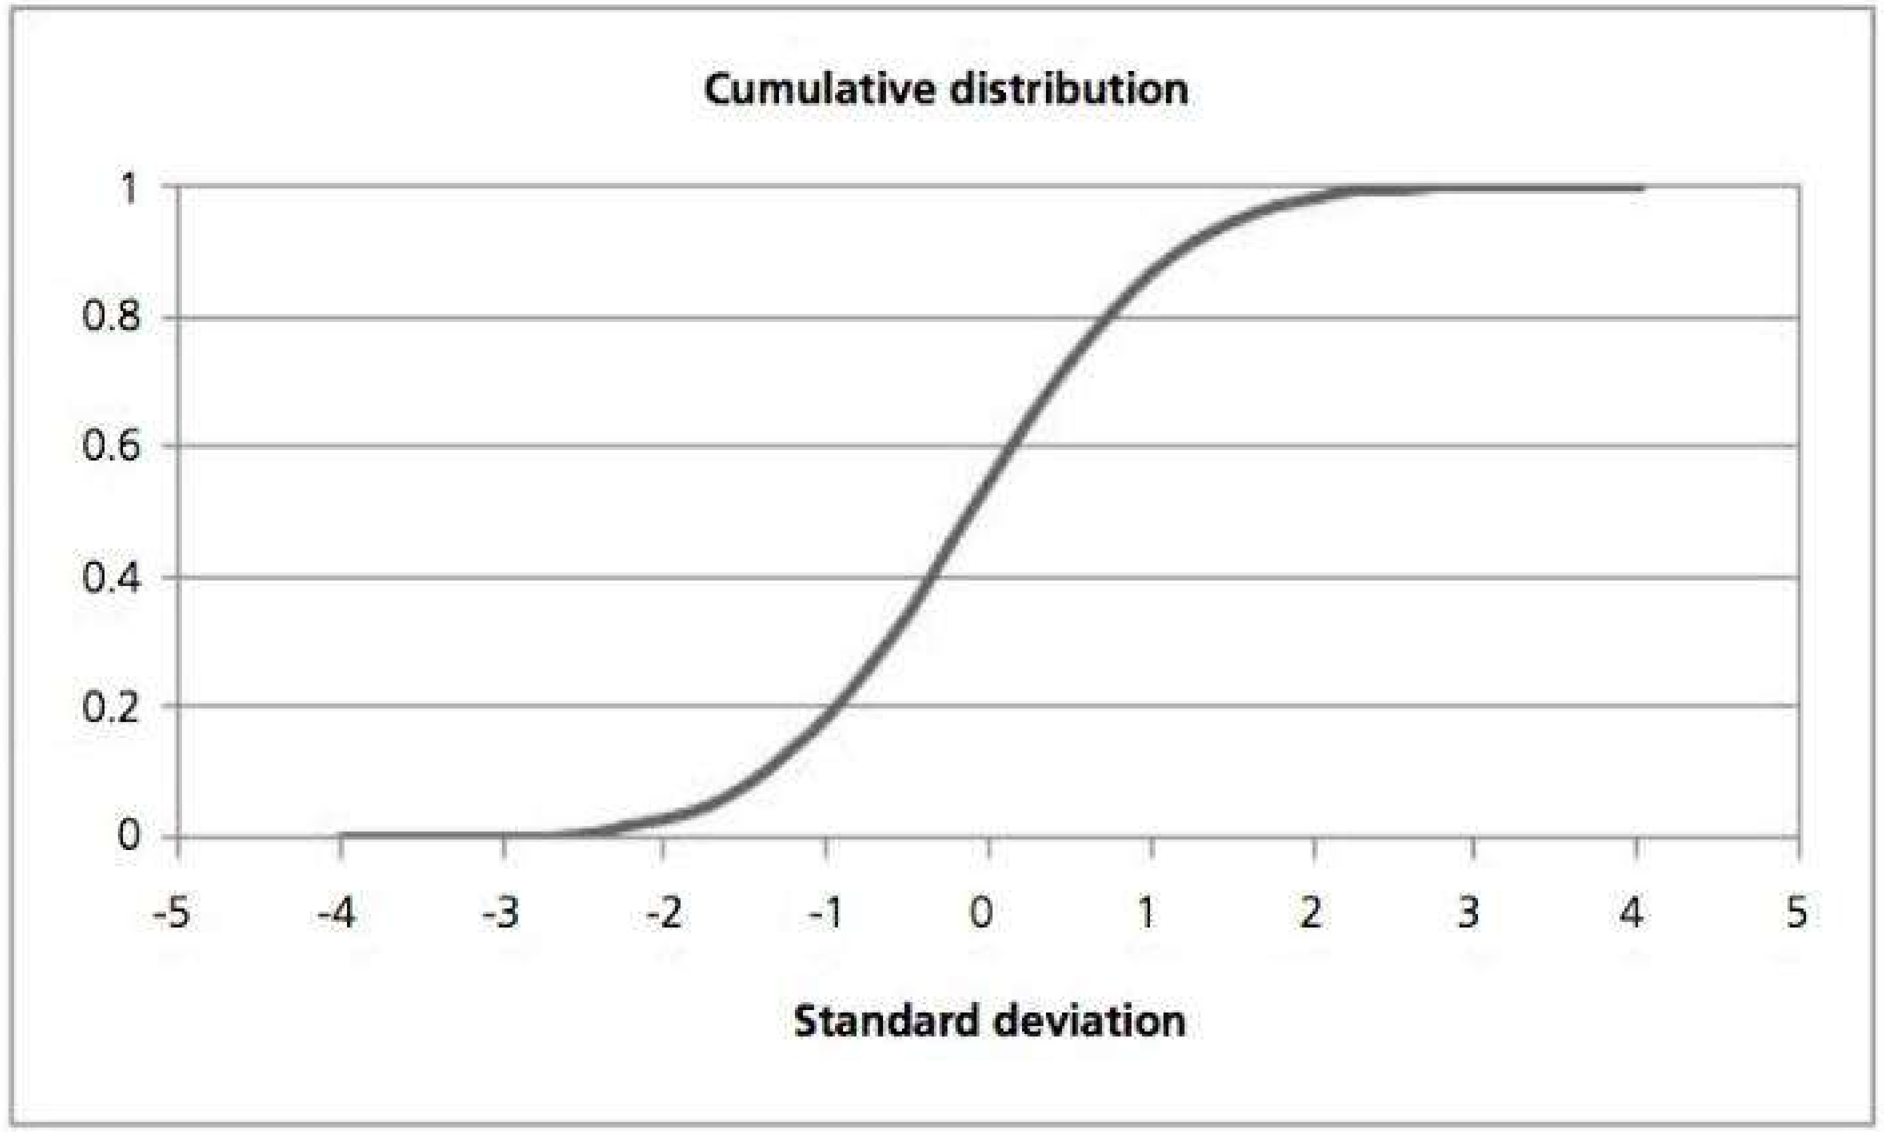


For any value along the x-axis, the corresponding y-value will give the probability that a new point from the distribution will be less than the value on the y-axis. By the time we reach the mid-point at 0 standard deviation, this probability is 0.5 (i.e. 50%), i.e. there is an equal chance of the new point being higher or lower than the mean value, which in this case is of course 0.

## THE NORMAL DISTRIBUTION OF STOCK PRICE CHANGES

We have already looked at the sequences of daily rises and falls in the AstraZeneca stock price. As far as the value of these changes is concerned, they range from a fall of 227p to a rise of 304p. The average value is -0.12, which for all intents and purposes can be taken to be 0. The standard deviation, which is a measure of the scatter of these changes around the average value, is 41.36.

Using the same method as for the normal distribution shown in Figure 4.3, i.e. just using the mean value and standard deviation and making the assumption that the data is normally distributed, we get the plot shown in Figure 4.5.

Since this is a plot of the normal distribution of AstraZeneca, its shape will be exactly the same as the distribution shown in Figure 4.3. The values on the x-axis will be different, since the mean is not 0 (it is -0.12) and the standard deviation is not 1 (it is 41.36).

**Figure 4.5 – The normal distribution of AstraZeneca daily price changes based solely on the mean value of -0.12 and standard deviation of 41.36.**

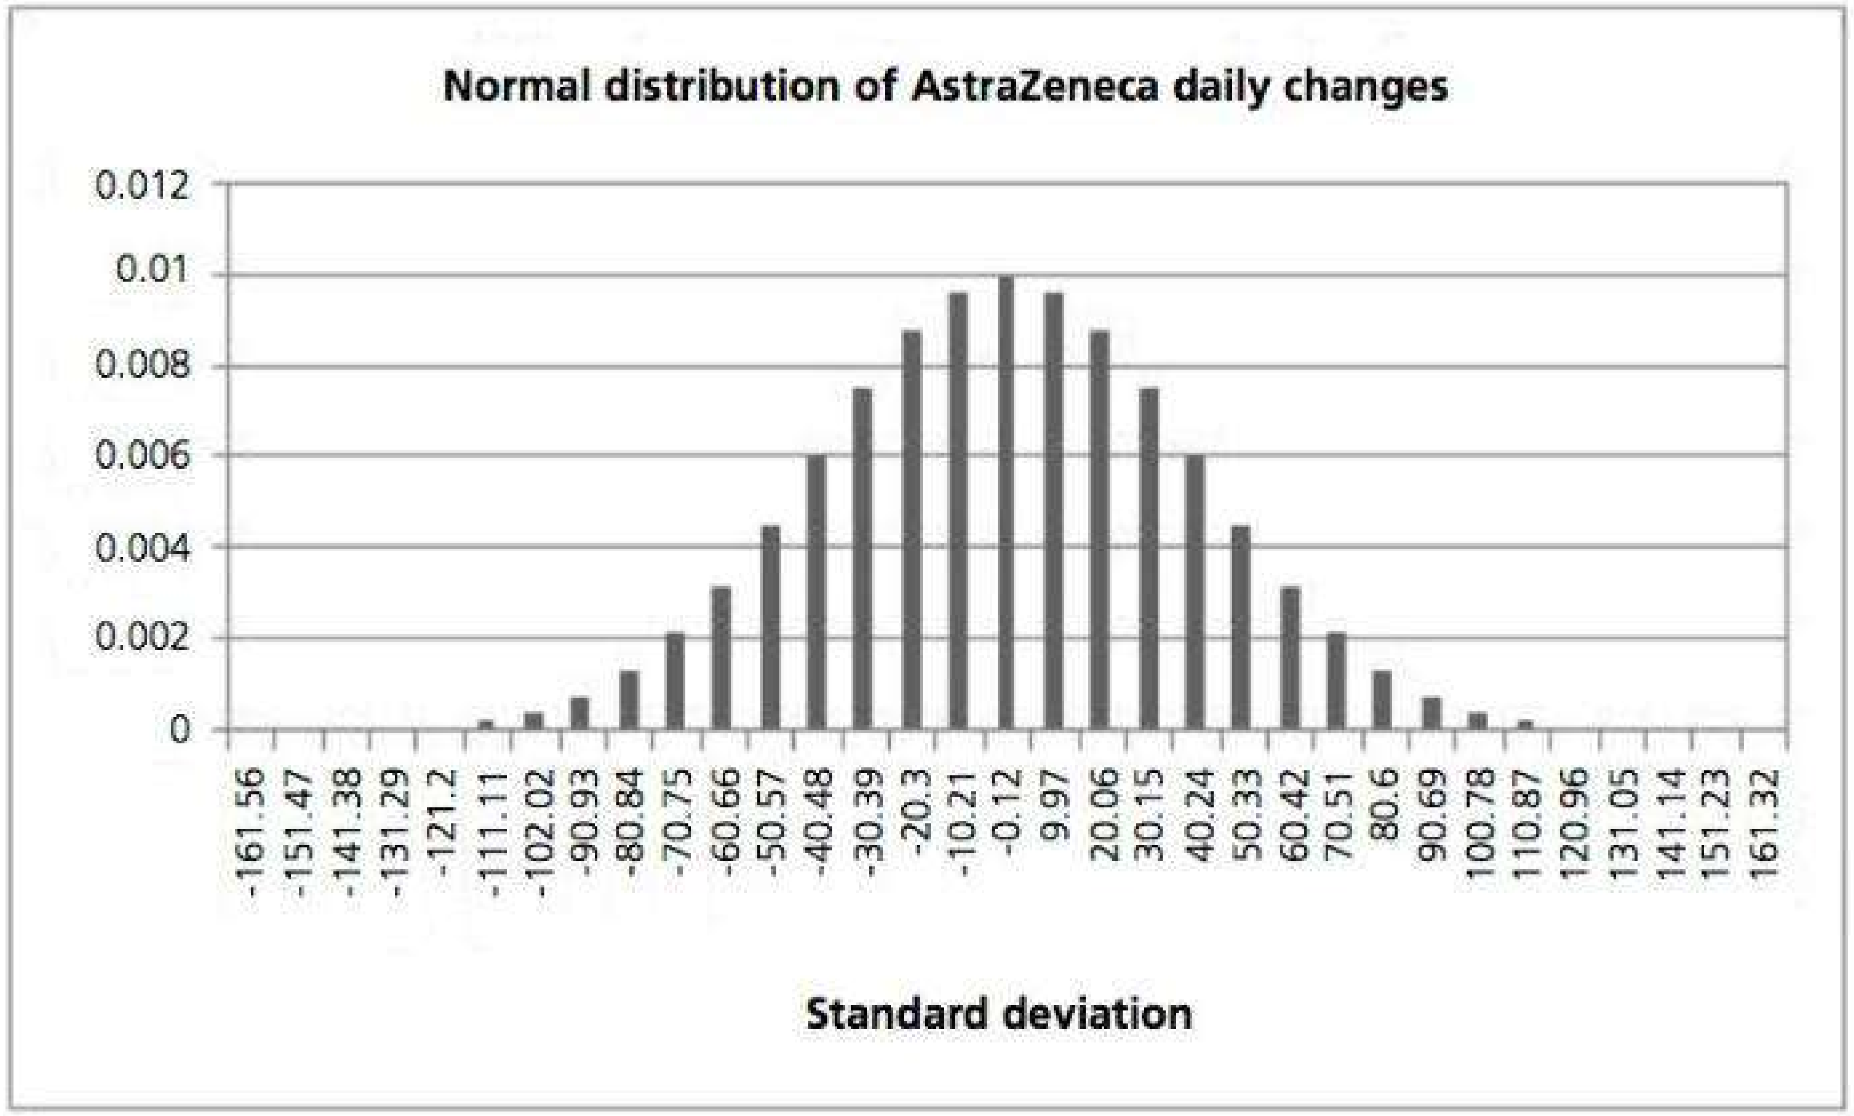

The cumulative distribution is the absolute key to determining the probability of a future point being less than or more than a certain value or indeed the probability of the future point being inside or outside of a range of values provided the data are normally distributed. The Channalyze program has a module which plots both the normal distribution and the cumulative distribution at the same time and allows the probability of any price change to be determined. This is shown in Figure 4.6. Here the values have been distributed into 60 bins and the cursor has been placed over the value for bin 26.

**Figure 4.6 – The normal and cumulative distributions from AstraZeneca daily changes. Moving the cursor over the cumulative plot gives the probability that the price change will be less than the indicated value.**

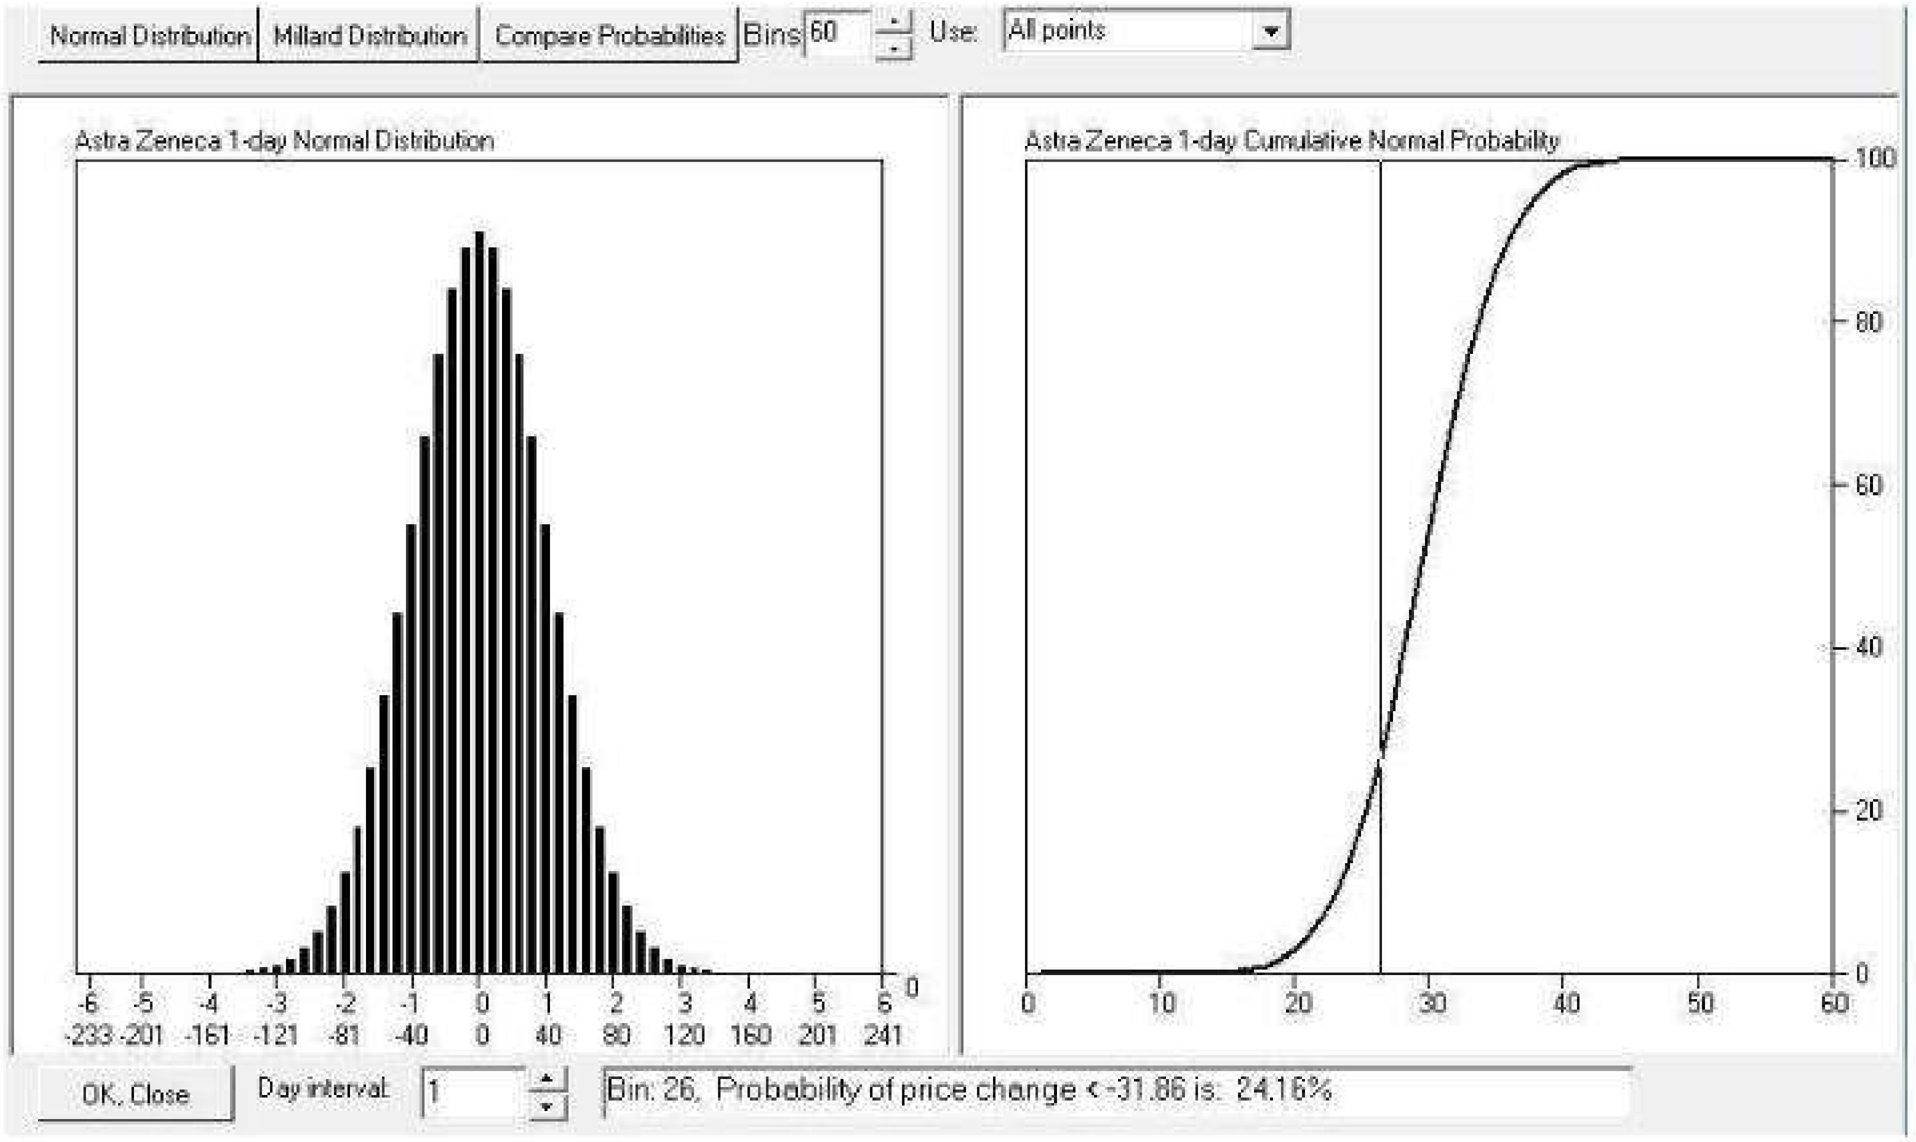


Bin 26 corresponds to a price change value of -31.86 and the corresponding probability is 24.16%. Since the values in the bin have been summed from the left-hand side, this means that the probabilities are for a future price change to be less than the indicated value. The probability that a future price change will be higher is obtained by subtracting this probability from 100, giving a 75.84% probability that the price change will be greater than -31.86.

## THE ACTUAL DISTRIBUTION OF STOCK PRICE CHANGES

### *Daily changes*

If we plot the distribution of the AstraZeneca daily changes by putting them into bins of width 10, running from -230p to 310p, we get the plot shown in Figure 4.7.

**Figure 4.7 – The actual distribution of daily changes in AstraZeneca. The changes are allocated to bins of 10p running from -230p (fall) to +310p (rise).**

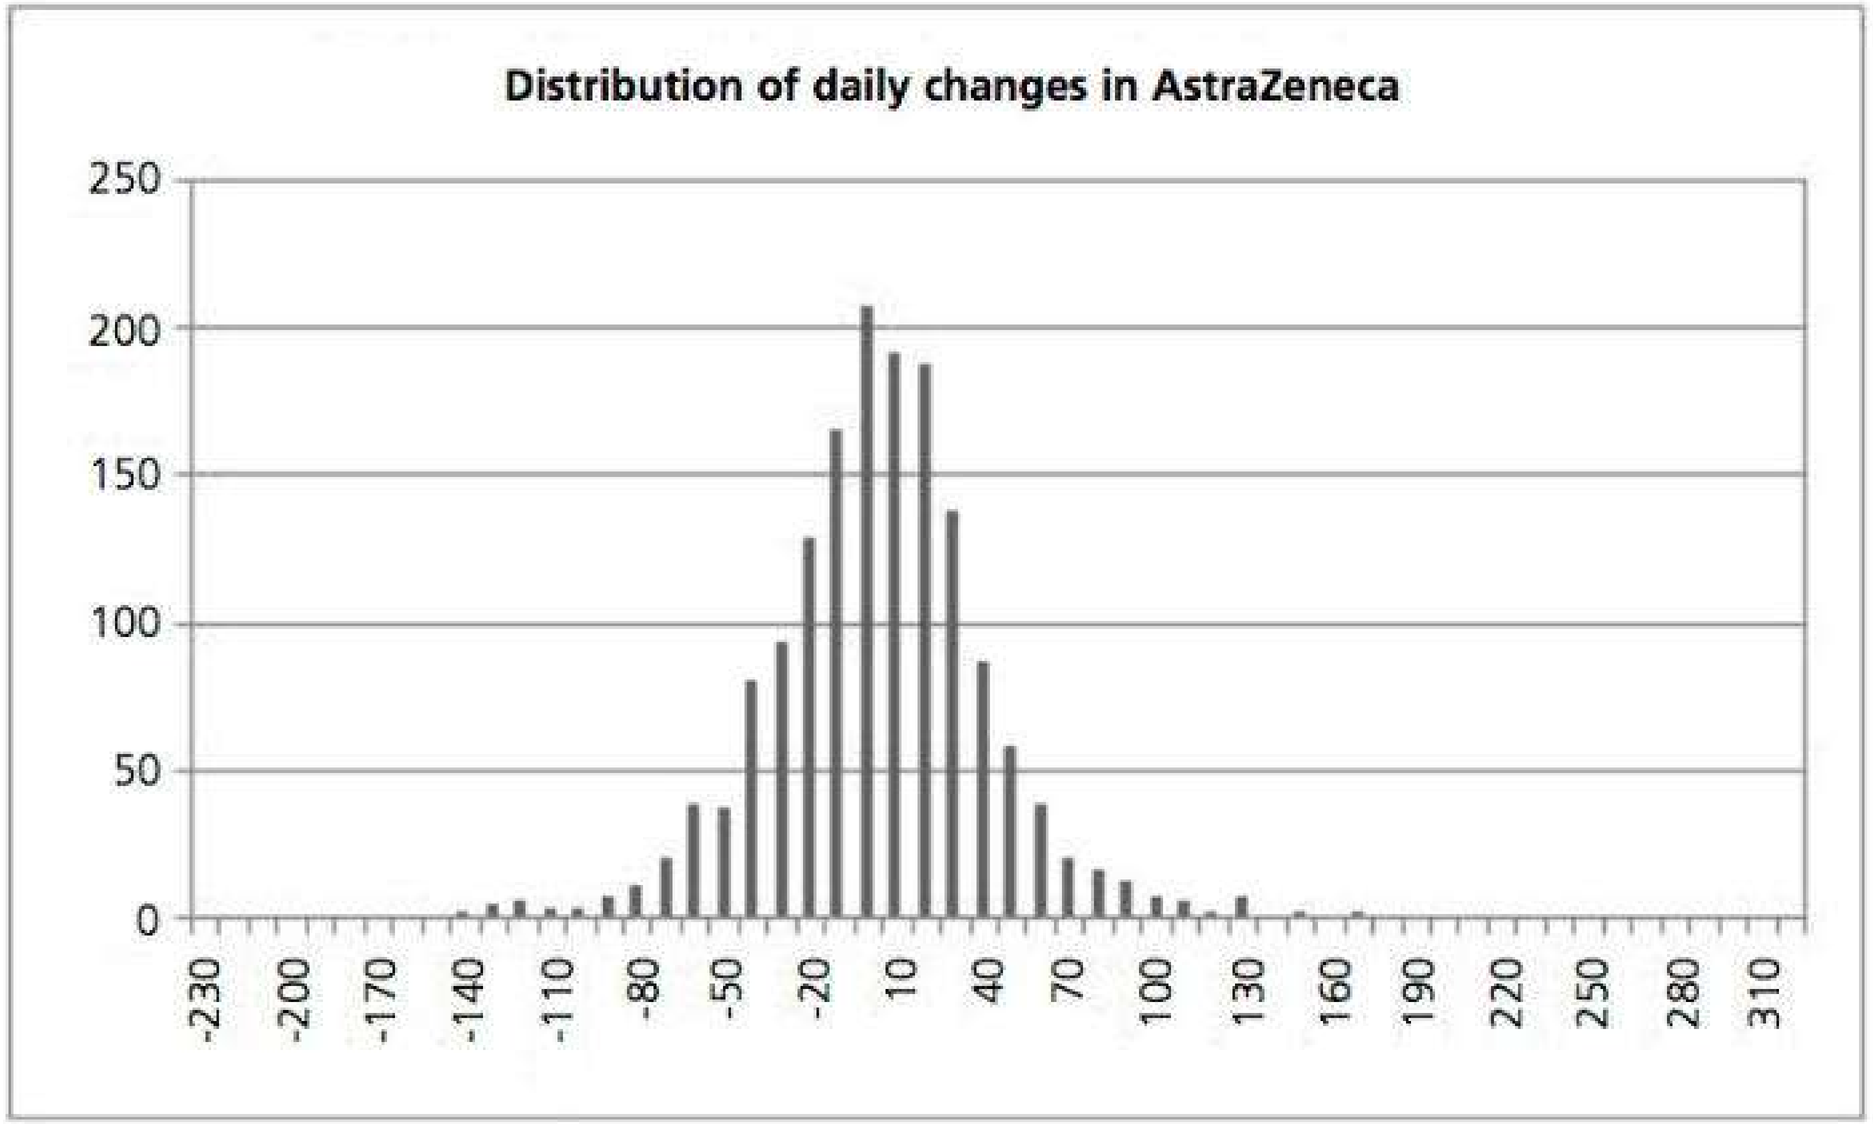


The question arises as to whether the distribution shown in Figure 4.7 is a normal distribution. The actual distribution is much higher in the middle relative to the tails than it should be for a normal distribution. Obviously it is not symmetrical and so we can say it is not exactly a normal distribution, but is it close to being a normal distribution? If we bring both of these plots from Figures 4.6 and 4.7 to the same scale and then take the difference between them we get the plot shown in Figure 4.8. This shows quite clearly the increased values from the middle of the distribution.

**Figure 4.8 – The differences between the actual distribution of AstraZeneca daily price changes and what would be expected if these were normally distributed.**

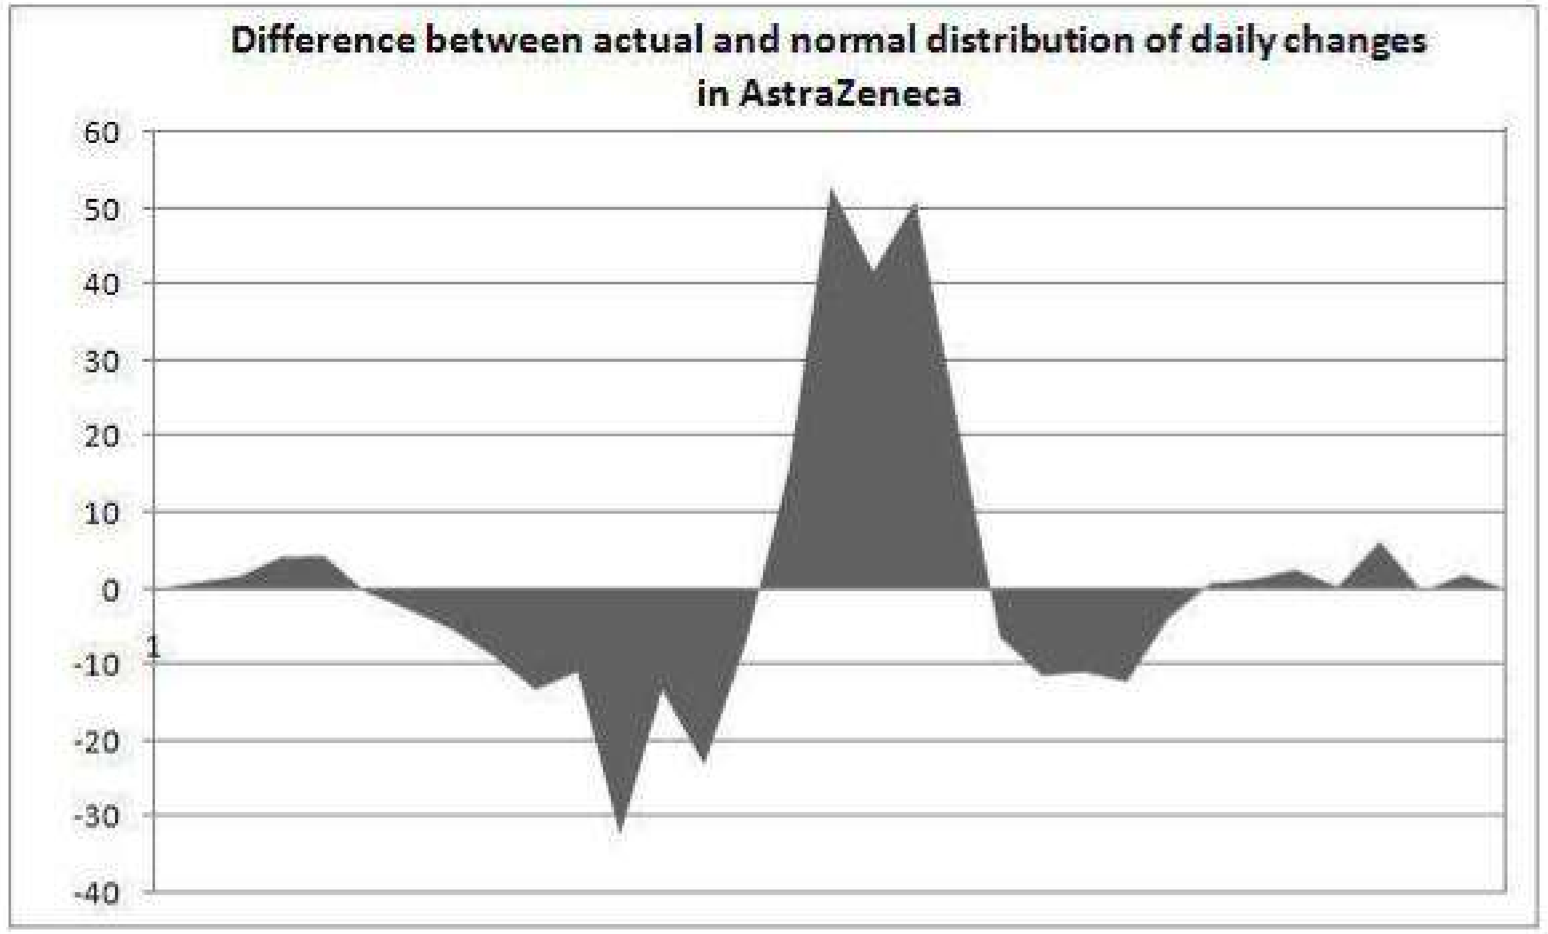


What this means is that the price changes in AstraZeneca are not randomly distributed. AstraZeneca was used as an example because it is a security quoted both in London and in New York. Other securities also behave in exactly this way. If prices are randomly distributed then the whole concept of technical analysis would be meaningless, since trend prediction would not be possible.

Now that we have decided that the actual price changes are not normally distributed, the question arises of whether or not we can still determine probabilities from this distribution. The answer fortunately is yes, since we can adopt the same approach that we used for the normal distribution of adding the values in each bin to give a cumulative plot. This gives the distribution shown in Figure 4.9.

**Figure 4.9 – The actual distribution and cumulative distribution from AstraZeneca daily changes. Moving the cursor over the cumulative plot gives the probability that the price**

**change will be less than the indicated value. In this case there is a 16.44% probability that the price change will be less than -29.66. This means there is a 70.34% (100-29.66) probability that the price change will be greater than -29.66.**

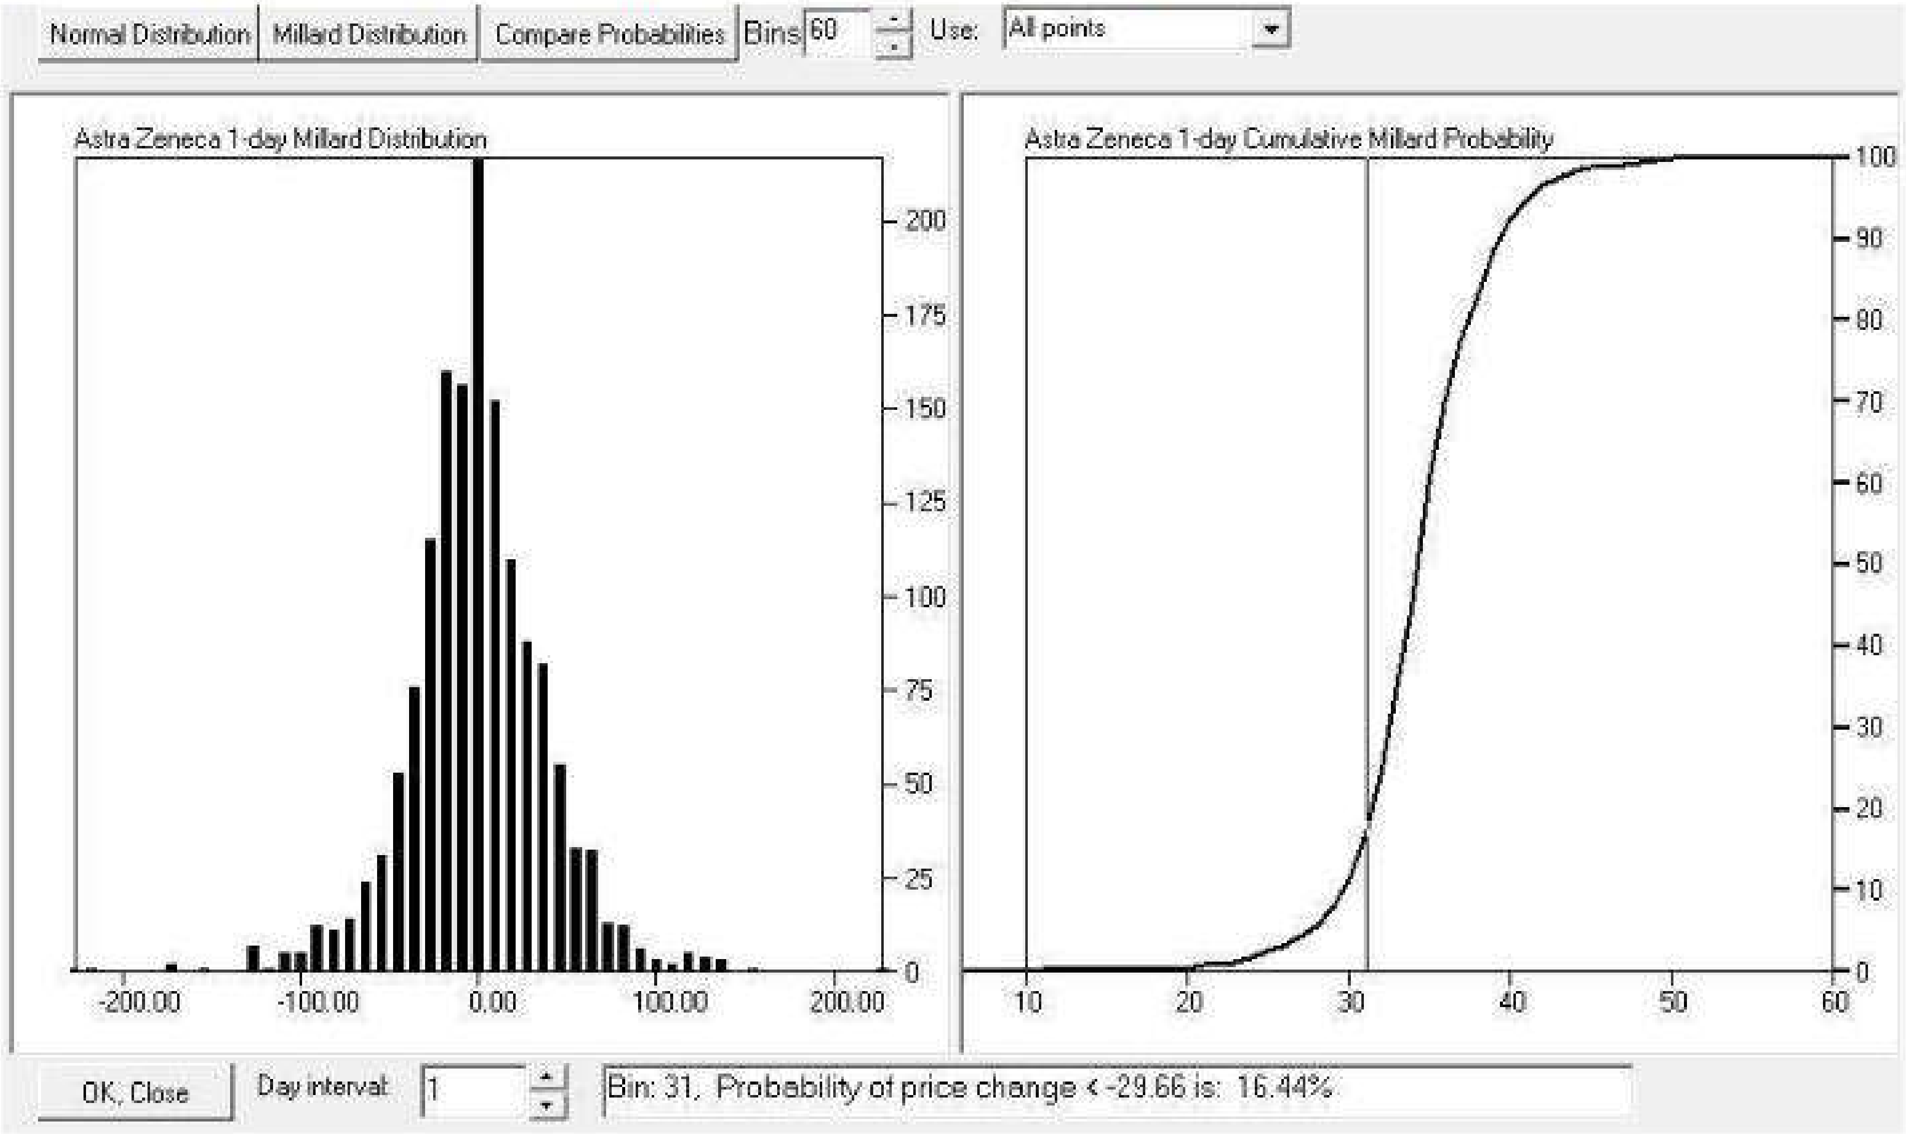


The curve is not as smooth as that shown for a normal distribution of the price changes, but that is of course because the plot of the actual distribution (Figure 4.8) does not have a smooth profile. The cumulative distribution simply reflects this fact.

***Changes of more than one day***

We can follow the principles used for calculating the distribution of daily changes to examine the distribution of changes which occur over periods longer than this. The data are obtained by the difference between each data point and the point n days further forward. As with daily changes, from this set of data a mean value and a standard deviation can be obtained. From this the normal distribution can be constructed via a spreadsheet. The normal distribution will have exactly the same shape as the normal distributions shown earlier in this chapter, as can be seen from Figure 4.10.

**Figure 4.10 – The normal distribution of five-day changes in AstraZeneca (left panel). The right panel shows the cumulative distribution from which probabilities can be obtained.**

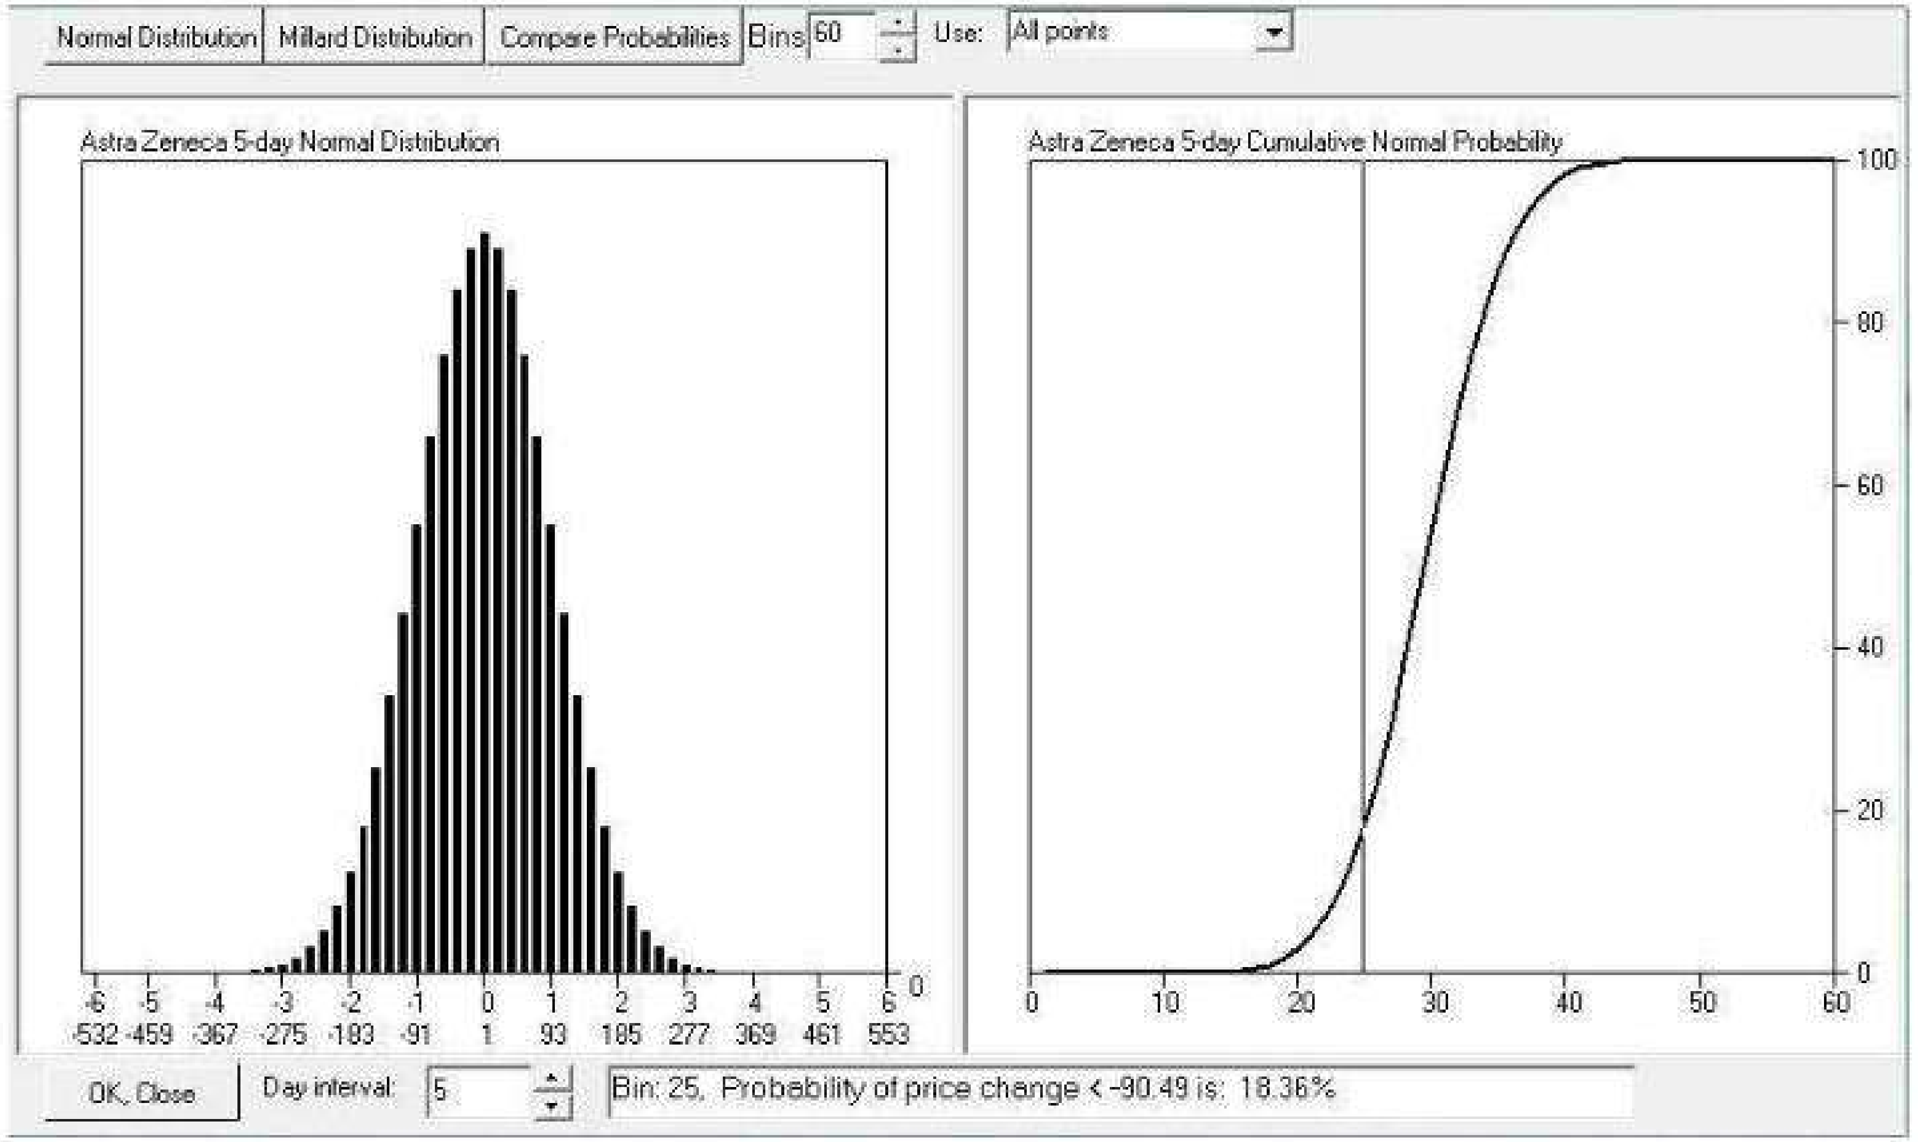


The cursor in this example has been placed at a point which shows the probability of a change of less than -90.49p over a future five-day period is 18.36%.

The actual distribution of five-day changes is shown in Figure 4.11. The interesting point here is that the actual distribution is much closer to the normal distribution shown in Figure 4.10 than is the case with one-day changes. However, the probability levels associated with the actual distribution are rather different. Thus the closest probability level that can be displayed to that of 18.36% in the normal distribution case is one of 19.83%. This is because only 60 bins have been used, giving a coarse control over the probability values. This level of 19.83% is the probability that the price change is less than -64.4p, considerably different from the value of -90.49 obtained for the normal distribution.

**Figure 4.11 – The actual distribution of five-day changes in AstraZeneca (left panel). The right panel shows the cumulative distribution from which probabilities can be obtained.**

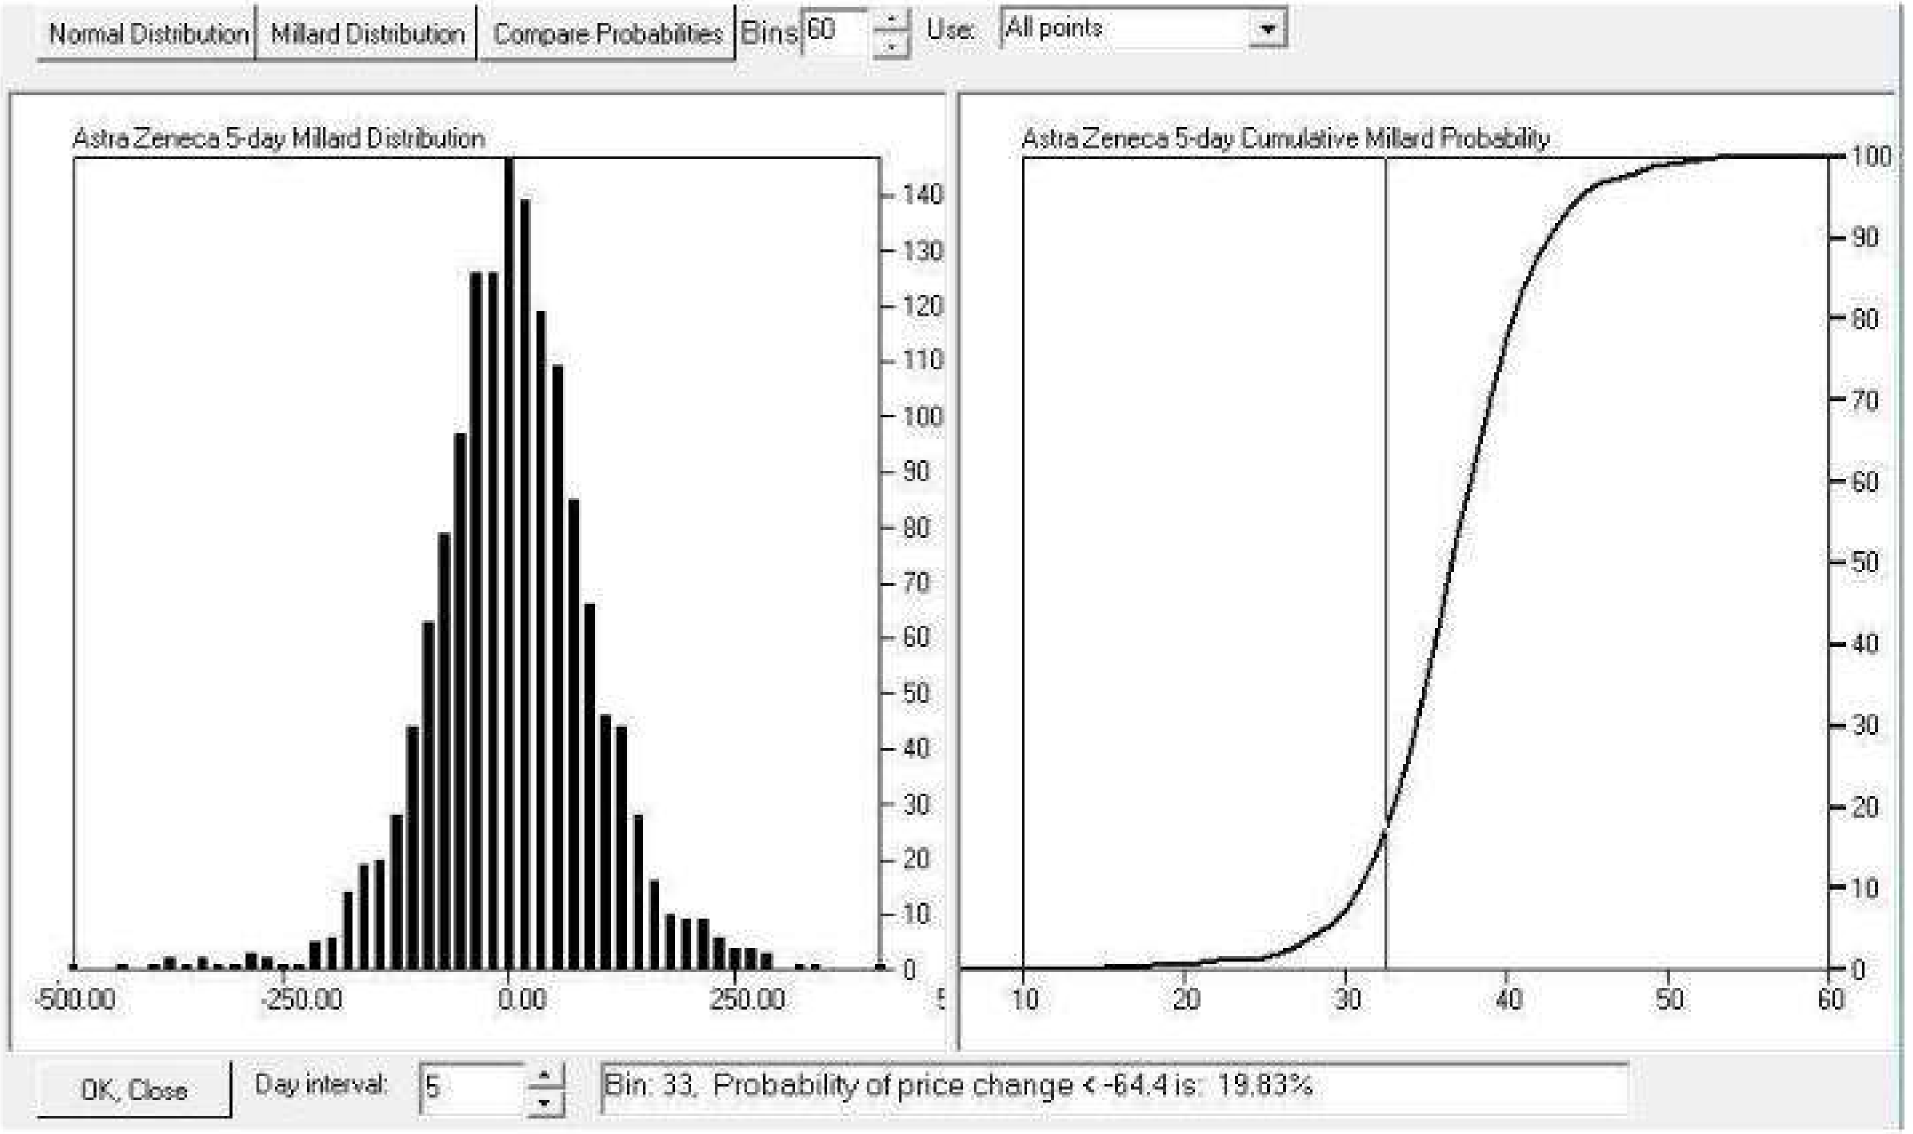


This method of estimating probable future movements is only valid for short periods of change, with five-day changes being a sensible upper limit. This can be demonstrated by taking an extreme value of change period and examining the distribution of the actual price changes. This is shown in Figure 4.12 for 151-day changes. It can be seen quite clearly that the structure has now disappeared, with a very large scatter about the mean. This of course means a large standard deviation, meaning that any prediction of a future price 151 days ahead will have such a large range as to be meaningless.

**Figure 4.12 – The actual distribution of 151-day changes in AstraZeneca (left panel). The right panel shows the cumulative distribution from which probabilities can be obtained.**

**It can be seen that now the distribution contains a wide range of changes of approximately similar probability.**

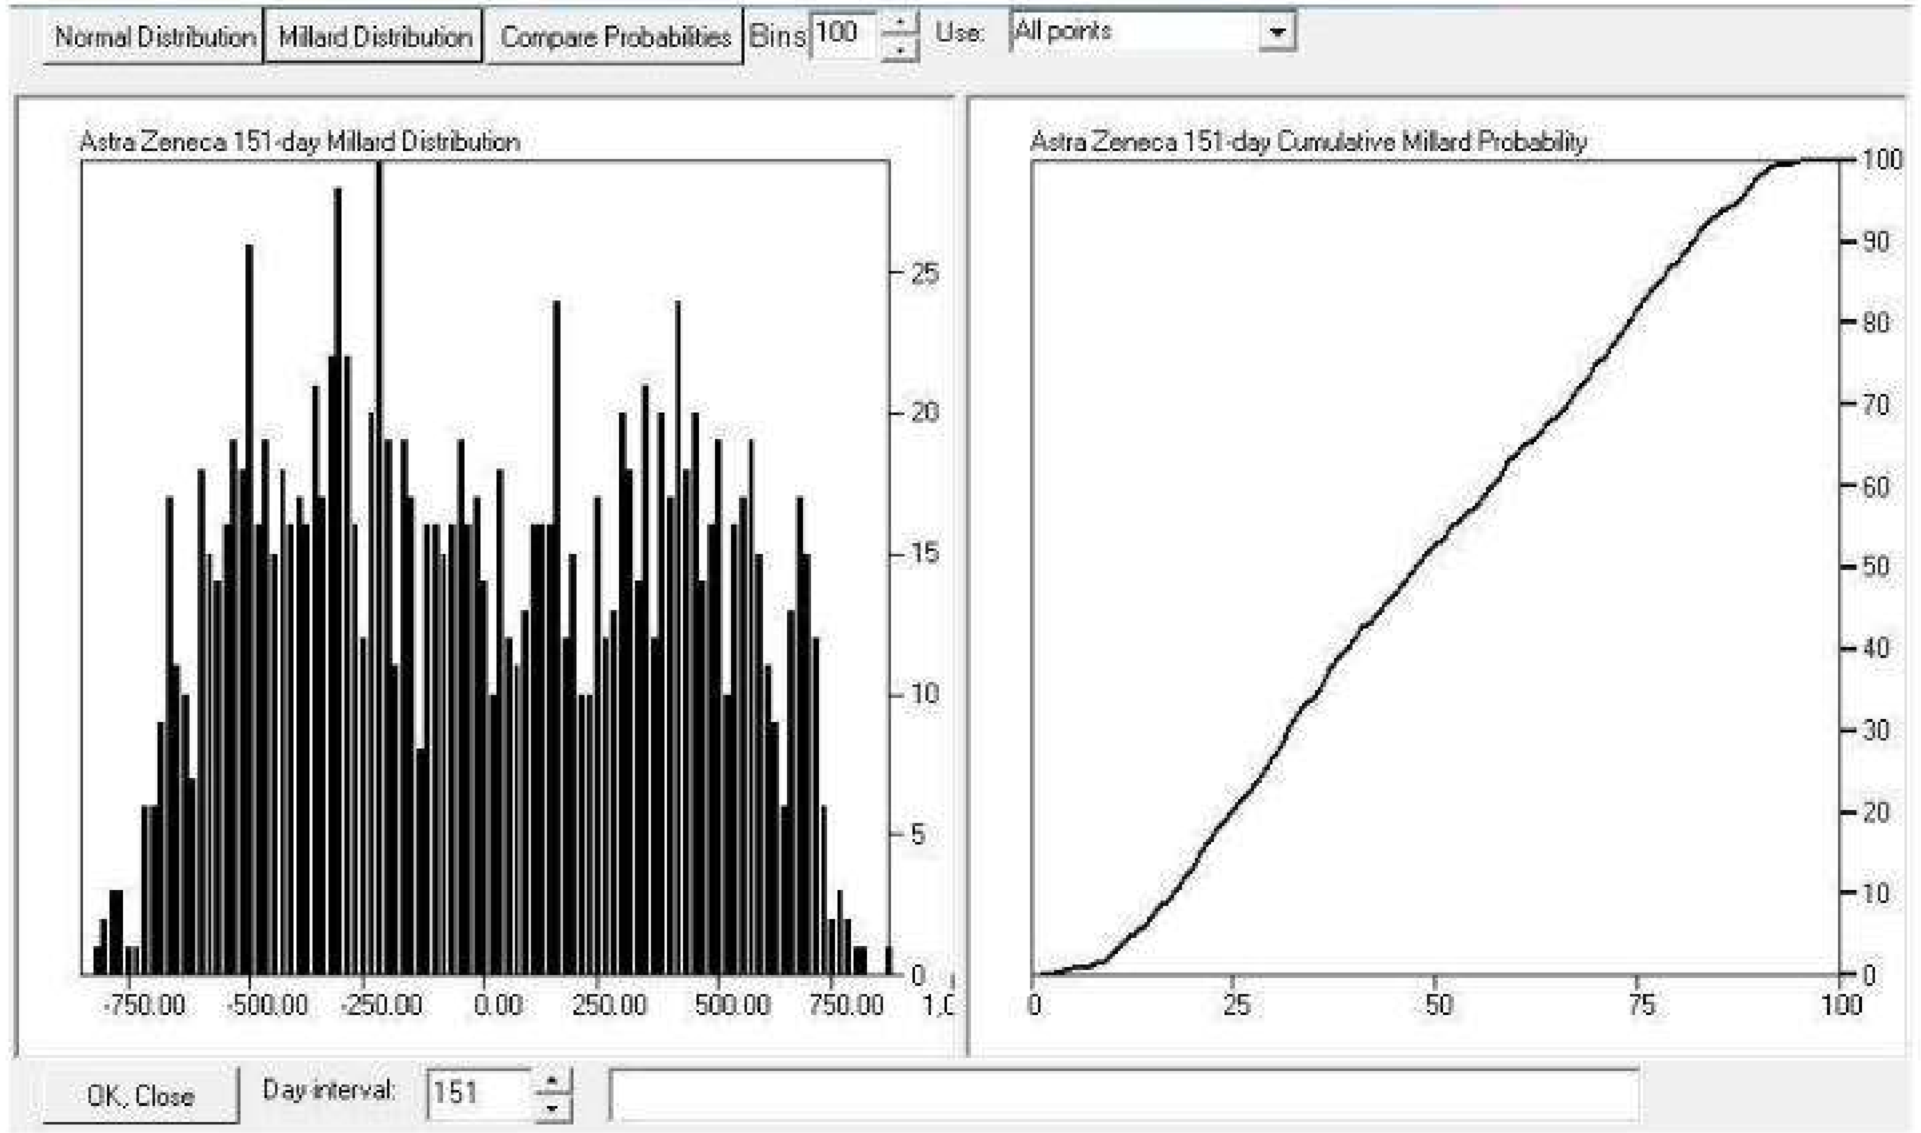


***Why large periods cannot be used***

The reason that the distribution of changes over periods of time greater than about five days loses its structure is because such changes use the same process as that used for moving averages. It will be seen in Chapter 10 that the value of a moving average depends on the value of the data point which is dropped and the value of the new data point which is brought into the addition. Thus the 151-day changes shown in Figure 4.12 are equivalent to examining the daily changes in a 150-day average. Since we know such averages to be much smoother versions of the original data, we can see why the distributions rapidly begin to differ from those when periods of up to around five days are used.

The overall message from this chapter therefore is that price changes of stocks over short periods of time are not random, since the distribution of price changes does not follow a normal distribution. Because of this, prediction of future prices over short time periods will give useful results with a limited probable price range. Once the period of time starts to exceed around five days, then the range becomes increasingly large so that predictions using this method are of limited value. However, this method is still extremely useful in determining the probability of a short-term trend changing direction, as is discussed in Chapter 10.

However, probabilities can still be derived over longer periods of change by using simulation, such as the Monte Carlo method. This is discussed in the next chapter.

## 5. Simulating Future Movement

We saw in the last chapter that using probabilities derived from the price change distribution was only valid for changes over a limited number of days. It is possible to construct probable movements over, say, a five-day period by taking probabilities from the individual daily changes as predicted from a normal distribution. Of course, we have already decided that daily changes are not normally distributed, but it is helpful to go through this exercise in order to clarify our thoughts about what is possible and what is not possible.

We pointed out earlier that one property of the normal distribution is that there is 95% probability that the next data point will lie within +- 2 standard deviations of the mean.

Thus, using this method on the AstraZeneca data, where the mean value is -0.12 and the standard deviation is 41.36, we arrive at the following range within which there is a 95% probability that the next price value will fall (the last price was 2393 on 23 February 2009):

1. lowest value = 2393 – (2 x 436) = 2310
2. highest value = 2393 + (2 x 41.36) = 2475

This therefore gives a range of four times the standard deviation for the next price, which is 7% of the last price. In other words there could be a rise of 3.5% or a fall of 3.5%.

Of course, if there is a 95% probability that the next day's change will lie between these two values, then there is a 5% probability that it will lie outside of these limits. Since there are two such extremes, then there is a 2.5% probability that the price will be below 2310 and a 2.5% probability that the price will be above 2475.

If we look two days ahead then the position would be:

1. lowest value = 2310 - (2 x 436) = 2227
2. highest value = 2474 + (2 x 41.36) = 2556

In general therefore, the 95% limits for a future price are:

1. lowest value = latest value - (2 x number of days x standard deviation)
2. highest value = latest value + (2 x number of days x standard deviation)

If the limits for AstraZeneca are plotted for a number of days into the future, then these are seen to be straight lines, as shown in Figure 5.1.

**FIGURE 5.1 – THE FUTURE PRICE FOR ASTRAZENECA WILL HAVE A 95% PROBABILITY OF STAYING BETWEEN THESE TWO STRAIGHT LINES IF A NORMAL DISTRIBUTION IS ASSUMED.**

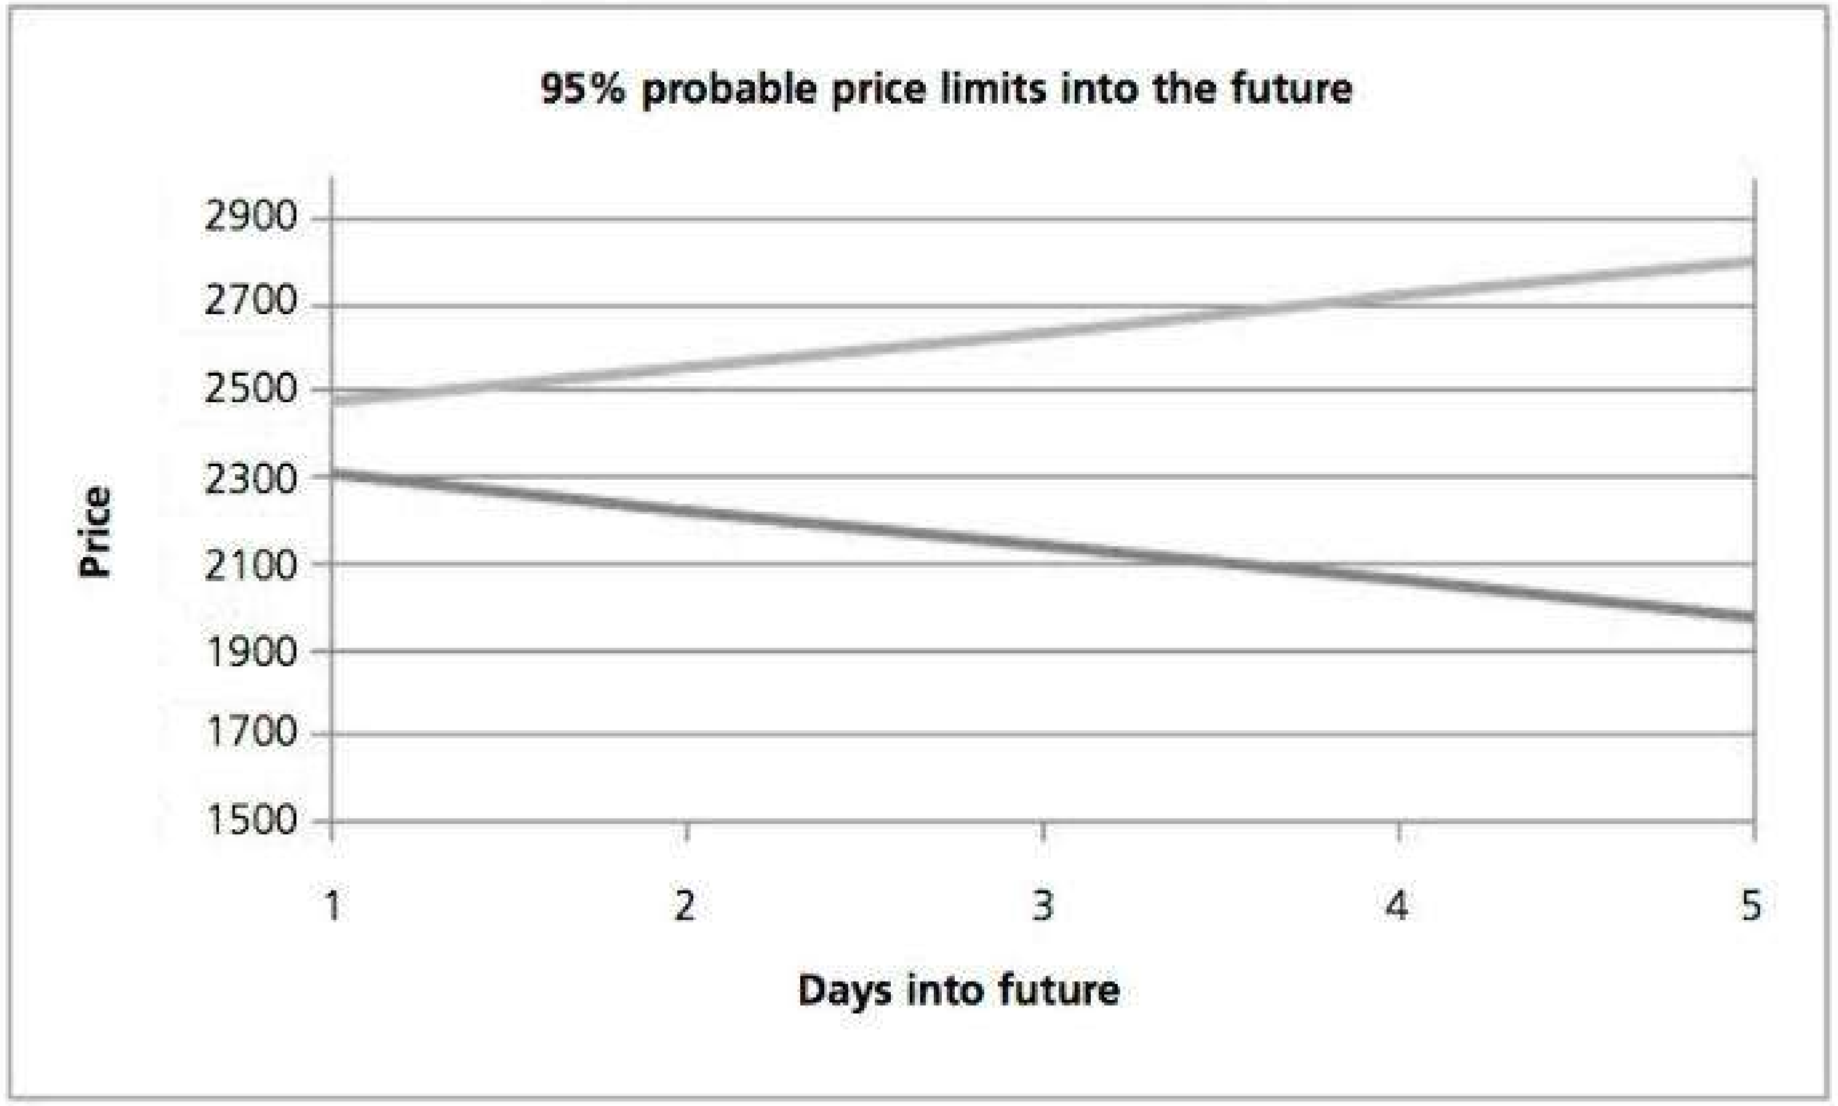


Although, as we have stated, the AstraZeneca price changes are not normally distributed, this way of looking at probabilities does give an approximate indication of the future movement. It should be pointed out that, although the range of price movement is approximated by this method, there is no indication of the direction of movement and it is this that is the subject of the remaining chapters in this book.

There is another way of tackling the issue of probability levels for future price movements over longer periods of time and that is by the simulation method.

## Monte Carlo Simulation

Monte Carlo simulation is the process of using a set of random numbers and a number or numbers randomly selected from this set to generate an outcome. This outcome will be just one of an infinite number.

A good example of the use of simulation is to determine the value of π. By performing a large number of simulations (trials) the overall outcome will tend towards the true result. Remembering Pythagoras' Theorem from our school days, this states that in a right-angled triangle the square of the hypotenuse is equal to the sums of the squares of the other two sides.

**Figure 5.2 – The radius of the circle OA can be calculated from the vertical distance AB from the origin (y-value) and the horizontal distance OB from the origin (x-value).**

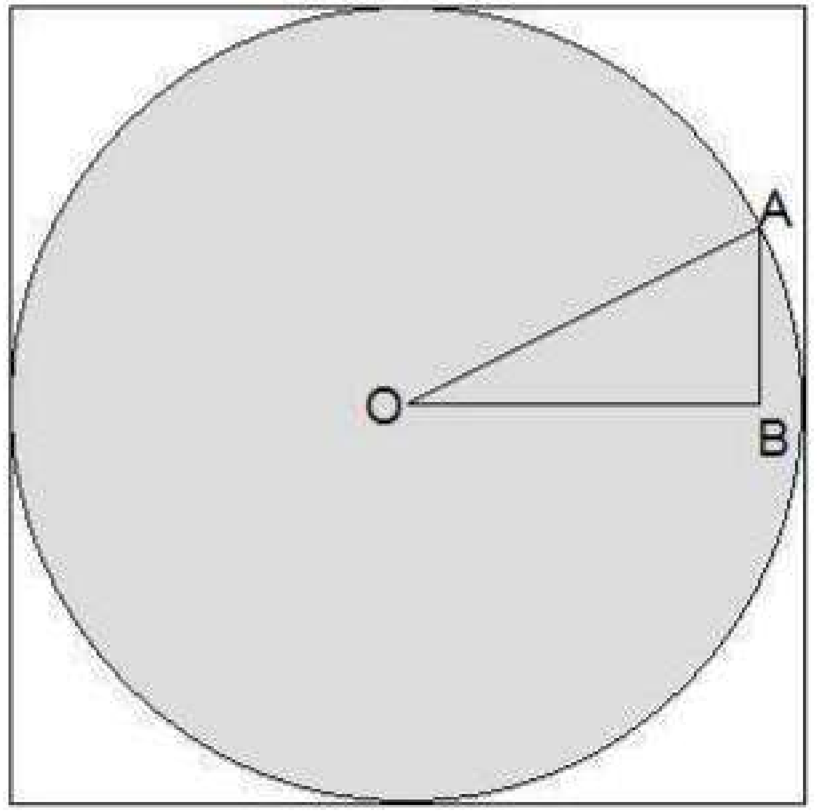


In Figure 5.2 we show a square with dimensions 2 units on each side, with a circle just touching the sides of the square. Thus the radius (r) of the circle is 1 unit. The area of this circle is πr2.

Also shown is a triangle with its longest side (the hypotenuse) equal to the radius of the circle. However, this is the triangle with the longest possible hypotenuse to keep it inside the circle. We can of course draw other triangles in which the hypotenuse is outside the circle but inside the square.

As shown, from Pythagoras we get:

(OA x OA) = (OB x OB) + (AB x AB)

If we use random values for OB and AB which lie between 0 and 1, then some of these will result in a value of OA which is equal to or less than the radius of the circle, while others will give a value of OA, which is larger than the value of the radius.

Now the area of the square is 4 units (2 units each side) and the area of the circle (with radius (r) = 1 unit) is πr2, i.e. π. Thus the ratio of the area of the circle to the area of the square is π/4.

Now the problem can be boiled down to:

1. Take two random numbers, each between -1 and + Call these randX and randY.
2. Calculate $randR^2$ = the sums of the squares of each of these, i.e. $randR^2 = randX^2 + randY^2$.
3. Calculate $rand^2$ = square root of $randR^2$.
4. Compare randR with the required radius R (which is = 1).
5. If randR is greater than 1 then the random point lies outside of the circle. Otherwise it is in the circle.
6. Calculate π, which is = 4 x number of points inside/total points.

The more such trials are carried out, the closer will the result come to the value of π, which is 3.142 (to 3 decimal places). It should be noted that this is not the method of choice for calculating π, since it has been calculated to many thousands of decimal places by powerful computer programs. However, as an introduction to the Monte Carlo method it is extremely informative. Note also that to get accurate values is very demanding of computer time, since the error in the result improves by the square root of the number of trials. Thus to reduce the error by a factor of 100 would require 10,000 times the original number of trials.

## MILLARD VERSION OF MONTE CARLO SIMULATION

The usual Monte Carlo method as applied to option pricing uses the normal distribution from which to select a random value to use in a calculation. In the Millard version of the method the numbers from which a selection is made are not a set of random numbers but the set of price changes which have actually occurred historically. In the last chapter we saw that these changes, for example in the case of AstraZeneca, were not normally distributed but had a bias towards the smaller price changes.

The method can be thought of as putting into a pot all the possible price changes which have occurred historically and then removing one at a time randomly from the pot. The first value taken is added to the last real price value to give the price after day one, the next value taken from the pot is used to give the price after day two and so on. This moves the estimated price into the future, for example, 30 days. This first exercise (one simulation) is then but one of the very large number of possibilities for price movement. In order to arrive at a more meaningful result, a large number of simulations (trials) are repeated, each of course with a different outcome for future movement. These outcomes (trials) will form a distribution, just as was the case in the last chapter where distributions of actual price movements for AstraZeneca were shown. While the actual price distribution is limited by the amount of historical data available, the Millard Monte Carlo distribution is only limited by the number of trials that are carried out. From these trials an estimate is made of the probabilities of the future price being within certain ranges, for example the 95% probability range as was calculated for AstraZeneca using the normal distribution method.

Figure 5.3 shows the result of simulating the five-day changes in AstraZeneca using 2000 trials. The 95% probable price change is from -298p to 311p. Thus taking the value of 2393p on 23 February 2009, the probable price range five trading days into the future (2 March 2009) is from 2095p to 2704p. For comparison, the actual distribution of 5-day changes in AstraZeneca is shown in Figure 5.4. In this case the 95% probable price change is from -127p to 146p. A better comparison between these two methods is obtained by calculating the 95% probable price range by each method up to nine days into the future. This is shown in Table 5.1.

**Table 5.1 – Comparison of the 95% probable price range for AstraZeneca price from one to nine days forward. The starting price on 23 February 2009 was 2393p. The column headed Millard is from the actual distribution, and the column headed Millard M.C. is from a simulation with 2000 trials.**

**Figure 5.3 – The distribution of five-day changes in AstraZeneca from a Monte Carlo simulation using 2000 trials.**

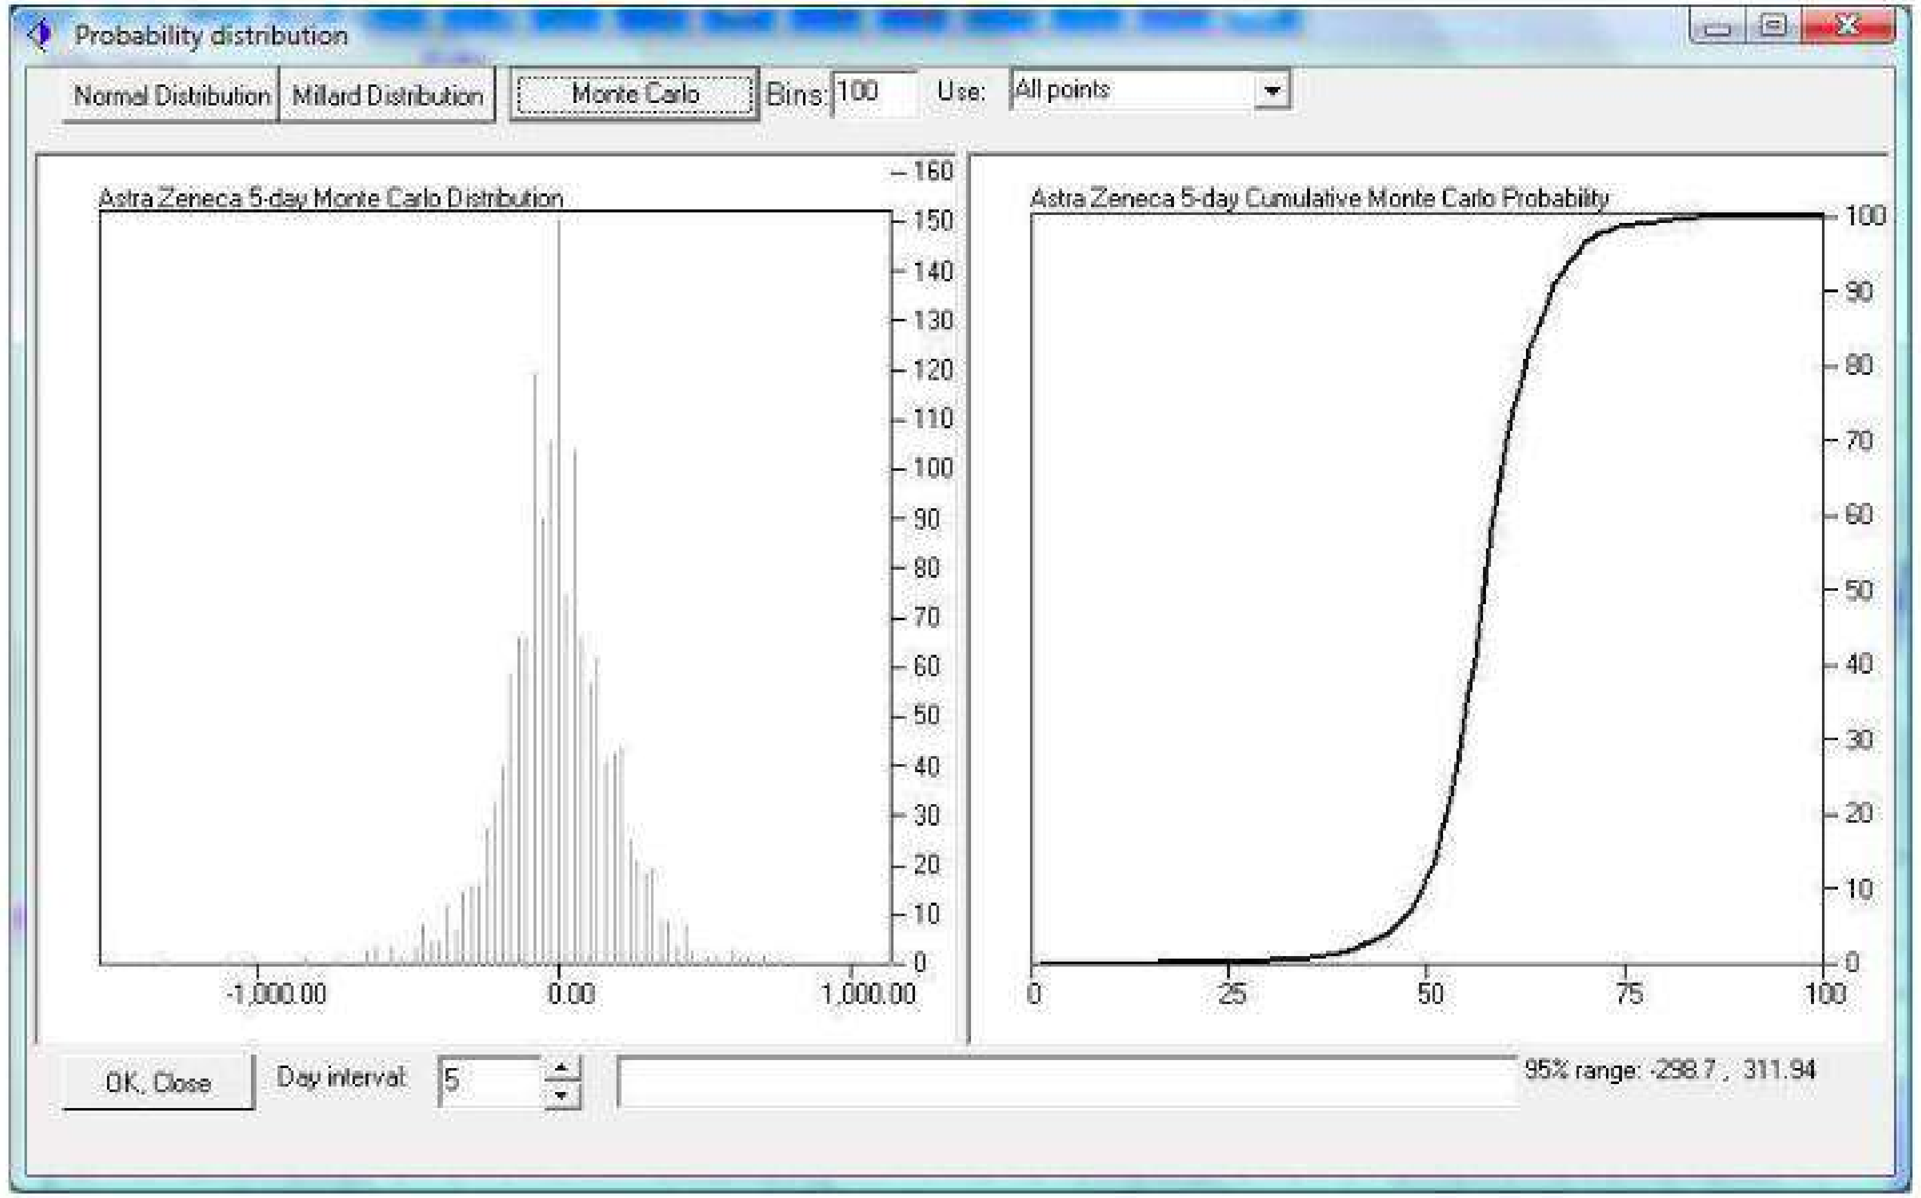


Figure 5.4 – The actual distribution of five-day changes in AstraZeneca.

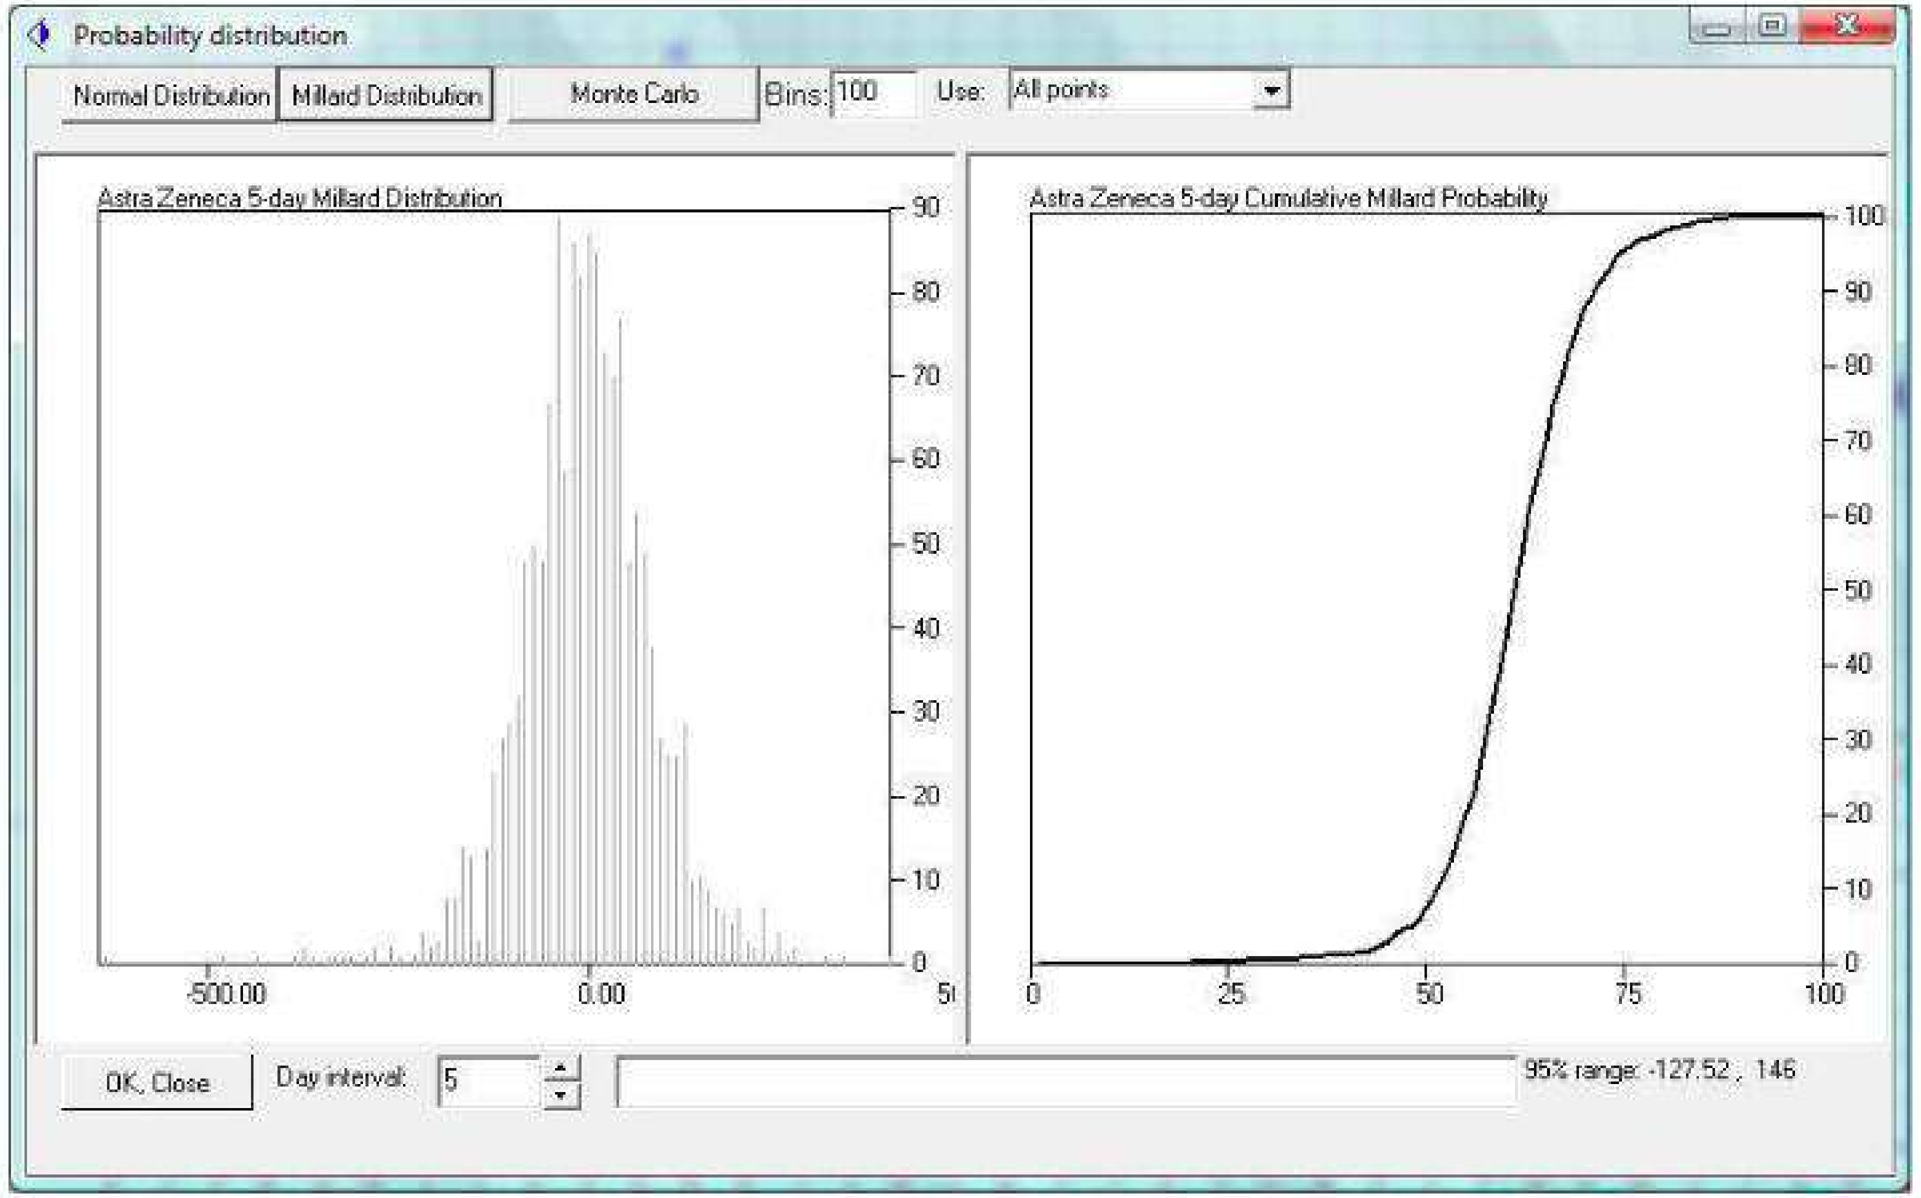


Table 5.2 – The movement of AstraZeneca nine days forward from 23 February 2009.

Figure 5.5 – The data from Table 5.1 are now plotted. The solid line is the actual movement of the AstraZeneca price since 23 February 2009. The dotted lines are the upper and lower limits of the 95% probable price range from the Millard distribution. The dashed lines are those derived from a Millard version of the Monte Carlo simulation. Day one is the starting day on 23 February.

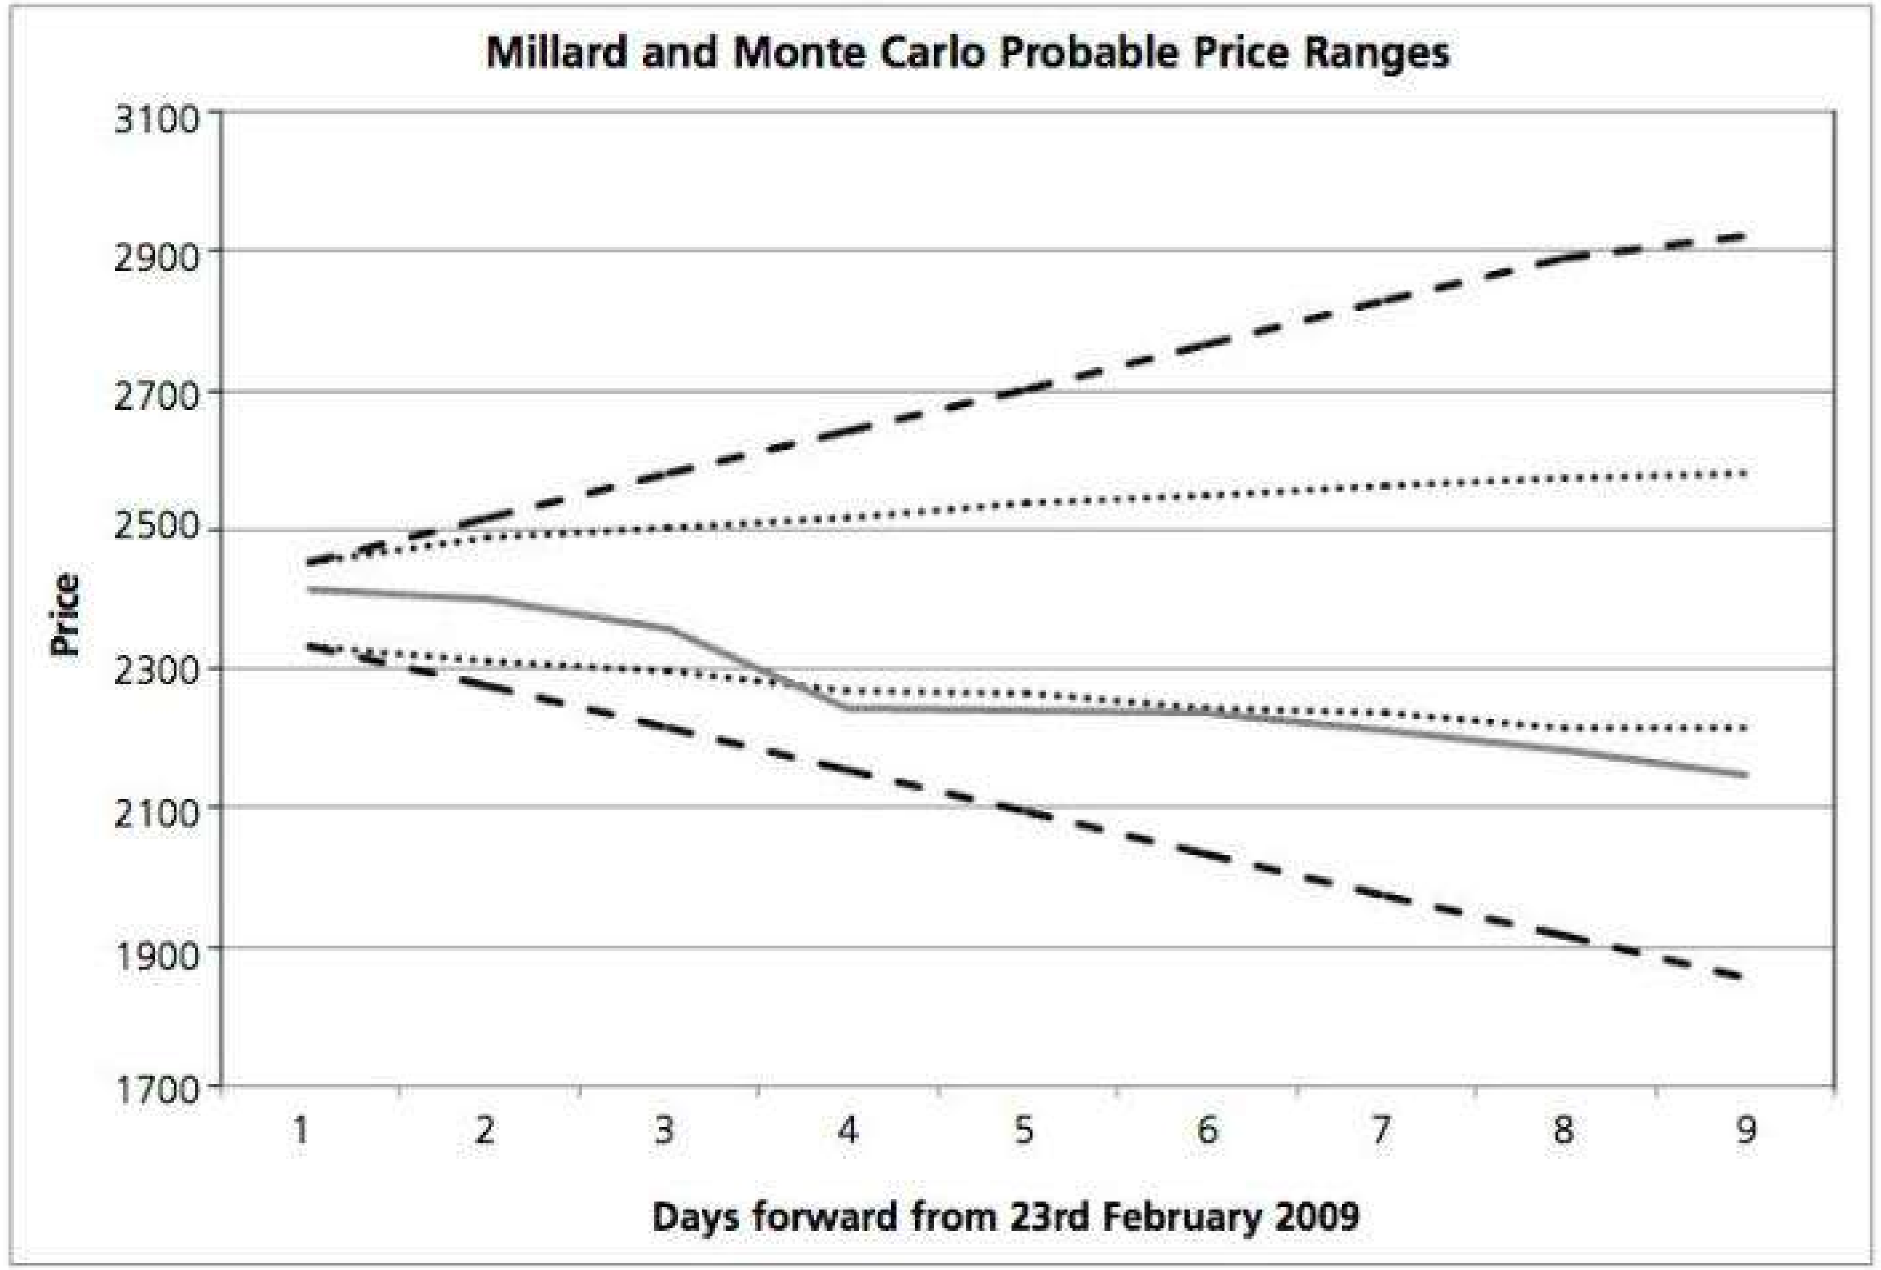


An examination of Figure 5.6 is very informative. We can see that using the 95% probabilities derived from the Millard distribution as described in the last chapter gives limits which are valid up to and including day four. However, after that the actual price falls below the lower limit and remains there to the last day (day 10 on the chart). The Monte Carlo 95% probabilities are valid throughout the chart, since the price remains comfortably within these limits, shown as a dashed line in the figure.

From these two different approaches to the calculation of the 95% probability range for future price movement, it can be seen that for the first four days, the Millard distribution method gives a tighter range for the predicted price and therefore is to be preferred. For longer periods into the future, the Millard version of the Monte Carlo simulation is to be preferred, even though it appears to give quite a wide range. We will see shortly how to address this issue of a wide range.

**Figure 5.6 – The distribution of five-day changes in American Express from a Monte Carlo simulation using 2000 trials.**

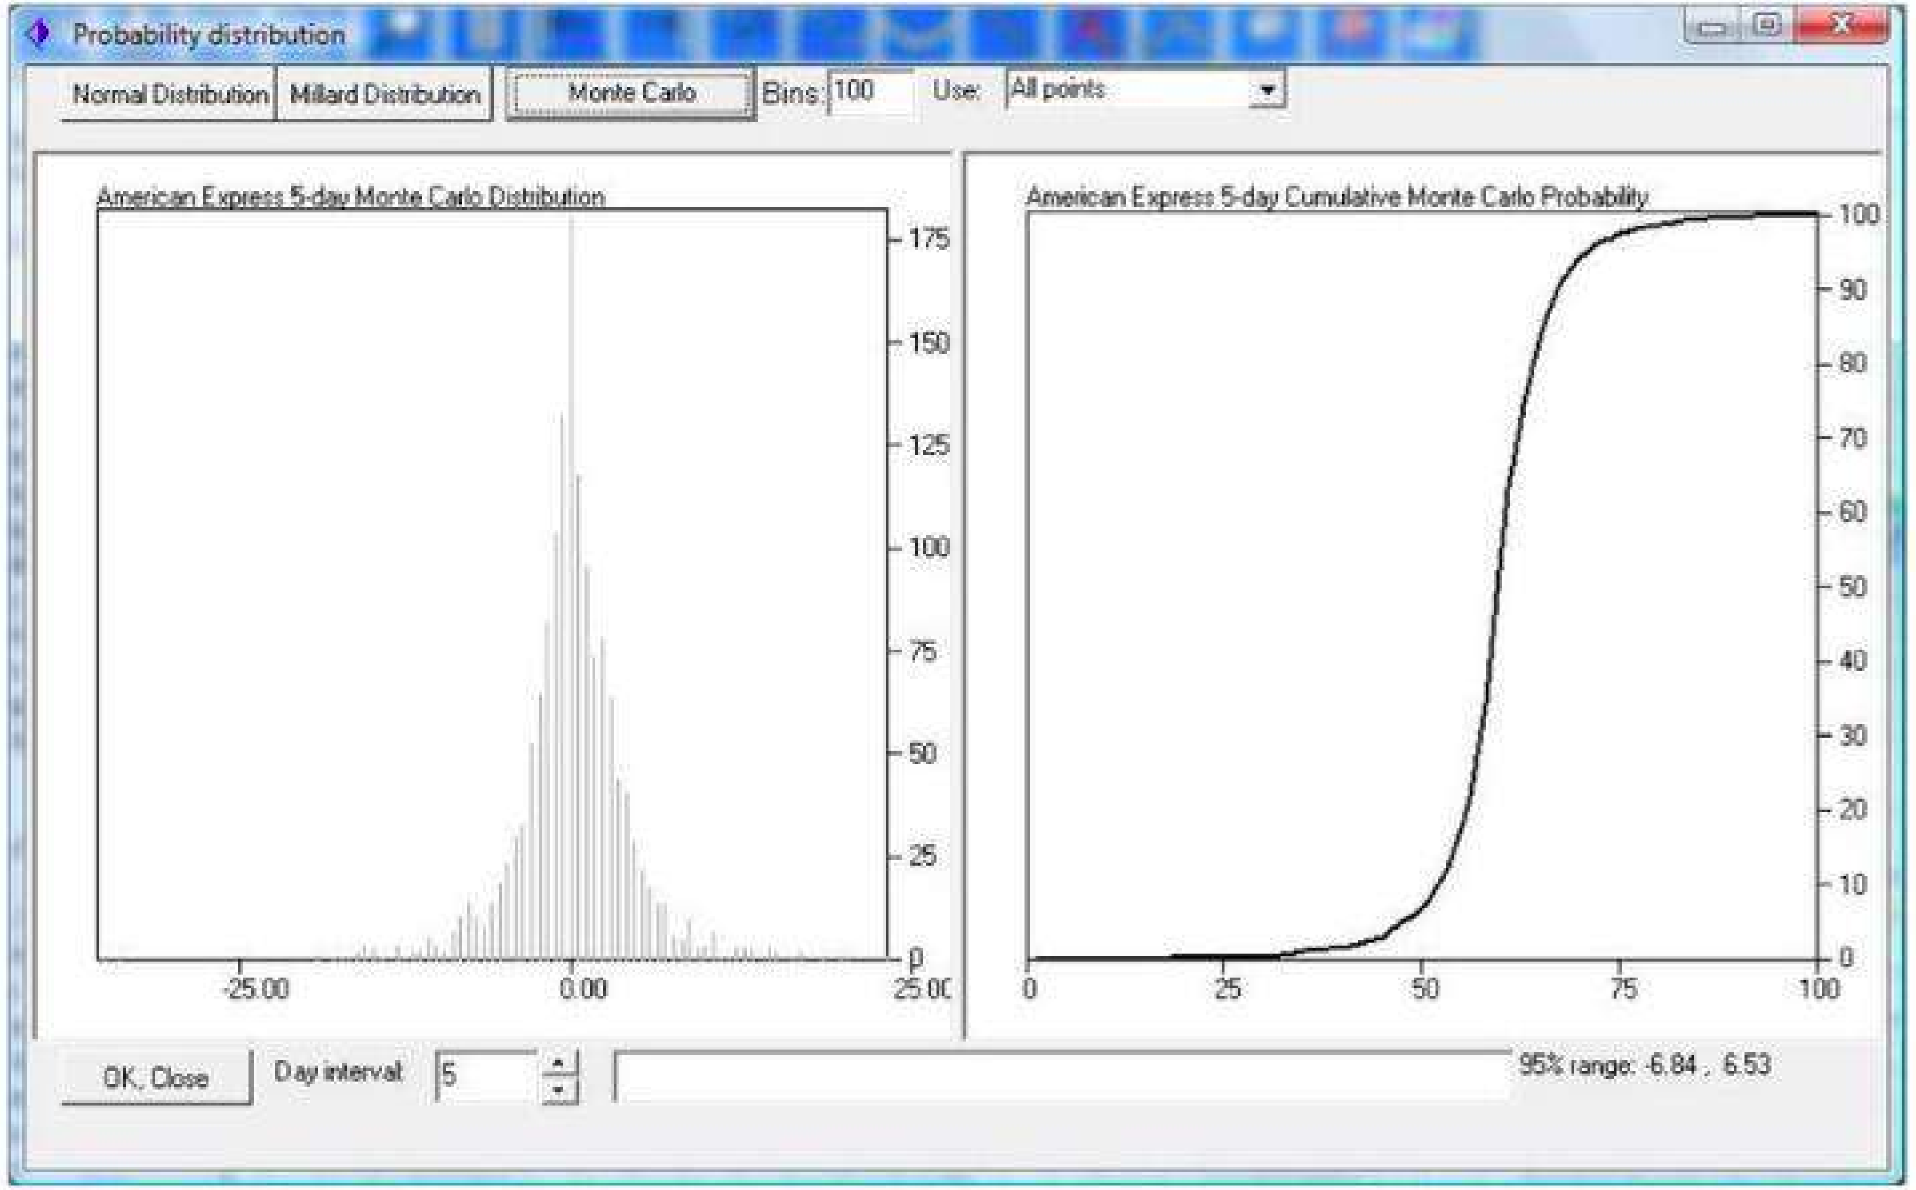


Just to confirm that these methods are applicable to other markets; in Figure 5.7 we show the Monte Carlo distribution of American Express. As was shown for AstraZeneca, the probability ranges derived from both the Millard distribution and the Monte Carlo simulations are shown in Table 5.3 for nine days forward.

**Table 5.3 – Comparison of the 95% probable range for American Express price from one to nine days forward. The starting price on 6 March 2009 was $10.26. The column headed Millard is from the actual distribution and the column headed Millard M.C. is from a simulation with 2000 trials.**

The data from Table 5.3 are plotted in Figure 5.7. Since the calculation of probable price ranges for the Millard version of the Monte Carlo method gives negative results for eight and nine days forward, these are replaced with a zero value. Unless American Express goes bankrupt, zero values have no meaning, but are used here to give a lower limit which can be plotted.

Again, we see that the actual price movement violates the probabilities derived from the Millard distribution at day five. The probabilities derived from the Monte Carlo simulation remain valid across the whole nine days.

**Figure 5.7 – The data from Table 5.3 are now plotted. The solid line is the actual movement of the American Express price since 6 March 2009. The dotted lines are the upper and lower limits of the 95% probable price range from the Millard distribution. The dashed lines are those derived from a Millard version of the Monte Carlo simulation.**

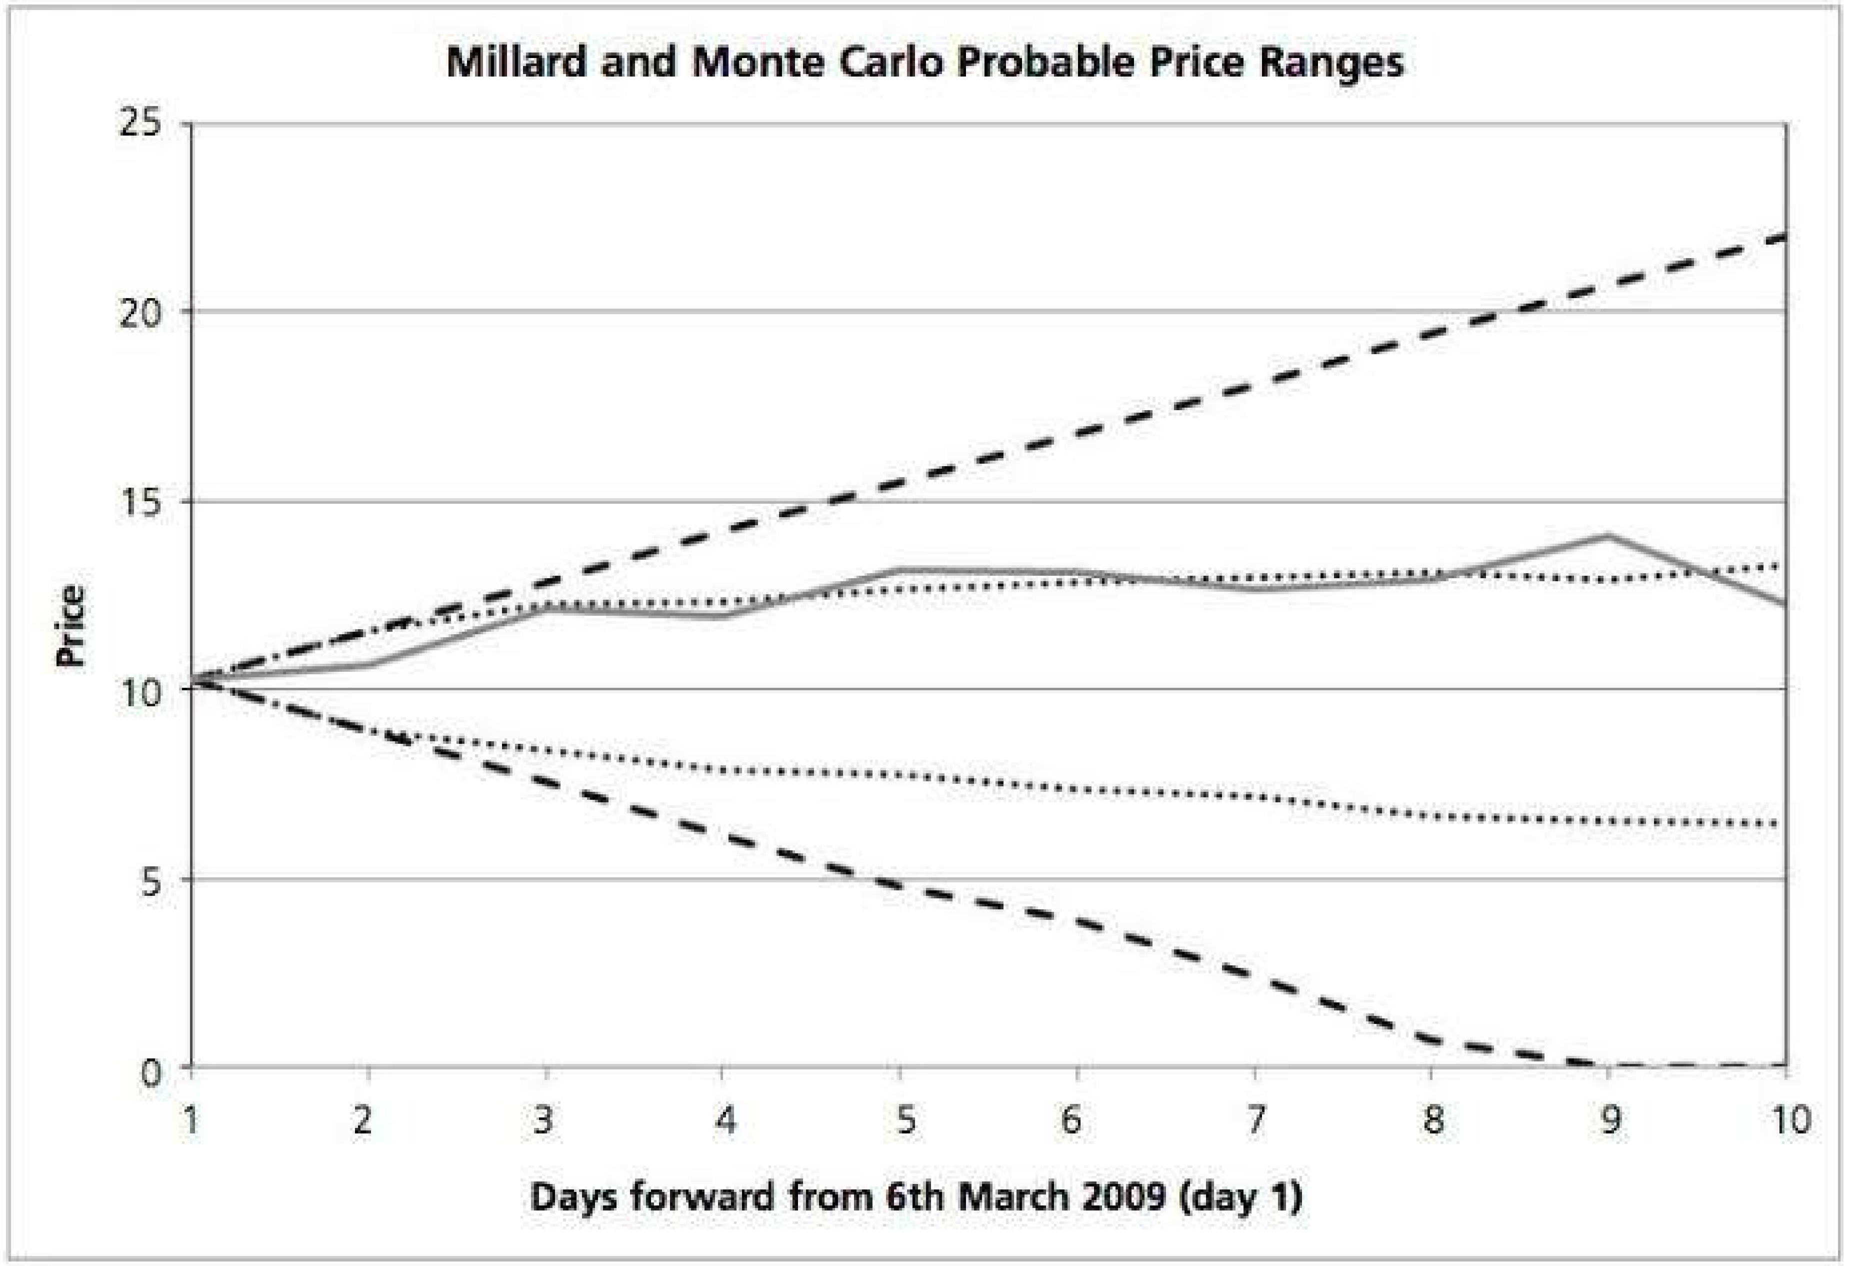


## NARROWING THE PROBABILITY BAND

Using the 95% probability bands obviously gives a wide range for the predicted future price movement that limits the usefulness of the method. A way of reducing this range is to use a lower probability, for example 90%. This had been done for AstraZeneca as shown in Figure 5.8. Quite clearly the dashed lines for the Millard version of the Monte Carlo simulation are now much closer together. Although the price level gets close to the predicted lower 90% level at day five (four days into the future), it is still contained between the predicted high and low 90% probability levels throughout the period shown in Figure 5.8.

**Figure 5.8 – The probability level for AstraZeneca has now been reduced to 90%. The solid line is the actual movement of the AstraZeneca price since 23 February 2009. The dotted lines are the upper and lower limits of the 95% probable price range from the Millard distribution. The dashed lines are those derived from a Millard version of the Monte Carlo simulation. Day one is the starting day on 23 February. The dashed lines are now closer together compared with those for 95% probabilities shown in Figure 5.5.**

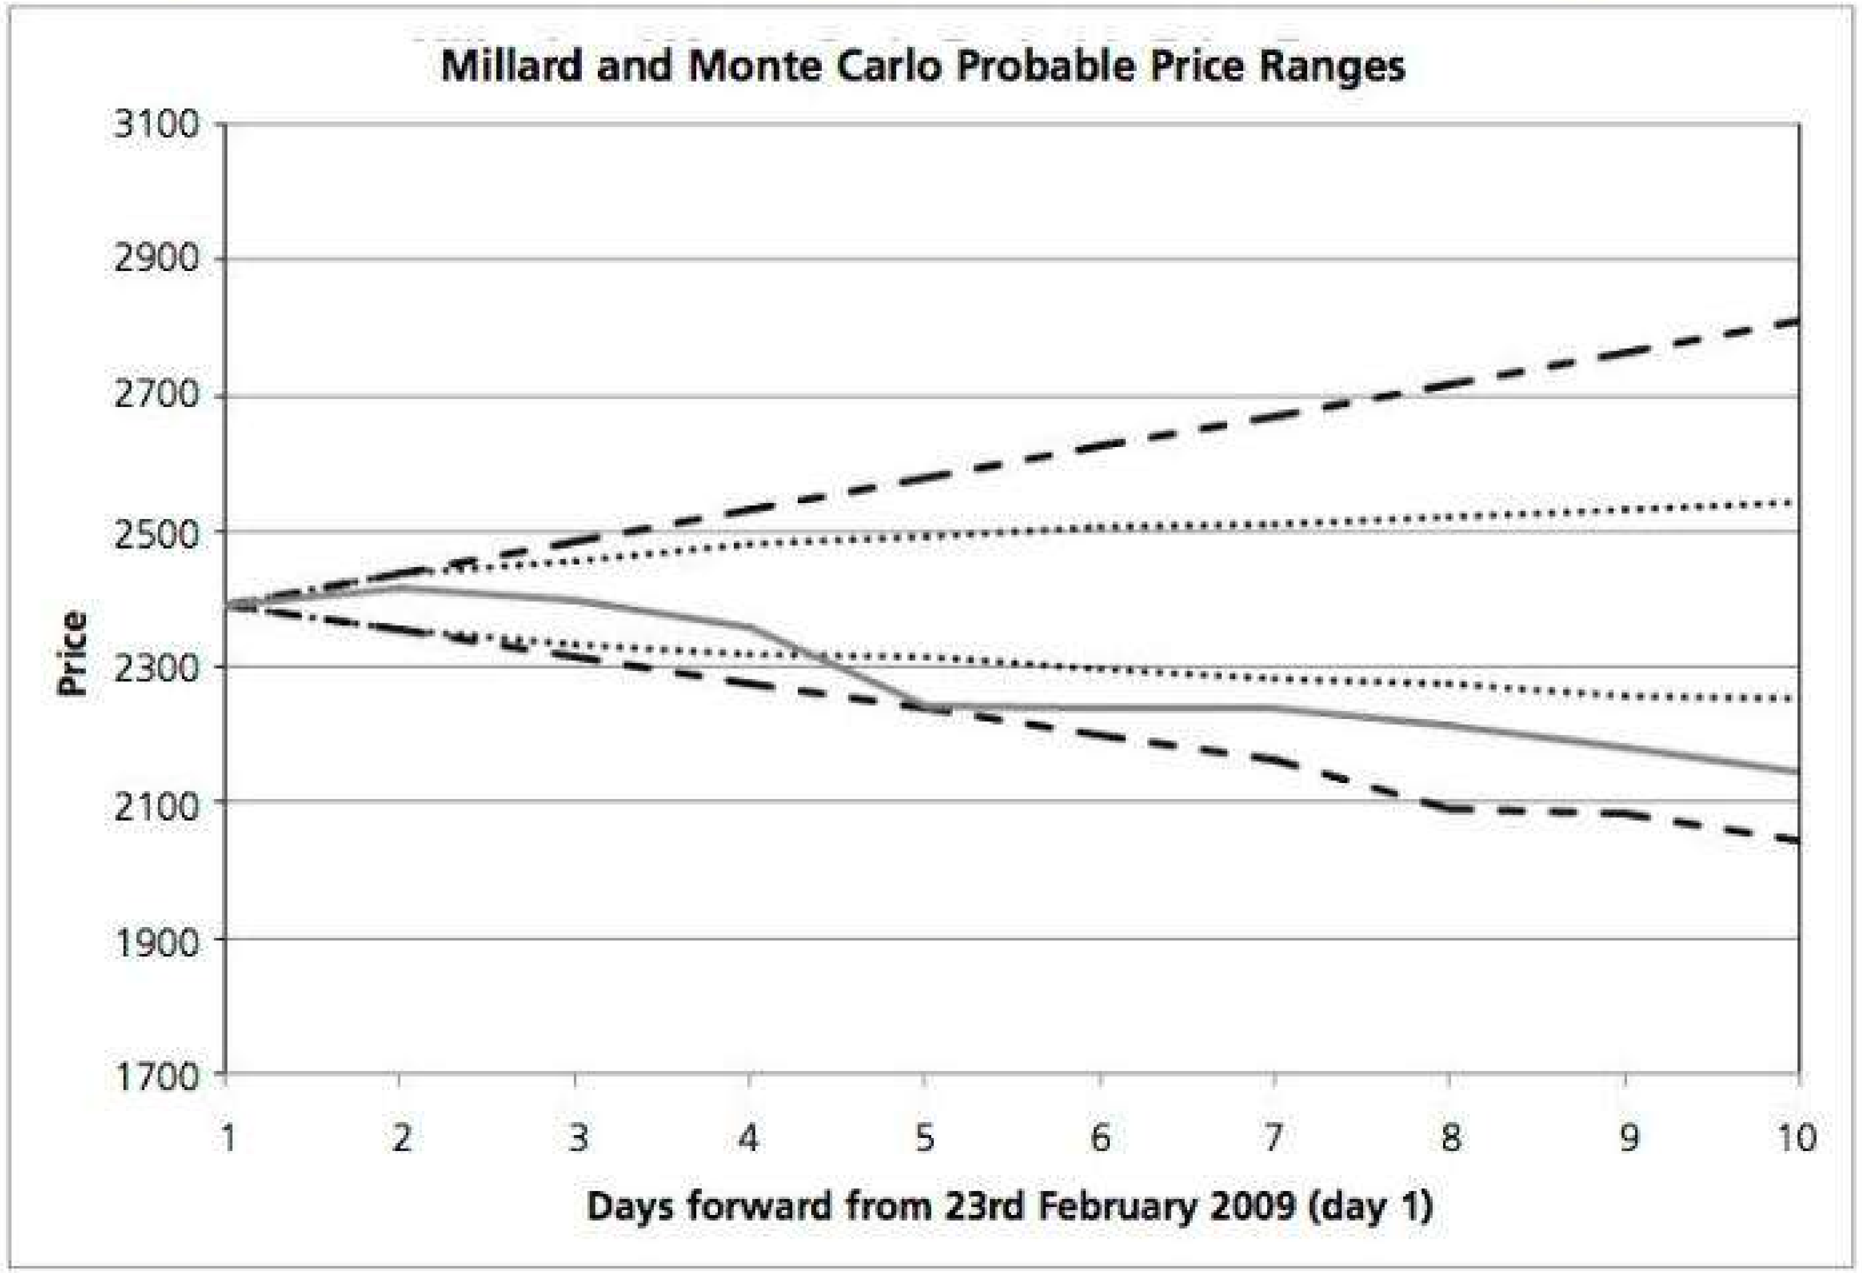


When drawing these levels on a chart in order to predict future price movement, it is best to draw a vertical line at the appropriate point in time which covers the predicted range. As an example, the predicted price range for AstraZeneca nine days into the future from the data point at 23 February 2009 is shown in Figure 5.9.

**Figure 5.9 – The predicted price range nine days into the future from the latest data point on 23 February 2009 is shown as a vertical line.**


If intuition is brought to bear on the situation shown in Figure 5.9, then it would seem that the predicted price range is ridiculously large for a point in time only nine days into the future. The price range is from 2046 to 2811, which is a range of 765p. This is 32% of the latest price value at 2393p. However, this is a case where intuition is totally wrong, as shown in Figure 5.10, where the subsequent price movement is shown. It can be seen that the price fell over the next nine days by a large amount to a value of 2147p, only 100p above the lower value of the predicted range.

**Figure 5.10 – Nine days into the future and the AstraZeneca price has fallen to a point not far above the low point of the predicted price range. The starting point on 23 February is indicated by the arrow.**


In order to reduce the predicted price range so that this low price of 2147p is close to the lower end of the predicted price range at day nine, the probability level would have to be reduced to 85%. This would give a predicted range of 2143p to 2715p.

It would be wrong to reduce the probability level to 85% and take this as a universally useful value, since this would be based on just this one example. In order to remain as a reliable method, it is best to use a probability level of 90% until future research shows that a better level is available.

What has been described in this chapter is another tool to use to examine probable price movement in the future. It is not intended to be a stand-alone method of prediction, but rather as part of the overall analysis of an individual stock which should always take place before a trading decision is taken.

## 6. Cycles and the Market

A cycle is a series of repeating events; there are many mathematical curves which exhibit this behaviour. Thus a particular section of the graph of such a curve can be seen as a pattern which is then repeated many times. For example, a cycloid is the curve formed by the vertical position of a point on a circle which is rolling on a flat surface. There are a number of interesting facts about a cycloid, such as the fact that the area under the curve is three times the area of the circle which generates it. Also the length of the cycloid is four times the diameter of the generating circle. An example of a repeating cycloid is shown in Figure 6.1.

**Figure 6.1 – A cycloid such as that shown here is formed by a point on the circumference of a rolling circle.**

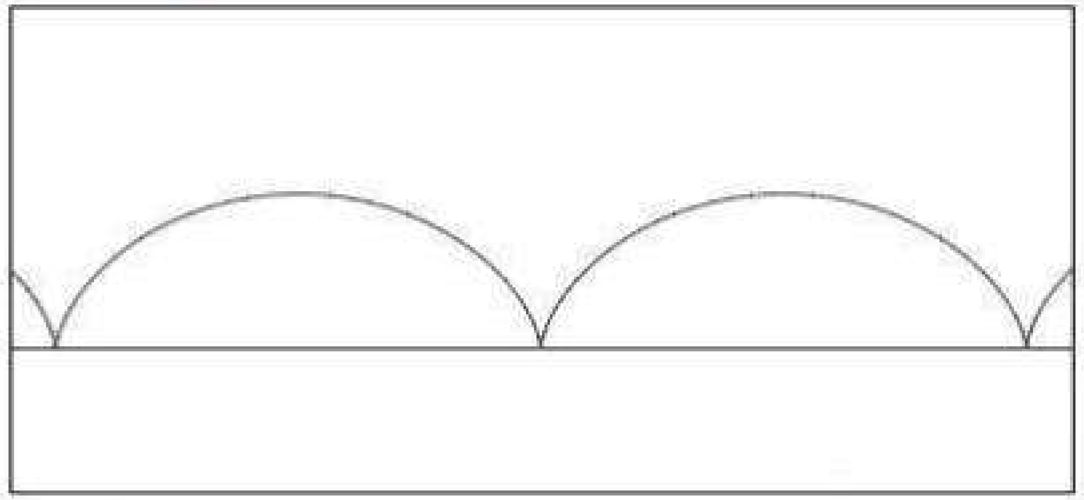

It is possible to find this type of curve in market data, but the pattern does not repeat more than once. Usually such a pattern in the market is characterised by an increasing rate of fall to a point from which it rises equally rapidly. An example is shown in Figure 6.2.

**Figure 6.2 – This short section of the closing values of the Honeywell stock price shows some characteristics of a cycloid with rounded tops and a sharp reversal from its lows.**


However, owing to its infrequent occurrence, the cycloid is of no significance in helping us to predict market movement in a more general sense. Of much more interest to us is the sine curve, an example of which is shown in Figure 6.3.

**Figure 6.3 – A sine wave.**

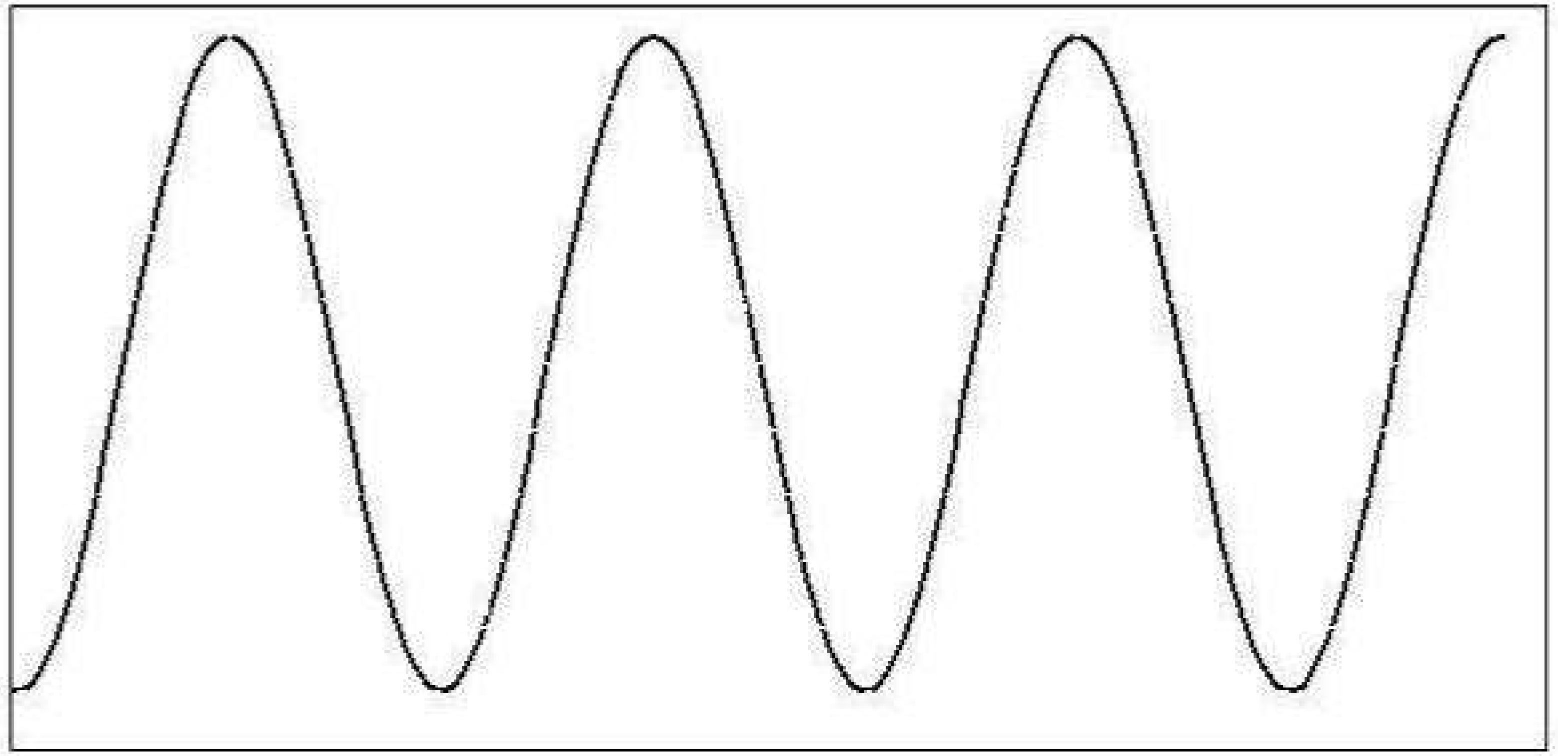

We will be able to demonstrate that the cycles which are present in market data are indeed of this type.

## Properties of Sine Waves

Sine waves such as that shown in Figure 6.3 are completely defined if we know three variables. These are:

- Wavelength or frequency
- Amplitude
- Phase

To help in future discussion we will define the section, from one point on the wave to the next corresponding point, as a sweep.

### WAVELENGTH

The *wavelength* is the distance between any point on the wave to the next corresponding point. For ease of measurement, this can be taken as the distance from one trough to the next trough or one peak to the next peak. A *sweep* therefore covers one wavelength. Wavelengths in the stock market are expressed in minutes, days, weeks or years. As far as this book is concerned, we will be using days, since we will only be concerned with daily closing values. We will not be using intra-day data, so that wavelengths in minutes will not be applicable. Two sine waves with different wavelengths but similar amplitudes are shown in Figure 6.4. One wavelength is three times the other.

***Figure 6.4 – This shows two sine waves of different wavelengths. One wavelength (thin line) is three times that of the other (thicker line). The sine wave of shorter wavelength terminates short of the right-hand edge.***

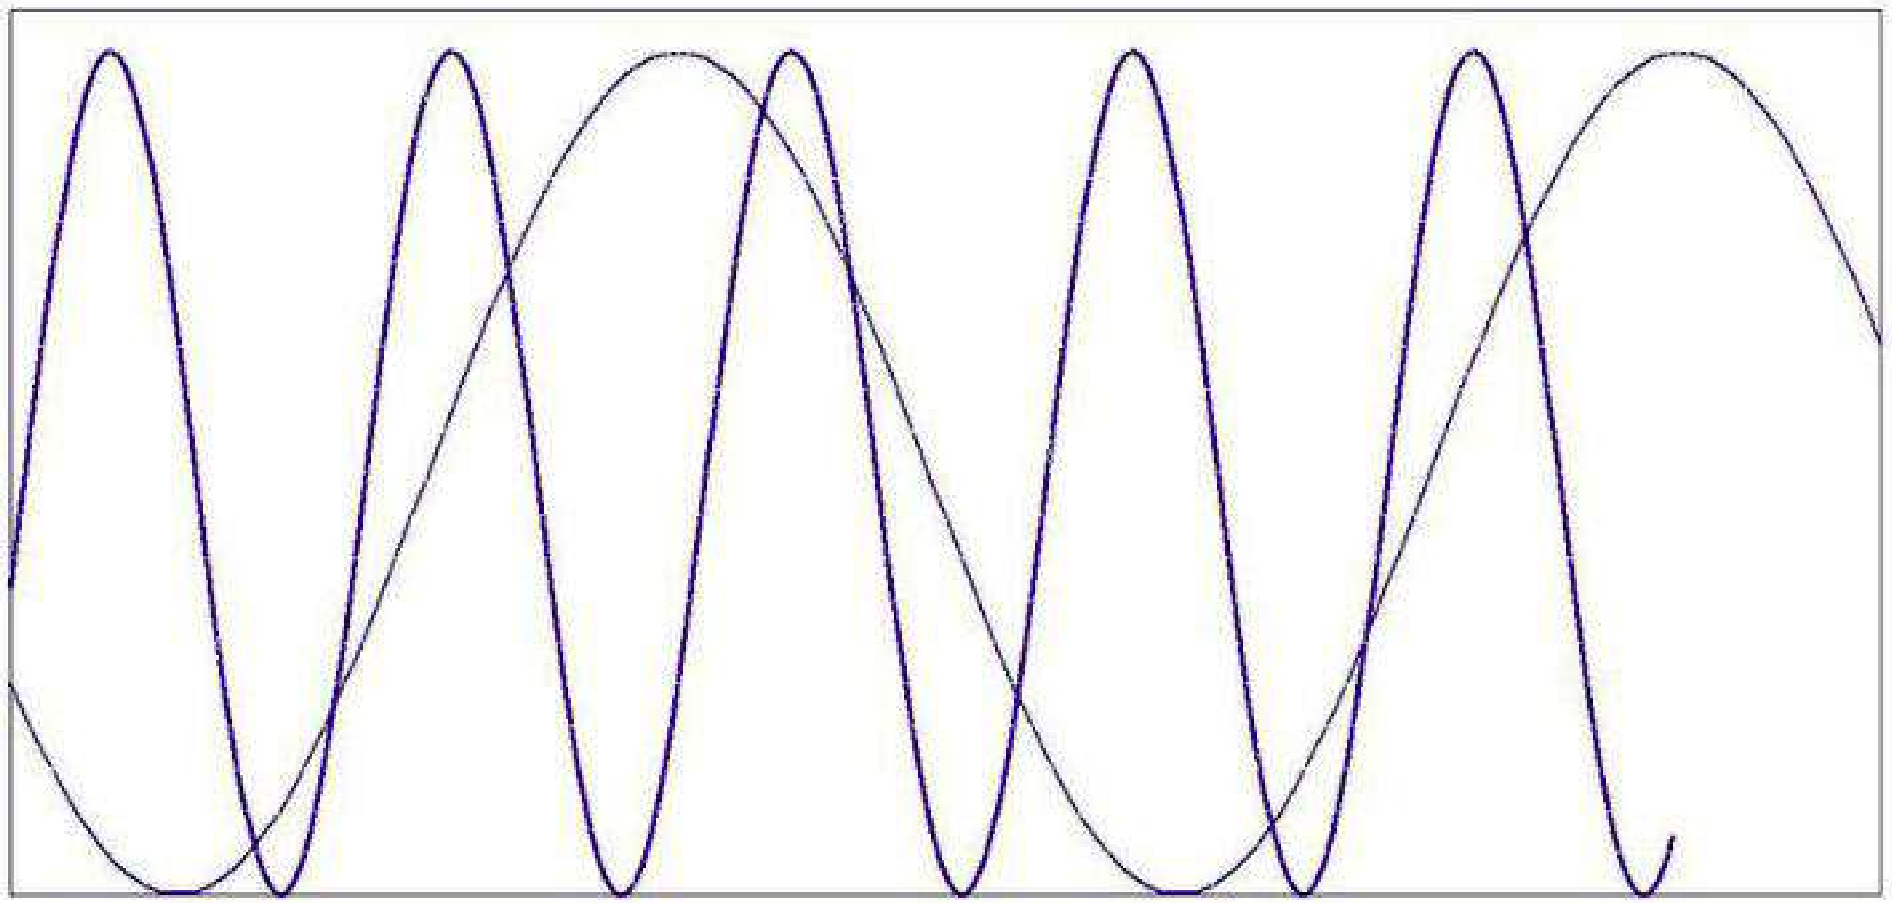

*Frequencies* are inversely related to wavelengths. Thus a wavelength of 10 days is equivalent to a frequency of 0.1 per day. A frequency of 50 days is equivalent to a frequency of 0.02 per day.

### AMPLITUDE

For our purposes we will take the *amplitude* of a sine wave to be the vertical distance from trough to peak. For stock market data this may be measured in a unit of currency, i.e. UK pounds/pence or US dollars/cents. For indices the measurement is simply a number, while for currencies it is of course a number which represents a ratio. A plot of two sine waves with the same wavelength but different amplitudes is shown in Figure 6.5. Since the wavelengths are the same and the phase is the same the peaks and troughs occur at the same positions.

***Figure 6.5 – This shows two sine waves of the same wavelength but different amplitudes.***

***The phase of each is identical.***

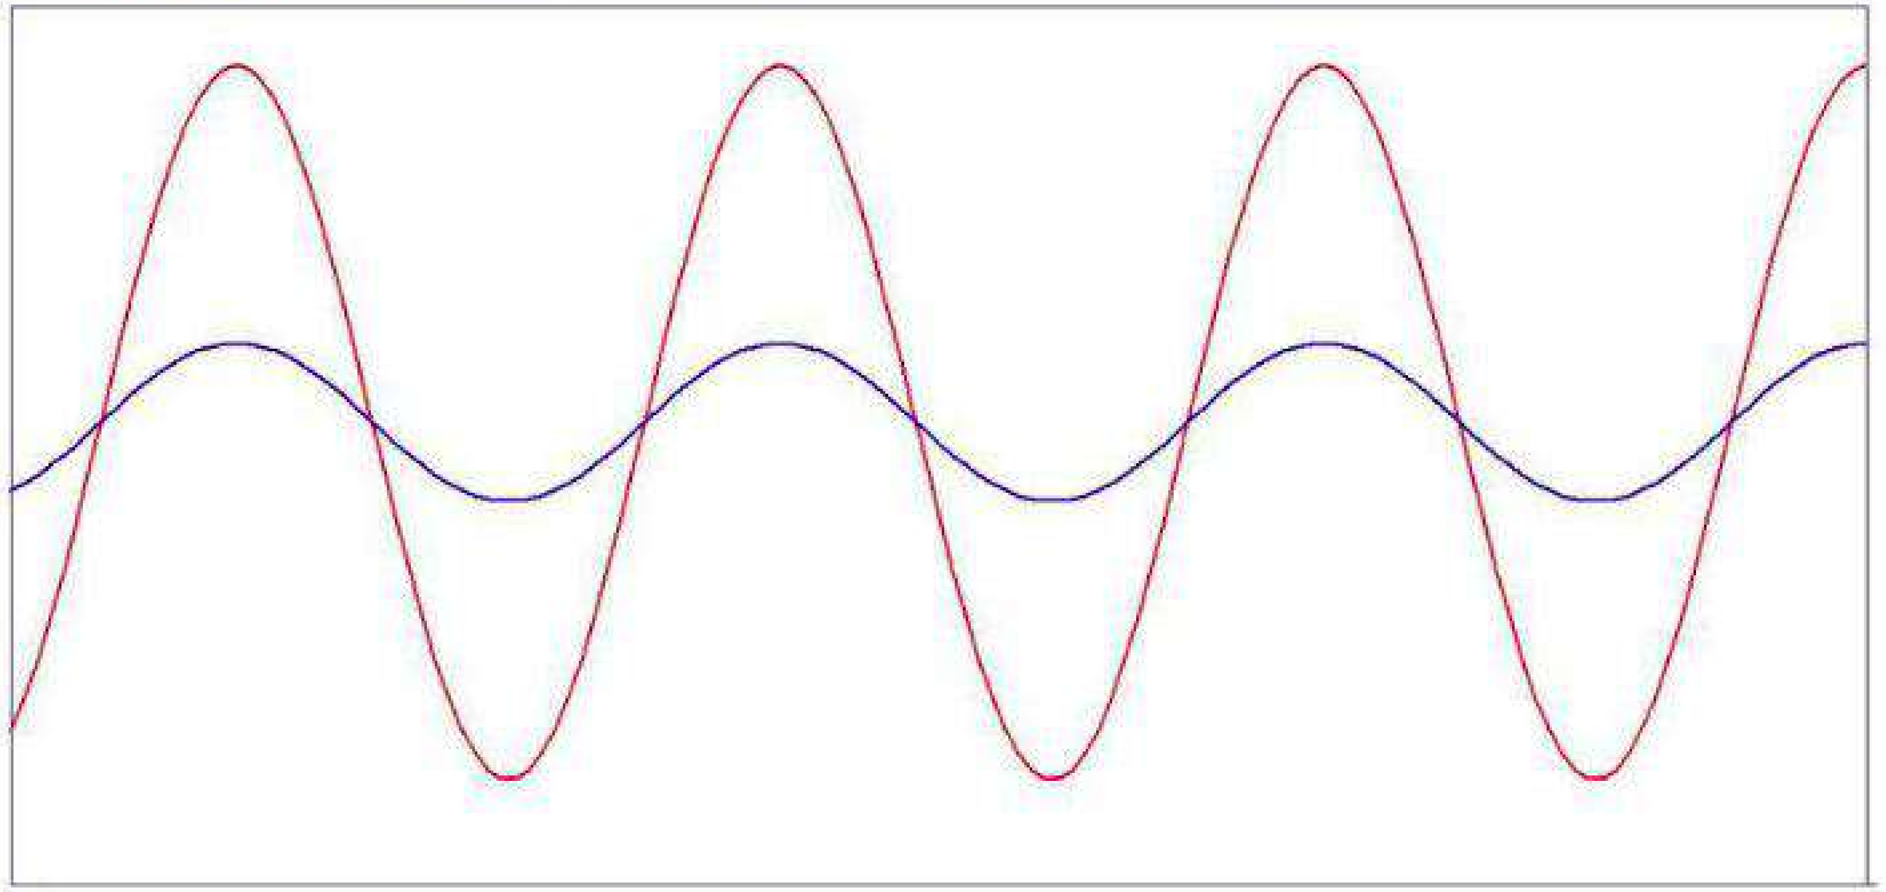

## PHASE

The *phase* is a more difficult concept, but it represents how far the sine wave is from some arbitrary starting point. The only way that this can be checked is to use a particular feature such as a peak, trough or mid-point.

Thus two sine waves can look exactly the same, but when plotted on the same chart the difference is obvious. This is shown in Figure 6.6, where two sine waves are super-imposed. They both have a wavelength of 41 days, which can be verified by checking the position of the peaks and troughs in each wave.

***Figure 6.6 – This shows two sine waves out of phase by half a wavelength. One has peaks at 1108, 1149 and 1190. The other has peaks at 1128 and 1169. The wavelength of both is 41 days and the amplitudes are identical.***

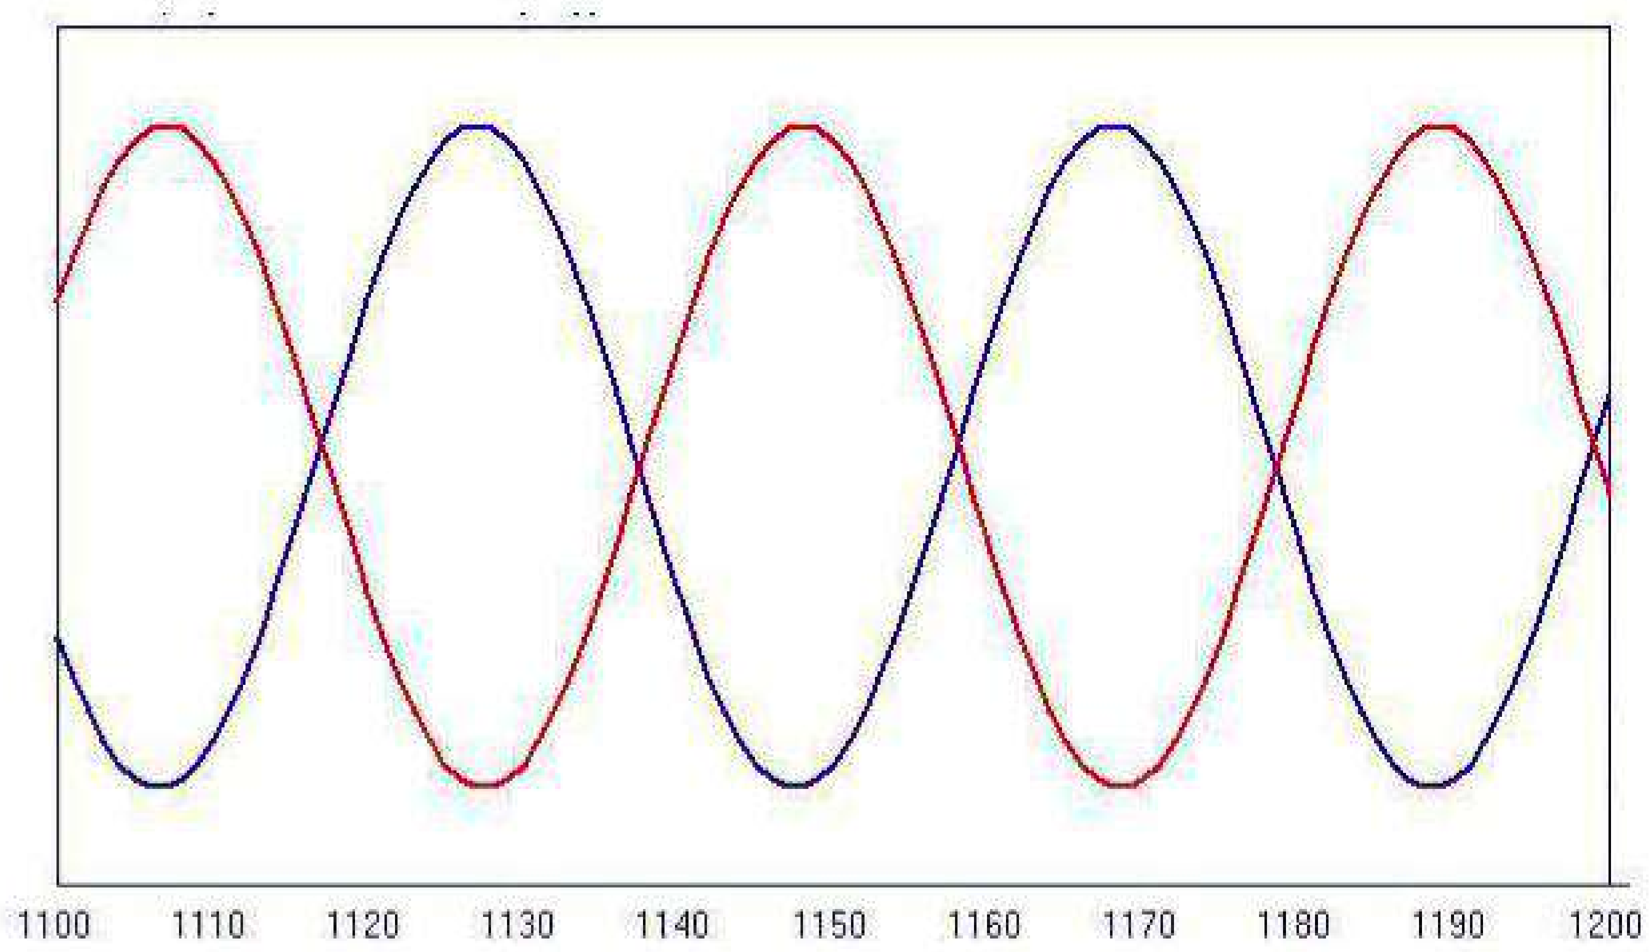


It can be seen that both of the sine waves also have the same amplitude.

The fact that they are exactly half a wavelength out of phase means that the peak of one is at the same position as the trough of the other. If the phase difference increases then the peaks (and troughs) of the two waves will separate even more. Of course, when the phases differ by an exact multiple of the wavelength, the peaks and troughs will coincide again. In such a case the two waves are essentially the same.

## HARMONICS

This is the term used to describe those frequencies which are an exact multiple of the original frequency. Thus the second harmonic has twice the frequency of the original, the third harmonic three times the frequency of the original, and so on.

Since wavelength is the inverse of the frequency, the second harmonic will be half of the wavelength of the original, the third harmonic one-third of the wavelength of the original and so on. The amplitude of the harmonics decrease as we move from second to third and so on. This decrease in amplitude as we move to shorter wavelengths is a fundamental property that will have implications for the relationship between the amplitudes of cycles of different wavelength in the same security.

Harmonics are mentioned here because it is quite possible that they exist in the cycles in market data. As will be shown later, many cycles are related by such a multiple. They also pass the test for change in amplitude.

## EQUATION OF A SINE WAVE

Now we know the three variables that make up a sine wave, we can put them together in the form of an equation relevant to the stock, foreign exchange or commodities markets.

This equation takes the following form:

Y = amplitude x sin (F + W x t)

Where:

Y = value in currency, index points, etc.

W = 2 x Pi/N

N = wavelength

Pi = 3.142

t = time in same units as N

F = a measurement of phase as discussed earlier.

Therefore in order to find the value for Y at some time (t) in the future for a given sine wave, we simply need to establish the wavelength (N), the amplitude and the phase.

If we have a section of sine wave sufficient to give us values for amplitude and wavelength we can calculate the equation which describes it. We can then plot the future path of this sine wave. We can adjust the phase until the calculated extension to the existing sine wave has a maximum or minimum exactly one wavelength forward from the last true maximum or minimum. The result of doing this is shown in Figure 6.7.

***Figure 6.7 – Once the wavelength and amplitude of the sine wave (solid line) are known, then its equation can be derived. It can then be extrapolated into the future (broken line).***

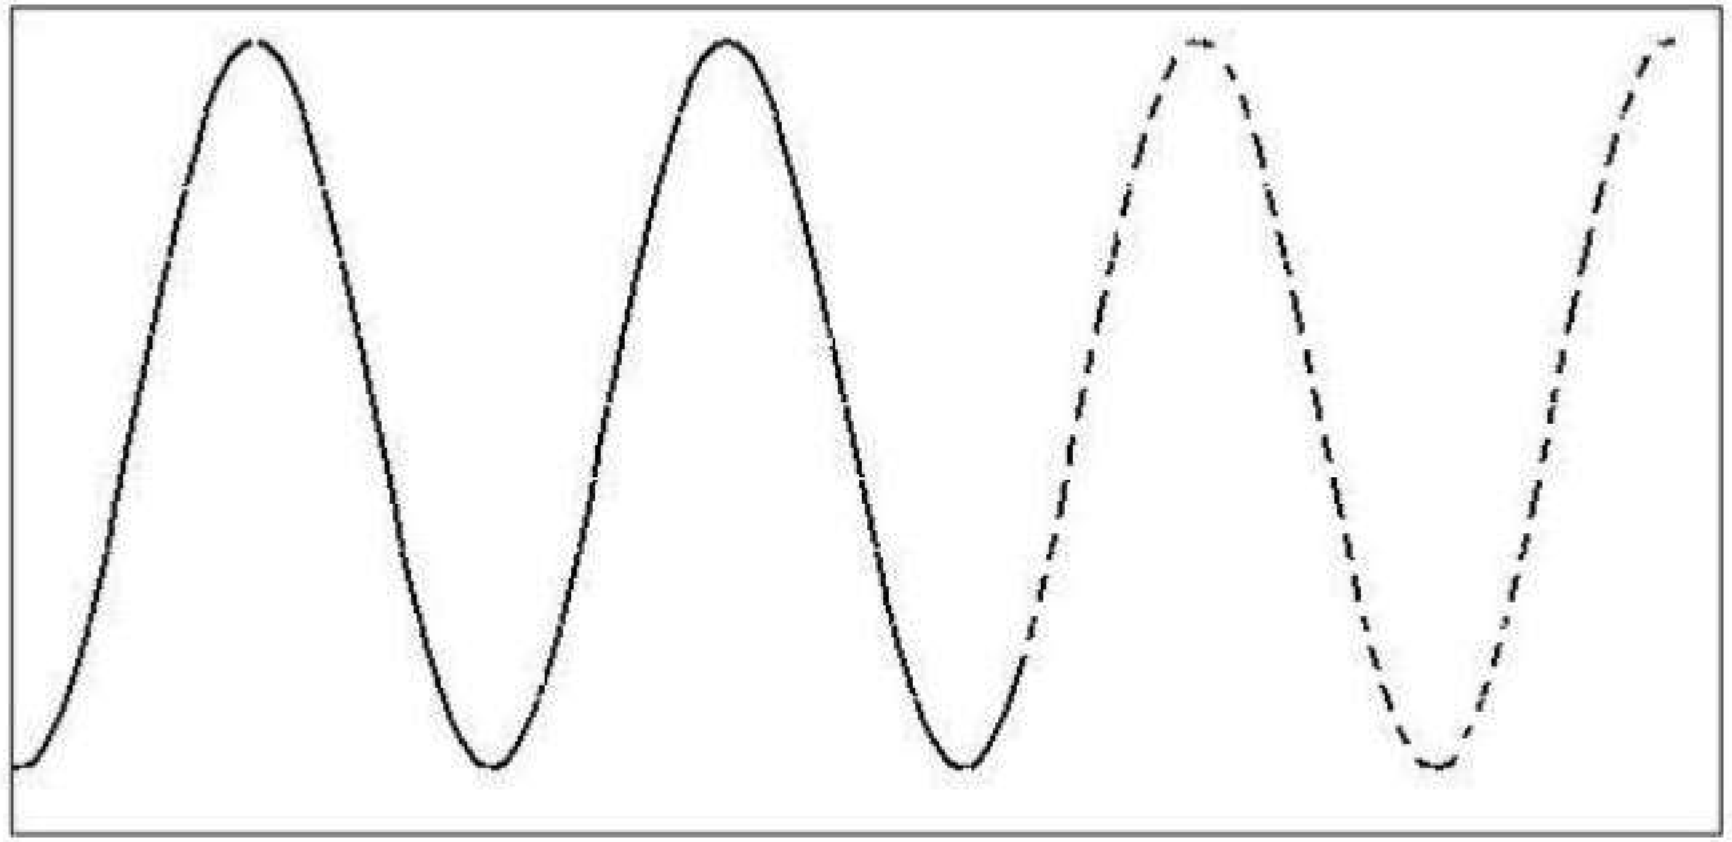

At this point we have covered the theory behind sine waves which can be derived mathematically. These sine waves are of course regular in the sense that the wavelength, amplitude and phase remain constant. This is why we can extrapolate any such sine wave into the future.

## Cycles in the Stock Market

When we look at waves that are present in the stock market, we will not see such regularity. We will find that each of the three components (wavelength, amplitude and phase) will not remain constant but will vary over the course of the evolution of the price data.

### NOMINAL WAVELENGTH

Since all three of the components of a sine wave in the market can change, it could be argued, philosophically, that the change in one or more components means that we are no longer observing the same sine wave. However, it helps the discussion to take the view that we can accept variations in the three components of a sine wave and consider it to be the same wave, as long as these variations are limited in nature. The way to do this is to use the term '*nominal wavelength*' to describe its average wavelength, around which limited change is allowable. We could also take the view that since a change in phase would shift the position of peaks and troughs from their originally estimated position, then an apparent change in wavelength may not be an actual change in wavelength but a change in phase. However, the cause of an apparent change in wavelength is immaterial; it is the fact of a change that is important.

This can be demonstrated by isolating a cycle of nominal wavelength by the methods discussed in later chapters; such a plot is shown in Figure 6.8.

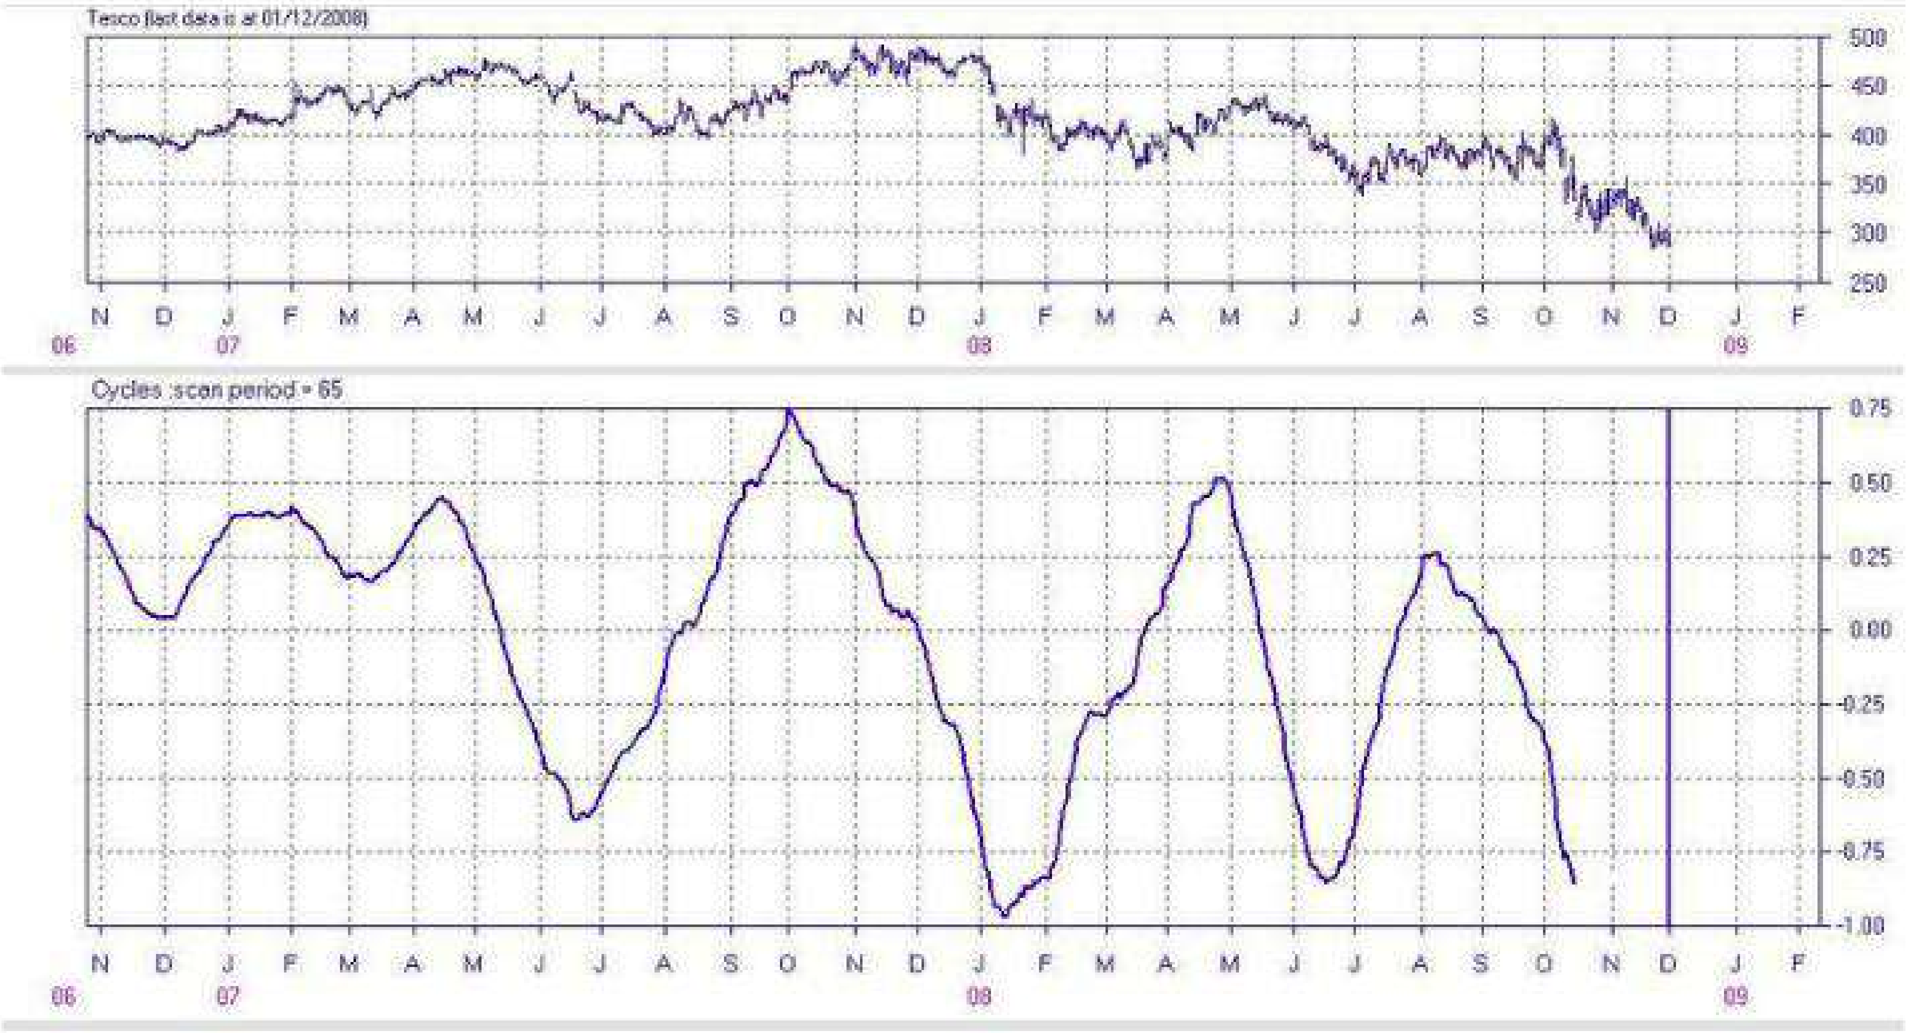


**Figure 6.8 – The cycle of nominal wavelength 130 days in the Tesco stock price (lower panel). It can be seen that the amplitude and wavelength are changing. The lumpy appearance is due to the influence of cycles of shorter wavelength. The stock price is shown in the upper panel.**

This shows quite clearly the variation in amplitude and wavelength of a cycle with a nominal wavelength of 130 days. The term '*nominal*' is used here because this is the average wavelength around which the increase and decrease is occurring. Quite clearly the two peaks at October 07 and May 08 are further apart than the peaks at May 08 and August 08. It can also be seen that the amplitude is decreasing gradually.

Working backwards, from right to left, a peak around February 2007 would have been expected if this gap between peaks was to be maintained. However, on inspection, it can be seen that the peak has been distorted by the addition of a cycle of much shorter wavelength, which is bottoming out at almost the exact place where the nominal 130-day cycle would have been topping out.

A much better appreciation of the variation in wavelength and amplitude can be seen from Figure 6.9. Here a longer history of the cycles can be seen. The time axis is now shown in days rather than dates. In this mode it is easier to establish the distance apart of peaks and troughs.

**Figure 6.9 – A longer section of cycles is shown here (lower panel) which highlights the more extreme changes in amplitude and wavelength which can occur. The stock price is shown in the upper panel.**

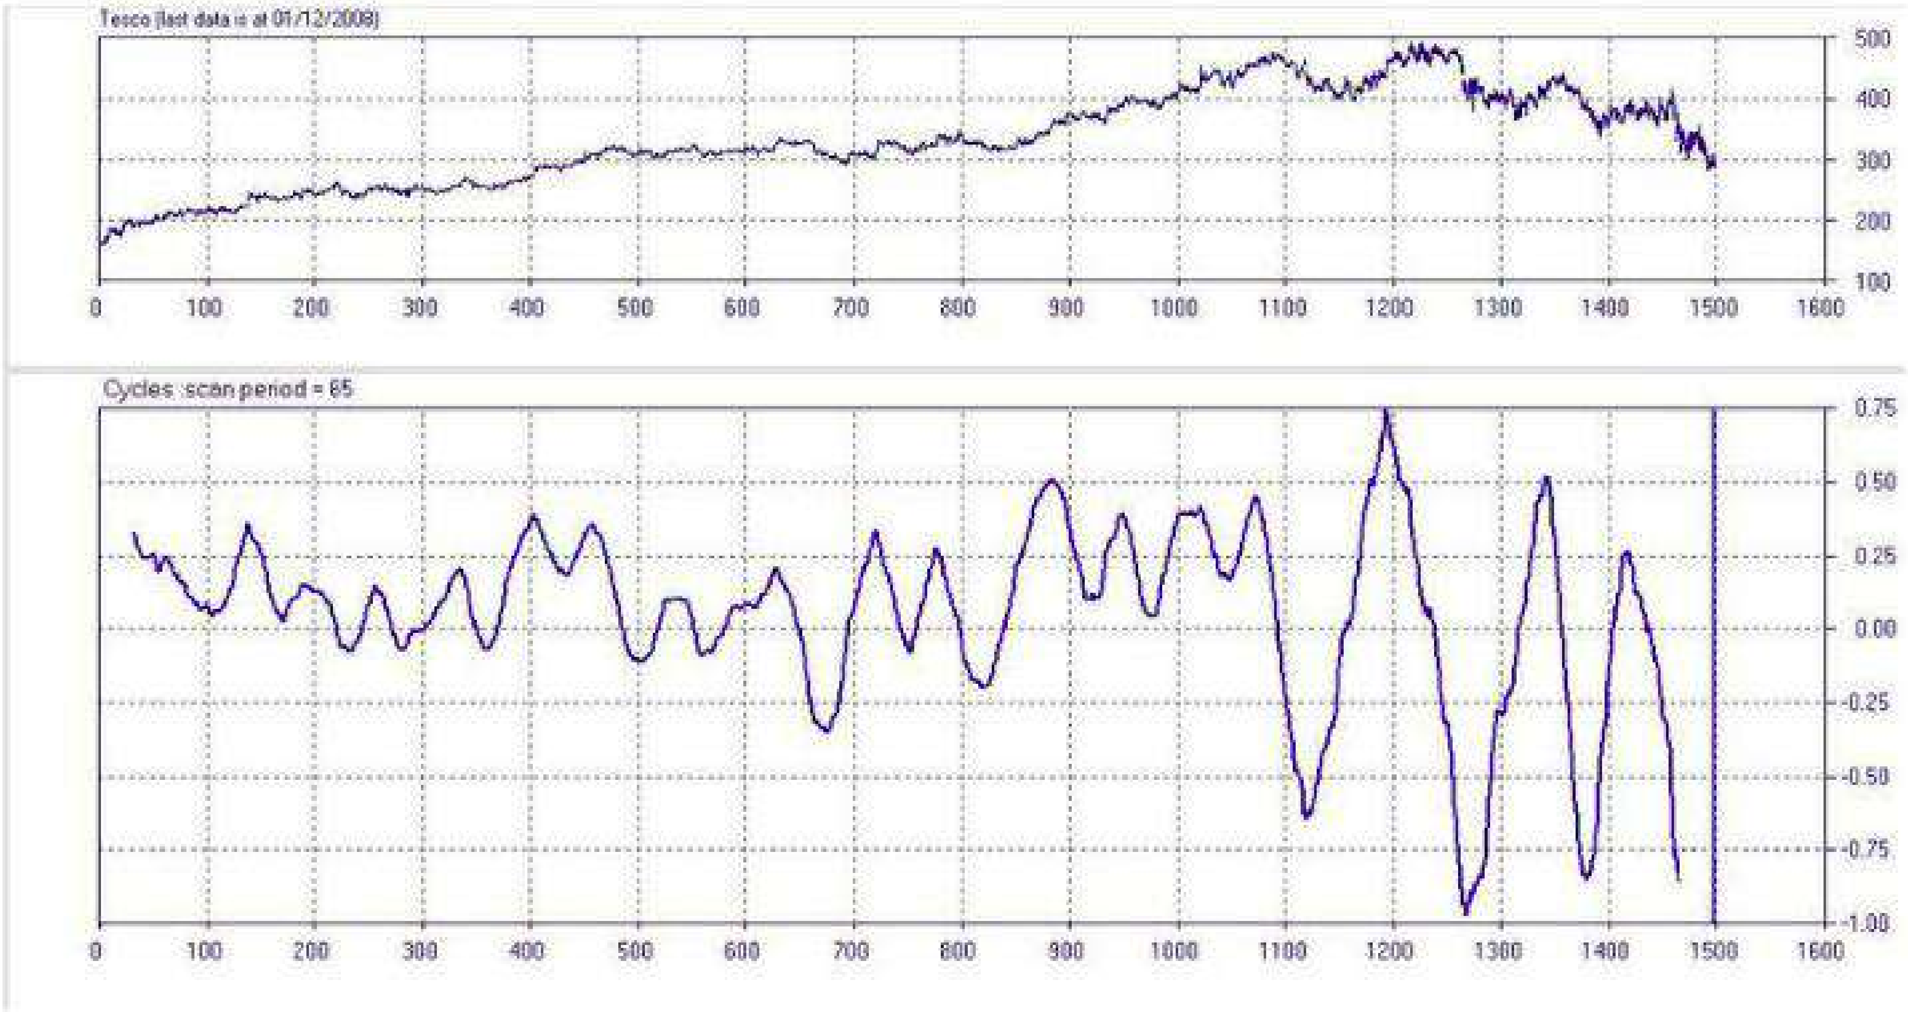


From the appearance of Figures 6.8 and 6.9 it might well be thought that using cycles in market data to predict future movement is fraught with such difficulty that it is not worth the attempt. However, by using better mathematics than used to produce Figures 6.8 and 6.9, the result shown in Figure 6.10 was obtained. Here there is a much smoother trace for the cycle. Because of that smoothness, the cycle produced by this method, although still varying in amplitude and wavelength at the right-hand section of the plot, shows a much less marked change.

**Figure 6.10 – The nominal cycle of wavelength 113 days in American Express (lower panel). The last half of a wavelength (broken line) has been extrapolated. The stock price is shown in the upper panel.**

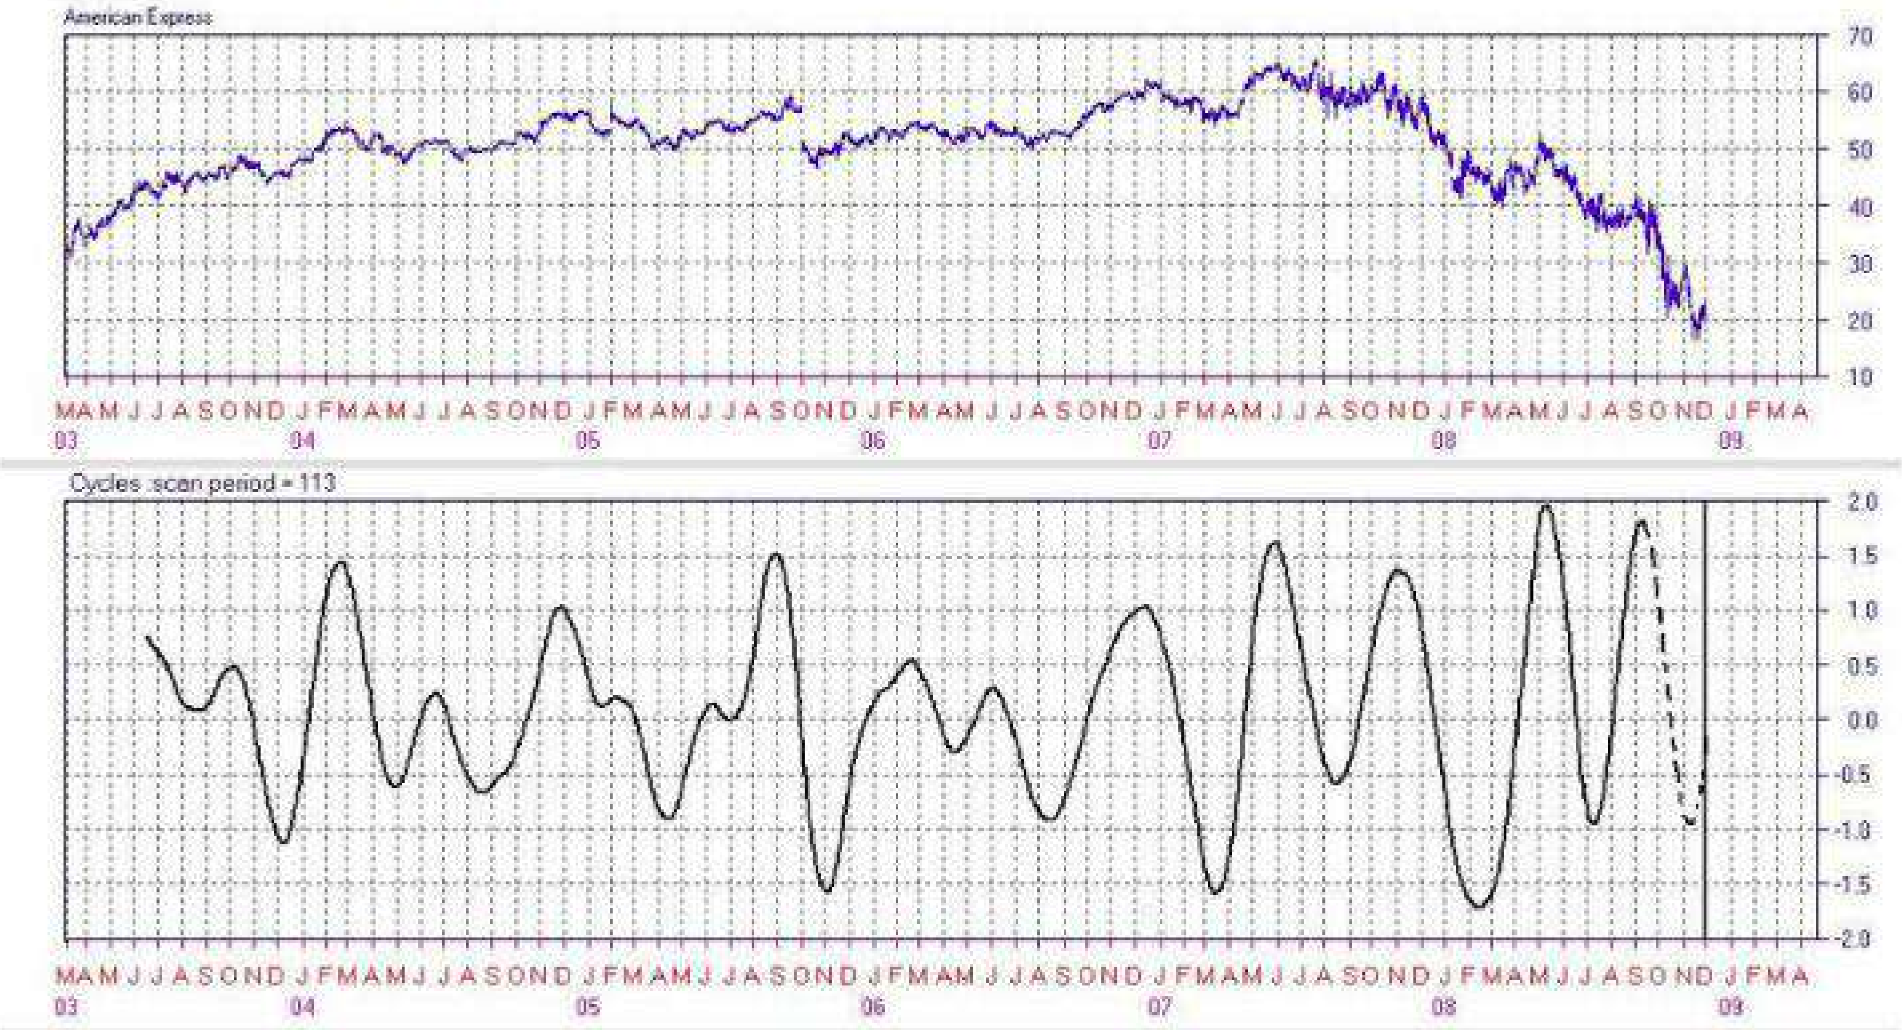


Note that the last half of a wavelength of the cycle is shown as a broken line. This is to show that this is an extrapolation from the last true calculated point, which terminates about half a wavelength back in time from the latest data point. The reason for this is explained in a later chapter.

The extrapolation is carried out by taking an average of the wavelength and amplitude over the last few sweeps of the cycle to calculate the equation of the equivalent sine wave. As will be seen later, such an extrapolation is only valid where the last few sweeps of a cycle fall within specified limits (e.g. 20%) of variation in both wavelength and amplitude.

The correctness of the extrapolation of this nominal 113-day cycle can be seen in Figure 6.11, where the data was taken several months later. The actual downward leg of the cycle agrees exactly with the predicted path from the earlier extrapolation.

**Figure 6.11 – Several months later, it can be seen that the extrapolation was exactly correct. The downward leg of the cycle agrees with the predicted position from Figure 6.10.**

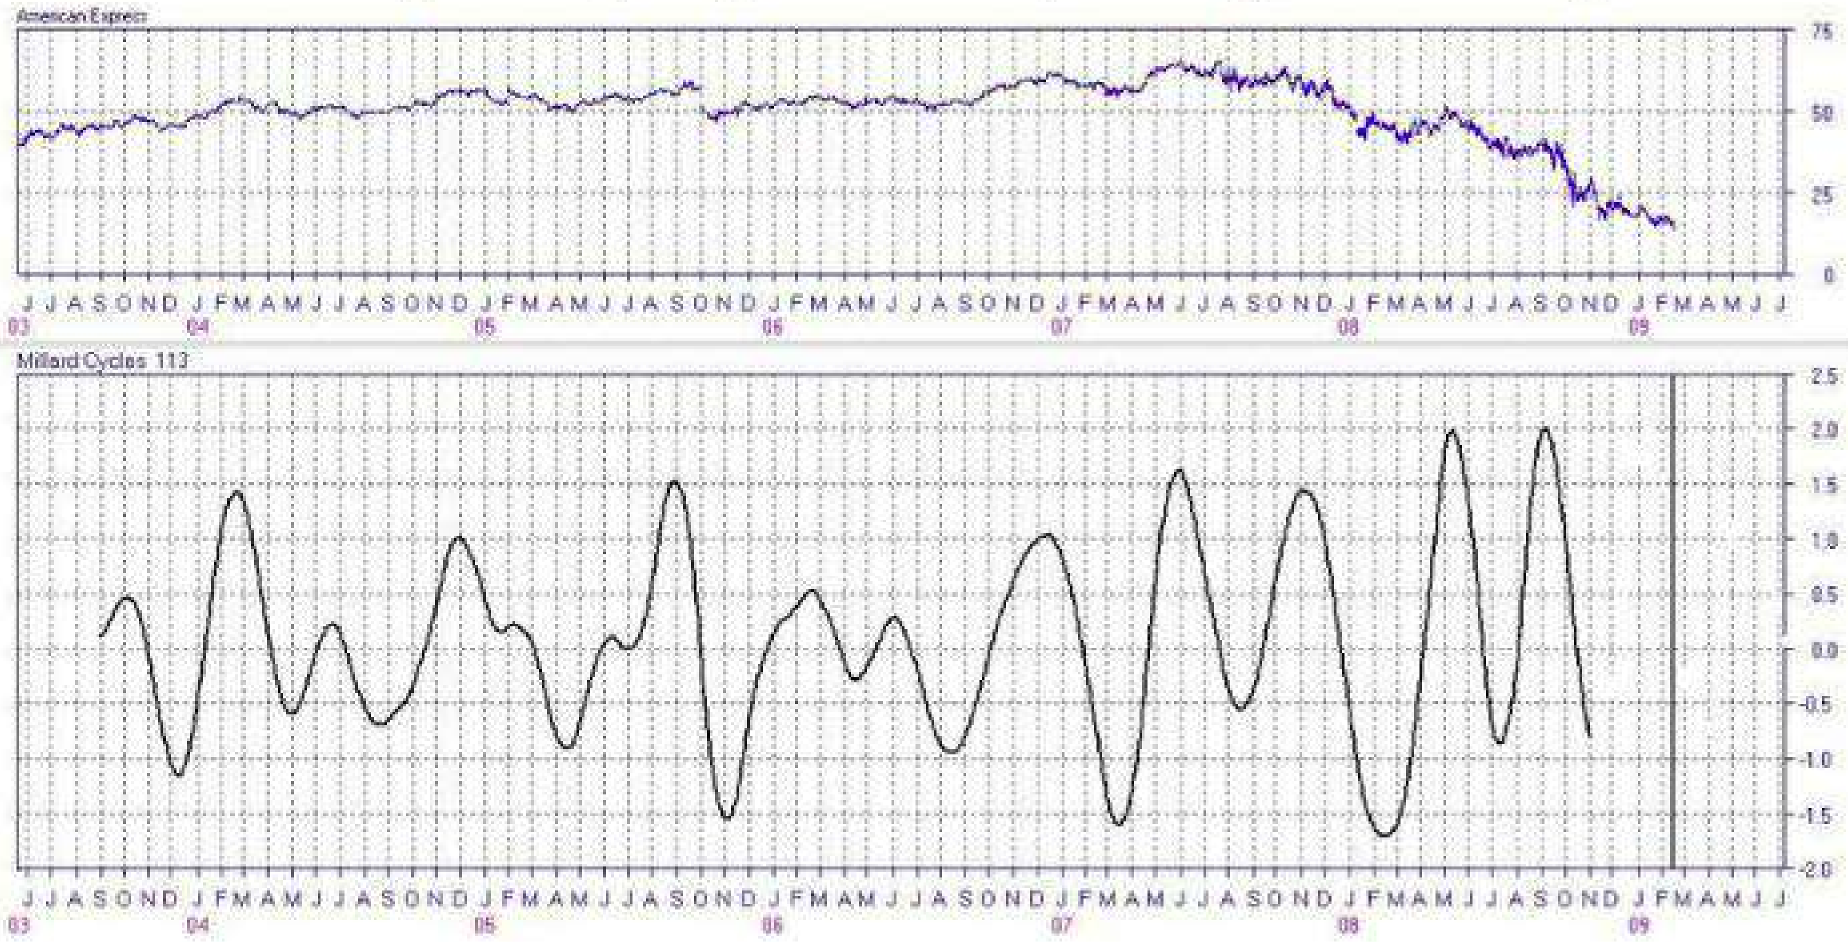


An example of a very regular cycle is found in the foreign exchange market in the Australian dollar versus the euro, as shown in Figure 6.12. The nominal wavelength in this case is 83 days. The most striking element of this chart is the regularity of wavelength. This stays within 20% across the whole plot.

**Figure 6.12 – A plot of the Australian dollar versus the euro (upper panel). The cycle shown is the nominal 83-day cycle (lower panel).**

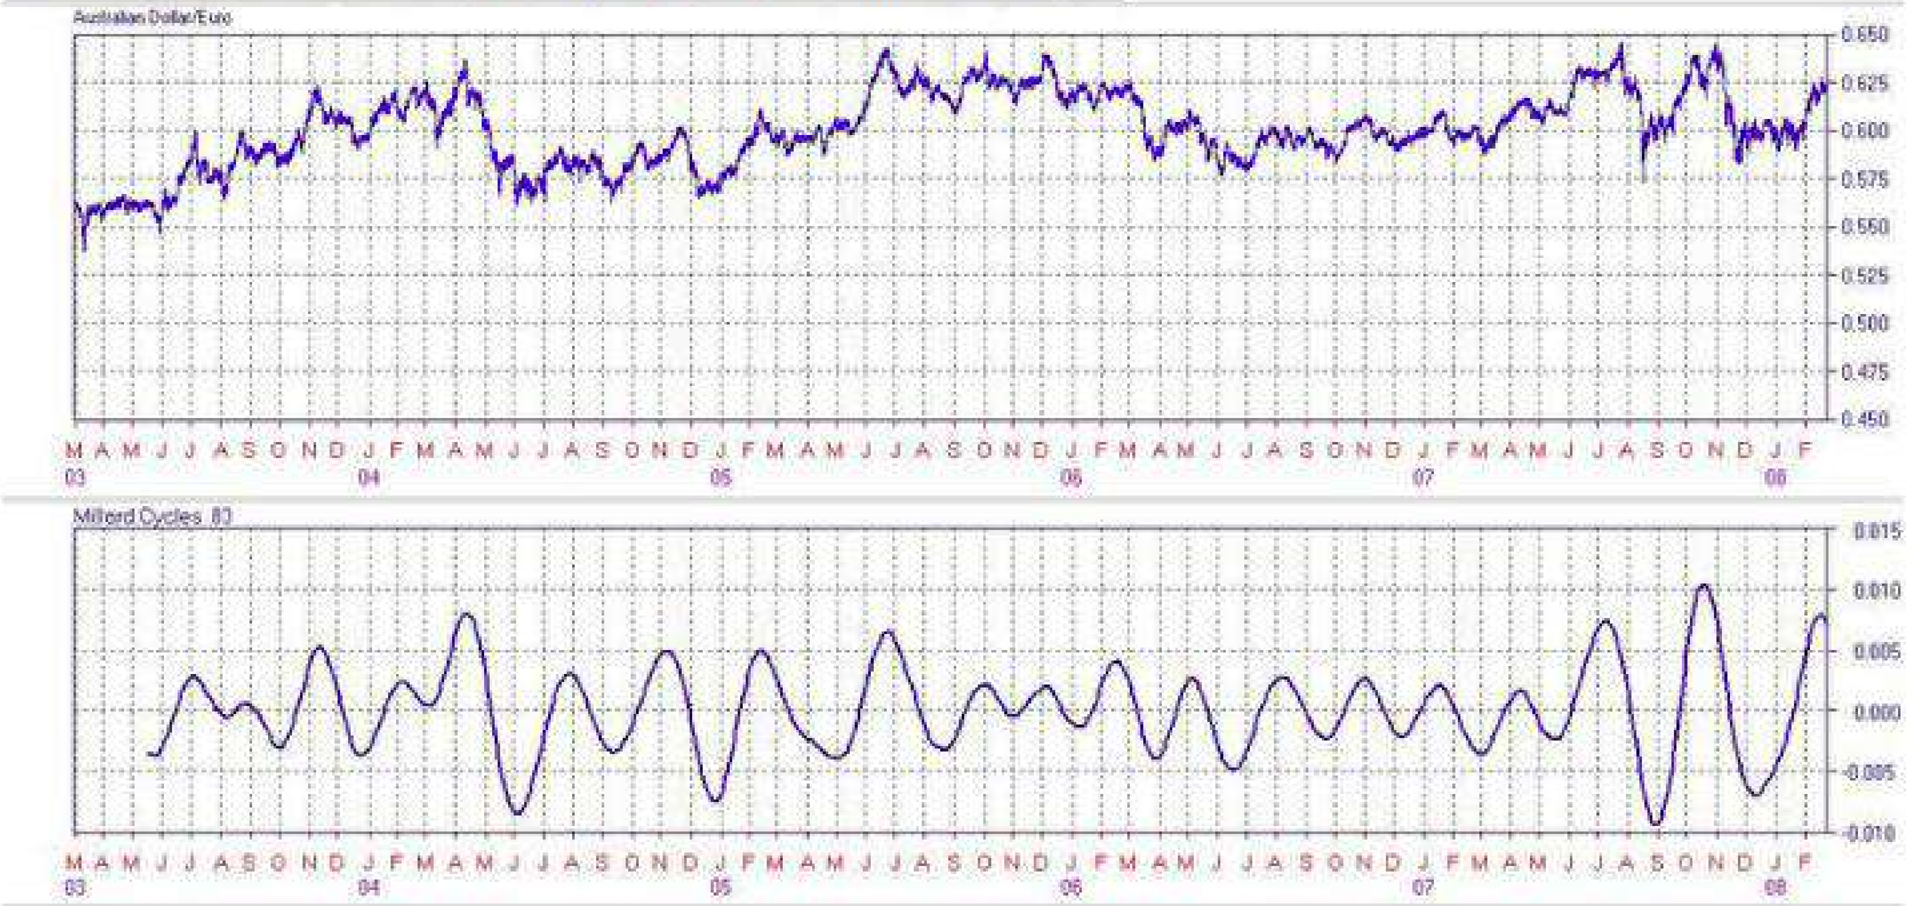


Thus in this case only the amplitude seems to vary to any great extent. Even so, from January 2006 to April 2007 the variation in amplitude was quite limited.

From these examples it can be seen that in many cases a cycle will pass through a state where it is relatively stable for a number of sweeps. During other time periods it is unstable in the sense that both the wavelength and the amplitude are varying by an amount which is outside the allowable limit of 20%.

*In order to predict trends into the future, only cycles which have been in this relatively stable state over the last few complete sweeps can be used.* The validity of this approach in the prediction of future trends will become apparent in later chapters, where predicted trends can be compared with actual trends. In some cases an extraordinary degree of prediction of trends will be evident.

## Research on Market Cycles

The existence of cycles in the market has of course been known for a long time. Thus there is the four-year presidential cycle in the United States, there is the Halloween indicator, in which it appears that buying around Halloween time and selling in May is more often profitable than not and so on.

In my book *Channel Analysis*, published in 1990 and reissued in 2009, I listed the known cycles in stock market data. This was based upon the work of many analysts throughout the last century.

***Table 6.1 – Known cycles in stock market data.***

*days are business days, i.e. five per week.

The amplitudes of these various cycles appear to be directly proportional to the wavelength for the shorter-term cycles, but the increase tails off as we move to higher wavelengths.

The cycles of wavelength greater than 4.3 years are of no practical use to the trader.

It is interesting to note that J.M. Hurst gave a different set of cycles in his book *The Profit Magic of Stock Transaction Timing*.

***Table 6.2 – Market cycles according to J.M. Hurst.***

Hurst stated quite categorically that under what he called the Commonality Principle these cyclic components were present in all securities.

In my research on those cycles that are stable over the last few sweeps I have arrived at a different conclusion from those which are illustrated in these two tables. In both the US and UK markets there seems to be very little correlation between stable cycles in individual securities. This is also true of the foreign exchange markets.

Of course, finding that at the time this book was being written (January to June 2009) there is no correlation does not mean that there was no such correlation in the past. Naturally the observation that cycles spend most of their history in a state of instability means that it makes sense that there is very little correlation, since an exact wavelength cannot be put on an unstable cycle.

Thus each individual security must have the wavelengths of its own cycles established and from this its own future trend determined without reference to any other security, other than, as discussed in Chapter 2, a consideration of the overall state of the market.

### COINCIDENCE OF CYCLE PEAKS OR TROUGHS

When two or more cycles with different wavelengths are present and they happen to coincide in the position of a peak or trough, then the question arises as to how far into the future this will happen again. This depends entirely upon the wavelengths involved.

The problem is solved by using the method for finding the lowest common multiple of the cycles. Take the case of two cycles with wavelengths, say, 20 days and 75 days that have coincident peaks today. We list multiples of each by multiplying each by 2, 3, 4, 5 and so on. We will get:

20, 40, 60, 80, 100, 120, 140, 160, 180, 200, 220, 240, 260, 280, 300

75, 150, 225, 300

From this we can see that the first common multiple is the value of 300. This means that the peaks will be only coincident again 300 days from now. During that time, the cycle of wavelength 20 will have seen 15 peaks and the cycle of wavelength 75 will have seen four peaks.

Unfortunately, because of the variation in wavelength in market cycles, this coincidence will be difficult to predict with any certainty.

# 7. Trends and the Market

The *Oxford English Dictionary* defines a trend as "general direction and tendency". *Wikipedia*, that rather less authoritative but more popular source of information, defines a market trend as "a prolonged period of time when prices in a financial market are rising or falling faster than their historical average, also known as 'bull' and 'bear' markets, respectively". What is meant by '*historical average*' is not explained.

*Wikipedia* also defines trend estimation as "the statistical analysis of data to extrapolate trends". It goes on to further define a market trend in these terms: "this principle incorporates the idea that market cycles occur with regularity and persistence. This belief is considered to be generally consistent with the practice of technical analysis". Note the use of 'trend' and 'cycles' in the same definition.

In this book we will not consider the traditional way of presenting uptrends and downtrends, as shown in Figure 7.1, to be representative of trends in the sense that we shall be using throughout this book.

**FIGURE 7.1 – THE TRADITIONAL WAY OF SHOWING UPTRENDS AND DOWNTRENDS IS SHOWN BY THE LINES ON THIS CHART.**


There are two main reasons for not proceeding any further with the traditional way of drawing trends as exemplified by Figure 7.1.

These are:

1. Trends drawn in this way do not satisfy the requirements of a mathematical trend.
2. There are no numbers upon which mathematical and statistical calculations can be performed which can lead to probabilities.

## Mathematical Trends

By '*mathematical trends*' we mean trends that are derived from a mathematical formula, such as a straight line, a polynomial, etc. In other words the trend is represented by a mathematical equation. On the other hand the straight lines shown in Figure 7.1, like any other straight lines, have intrinsic equations to describe them, but since they have been drawn on the chart using a ruler or its equivalent, these are not immediately known. It is not necessary for the reader to worry about the way in which such trends are calculated, since spreadsheet programs will do this. In the case of moving averages, all technical analysis software packages are able to carry out such a calculation. However, users of moving averages might not be aware that a moving average represents a trend, as discussed later.

### STRAIGHT LINE

A straight-line trend is shown in Figure 7.2. This is calculated on the closing prices over 100 days of the London Stock Exchange. The major way in which this straight line differs from those shown in Figure 7.1 is that the data oscillates around the line, i.e. some data points appear above the line and others below the line. This is a fundamental property of trends calculated in this way and is the reason that in this book it is this fundamental type of trend with which we will be concerned.

***Figure 7.2 – This shows a straight-line trend drawn on the closing prices of the London Stock Exchange. The chart has been produced by the Excel spreadsheet program.***

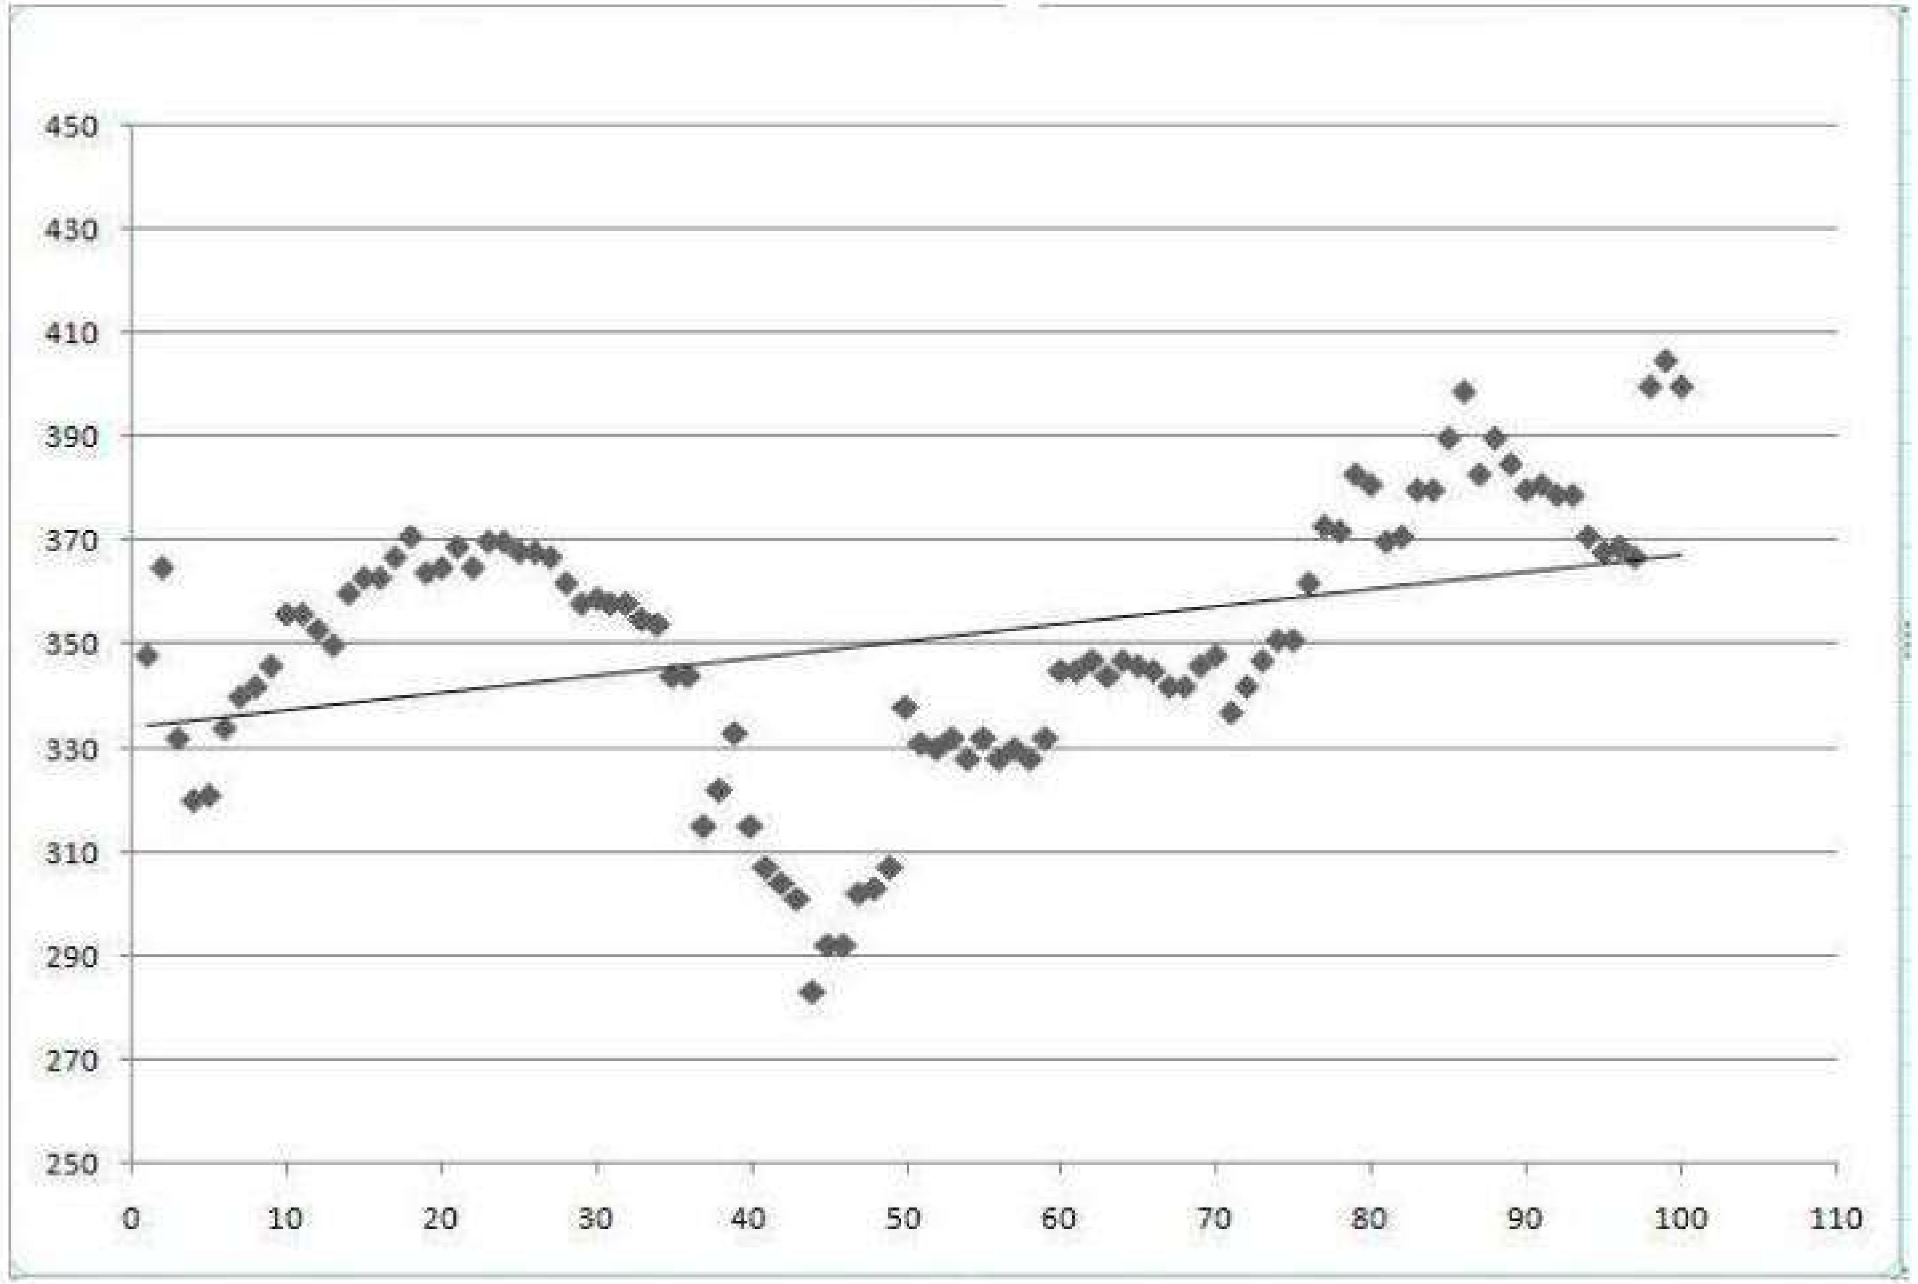


Since the trend is based on the known equation of a straight line, then obviously it can be projected into the future.

Whether the future trend can still be represented by this straight line is a different matter and the point at which this straight line fails to represent the trend of the data is not known in advance. However, various statistical tests can be applied, as each new data point arrives, that can help to determine whether or not the straight trend line for the new data is still valid.

### POLYNOMIAL

There are many other types of trend lines that can be applied to the data, for example a polynomial. The variable in this case is the order of the polynomial. Thus in Figure 7.3 we show a second order polynomial trend line (quadratic) drawn through the same data. Again we see the fundamental property that the data lies above and below the trend line. But, as is the case with the straight line, the quadratic equation of this line is known and therefore it can be projected into the future. As with the straight line, whether or not the trend line is still valid for future data points cannot be determined in advance.

***Figure 7.3 – A quadratic trend line has been drawn (using Excel) through the same set of data as that used in Figure 7.2.***

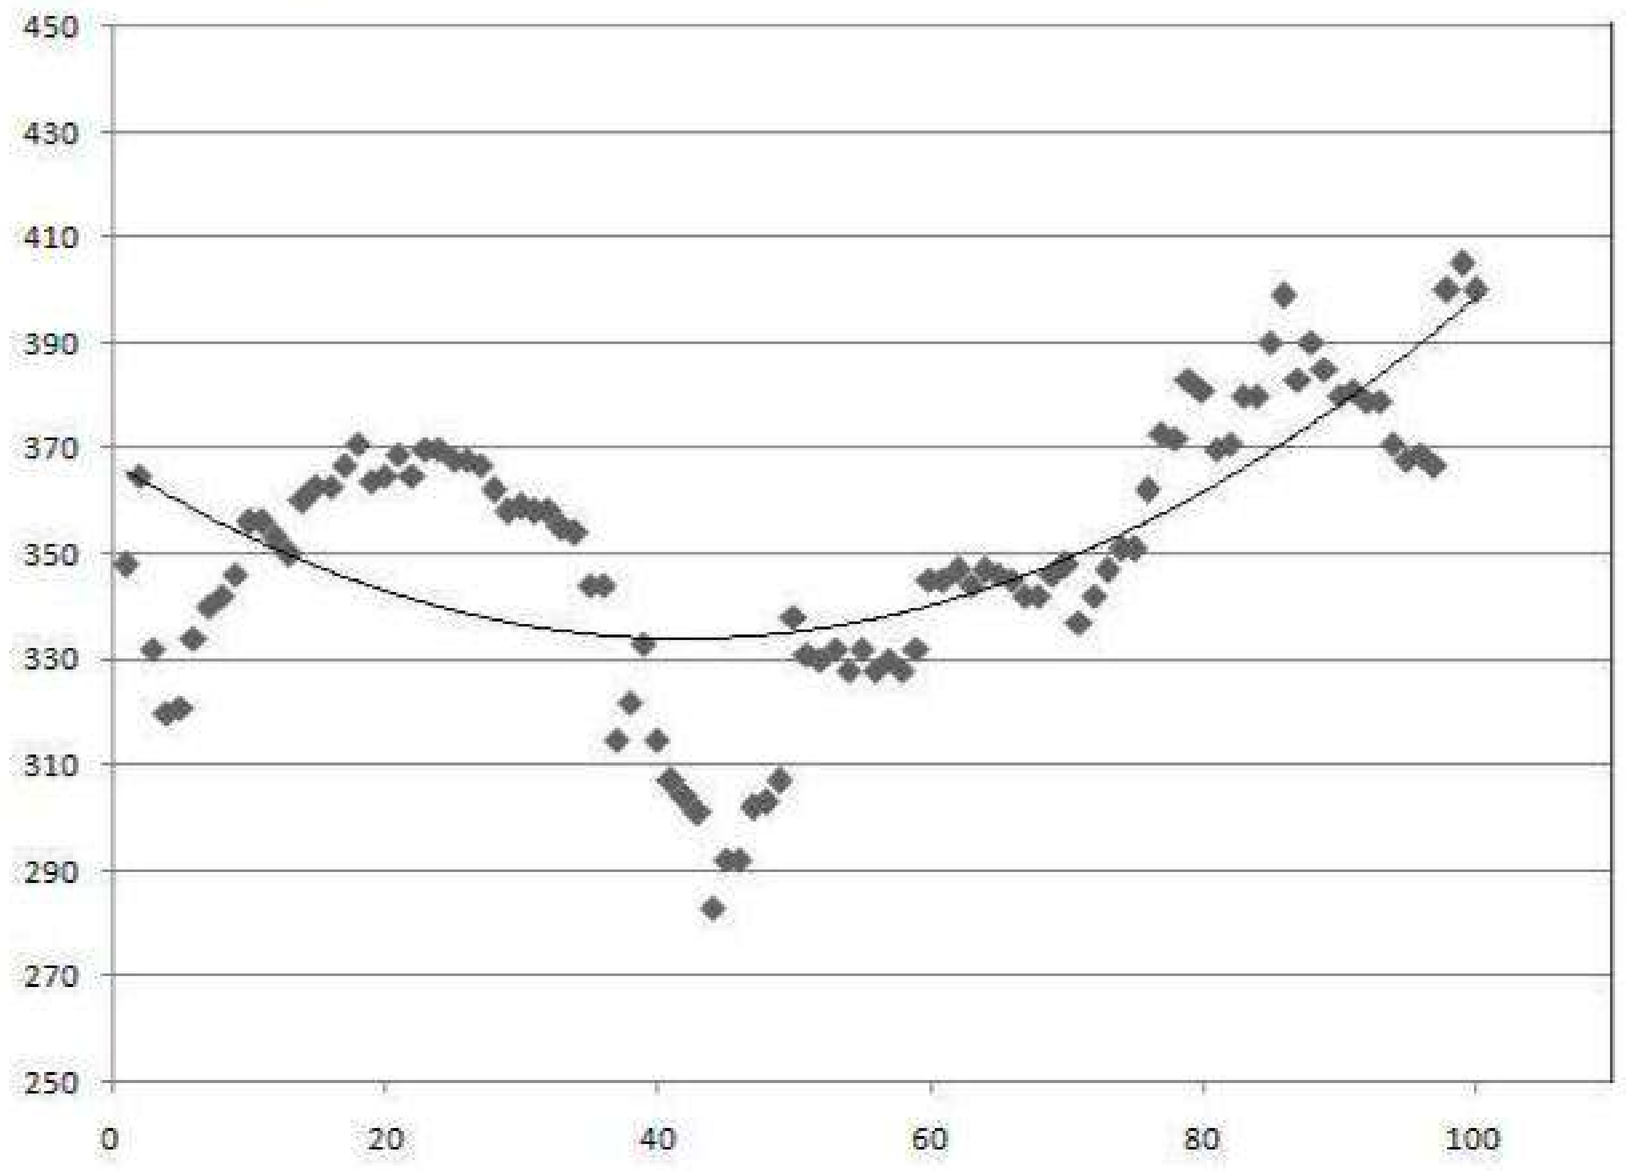


## CHANGE IN TREND LINE

It must be clearly understood that a trend line calculated from a data set will almost certainly not be the same as a trend line calculated from a portion of that data set. This can be seen in Figure 7.4, where only the first 50 points of the London Stock Exchange data have been used to calculate a quadratic trend line. Quite obviously this line does not marry up with the longer line which was plotted in Figure 7.3.

***Figure 7.4 – A quadratic trend line has now been drawn through the first 50 points of the data used in Figure 7.3.***

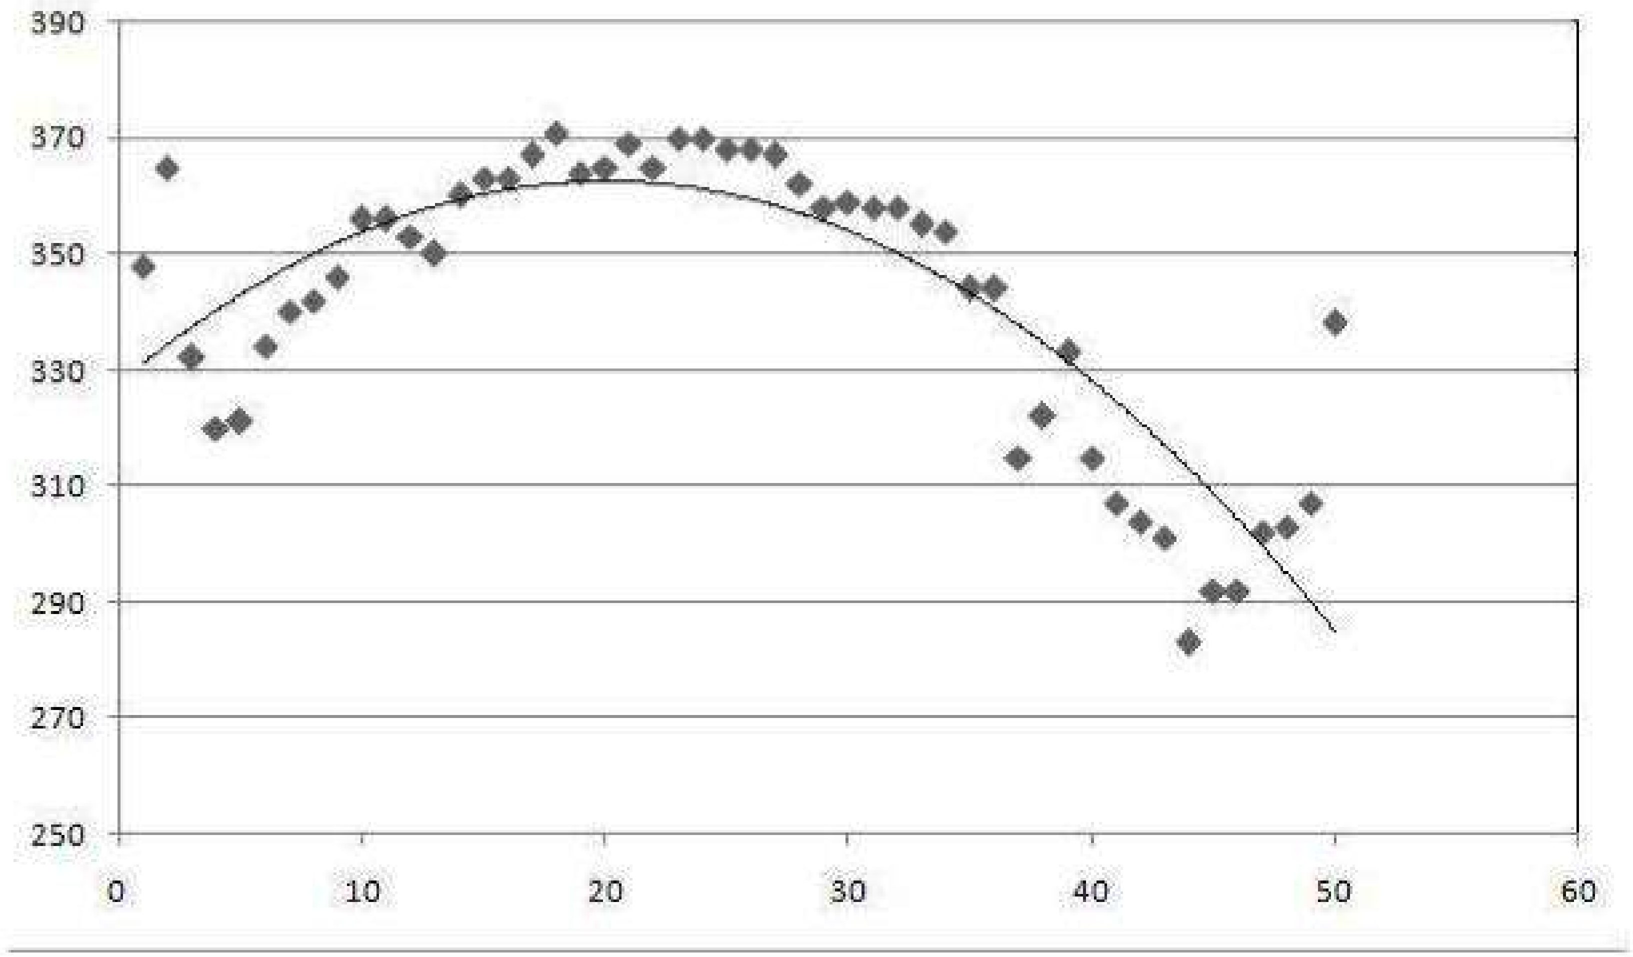


***Change with one data point***

As a new data point arrives, the only way in which the trend line will remain the same is if the new data point lies on the trend line. Any other position for the new data point will change the equation of the trend line and hence its position.

***Change with several data points***

Although the only way the equation of the line can remain the same is if the first new data point lies on it, this does not imply that the only way for a trend line to remain valid will be if all subsequent points to arrive also lie on the line.

If the first new point does not lie on the line, then the next point to arrive must lie on the other side of the projected trend line if there is to be no change from the original trend. Not only that, but there will be only one position for this new point that will result in the recalculated trend being the same as the original. In the case of market data this is of course unlikely.

Once a number of new points have arrived then a recalculation may well result in the equation of the new line being the same as that of the original. Again, this is unlikely in the case of market data, although there might not be a large change.

It has to be accepted therefore that predicted trends in market data will change as new data points arrive. This is why it is extremely important that trends are recalculated each day. If the recalculated trend lies within the expected range then all is well. If not – if the trend is seen to move in the wrong direction – then action must be taken.

The sequence of Figures 7.5 to 7.7 shows the effect that each new data point has on the quadratic trend which was drawn on the chart in Figure 7.4. A careful inspection of the position of this trend line relative to the last six points shows that it is gradually moving higher as each new point is added to the data set. This is of course because each new point is lying at a higher value than the last calculated value of the trend line.

***Figures 7.5 to 7.7 – Note the small change in the latest position of the trend.***

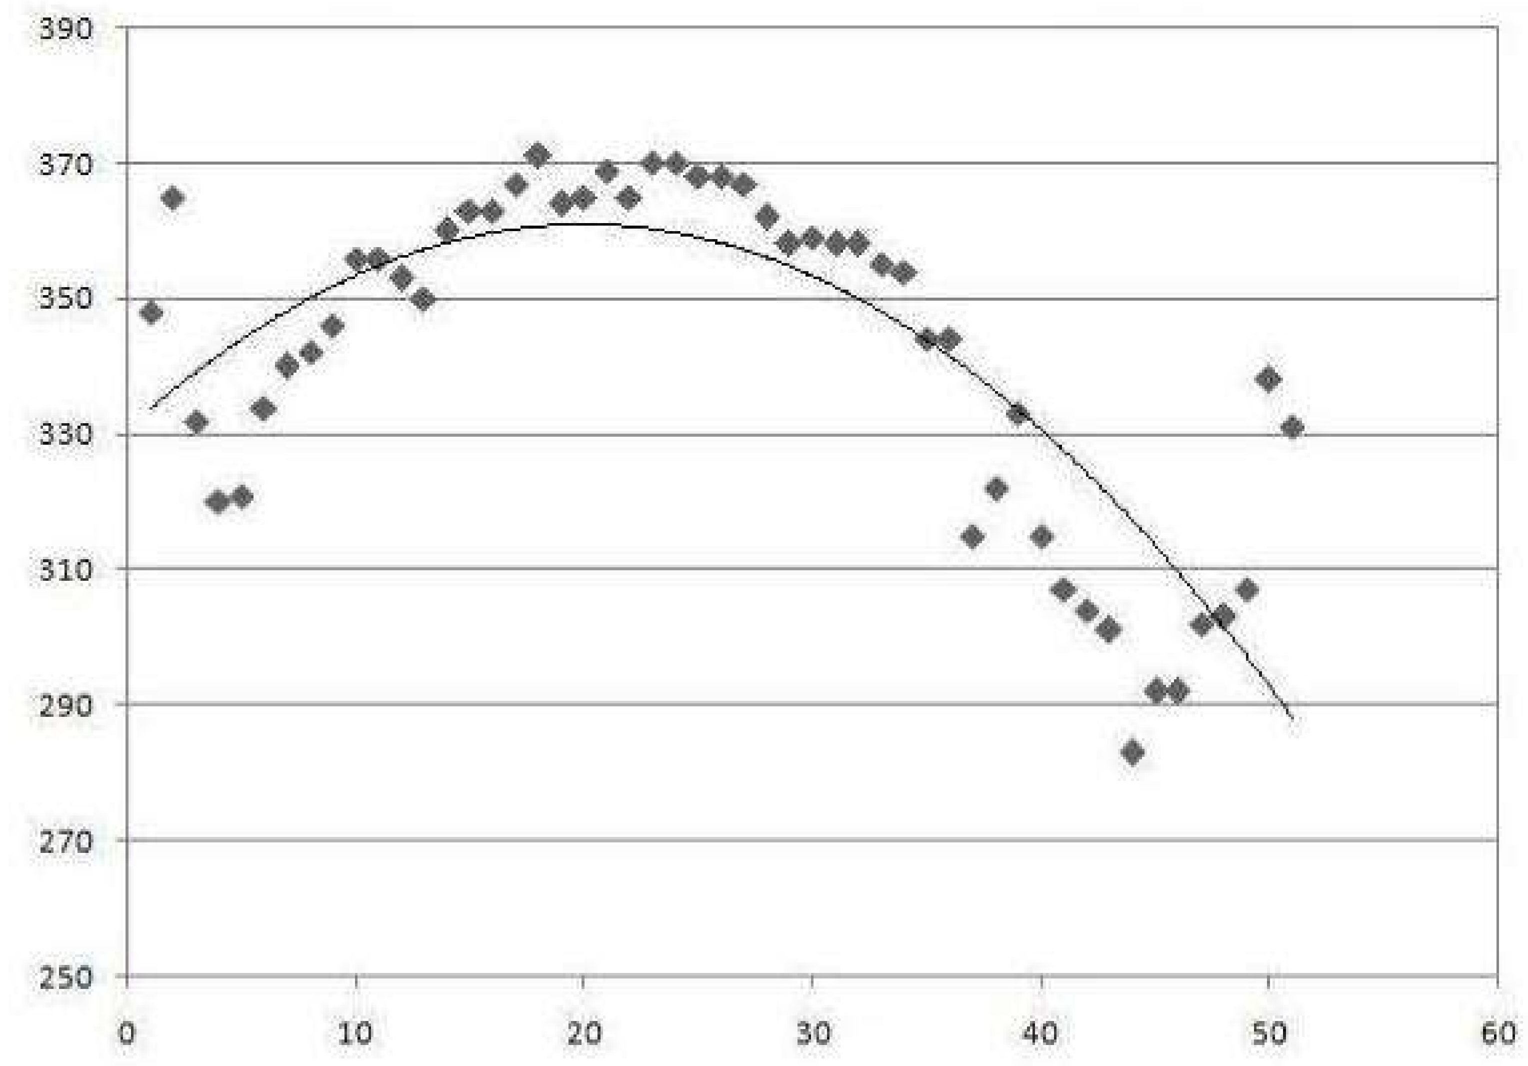

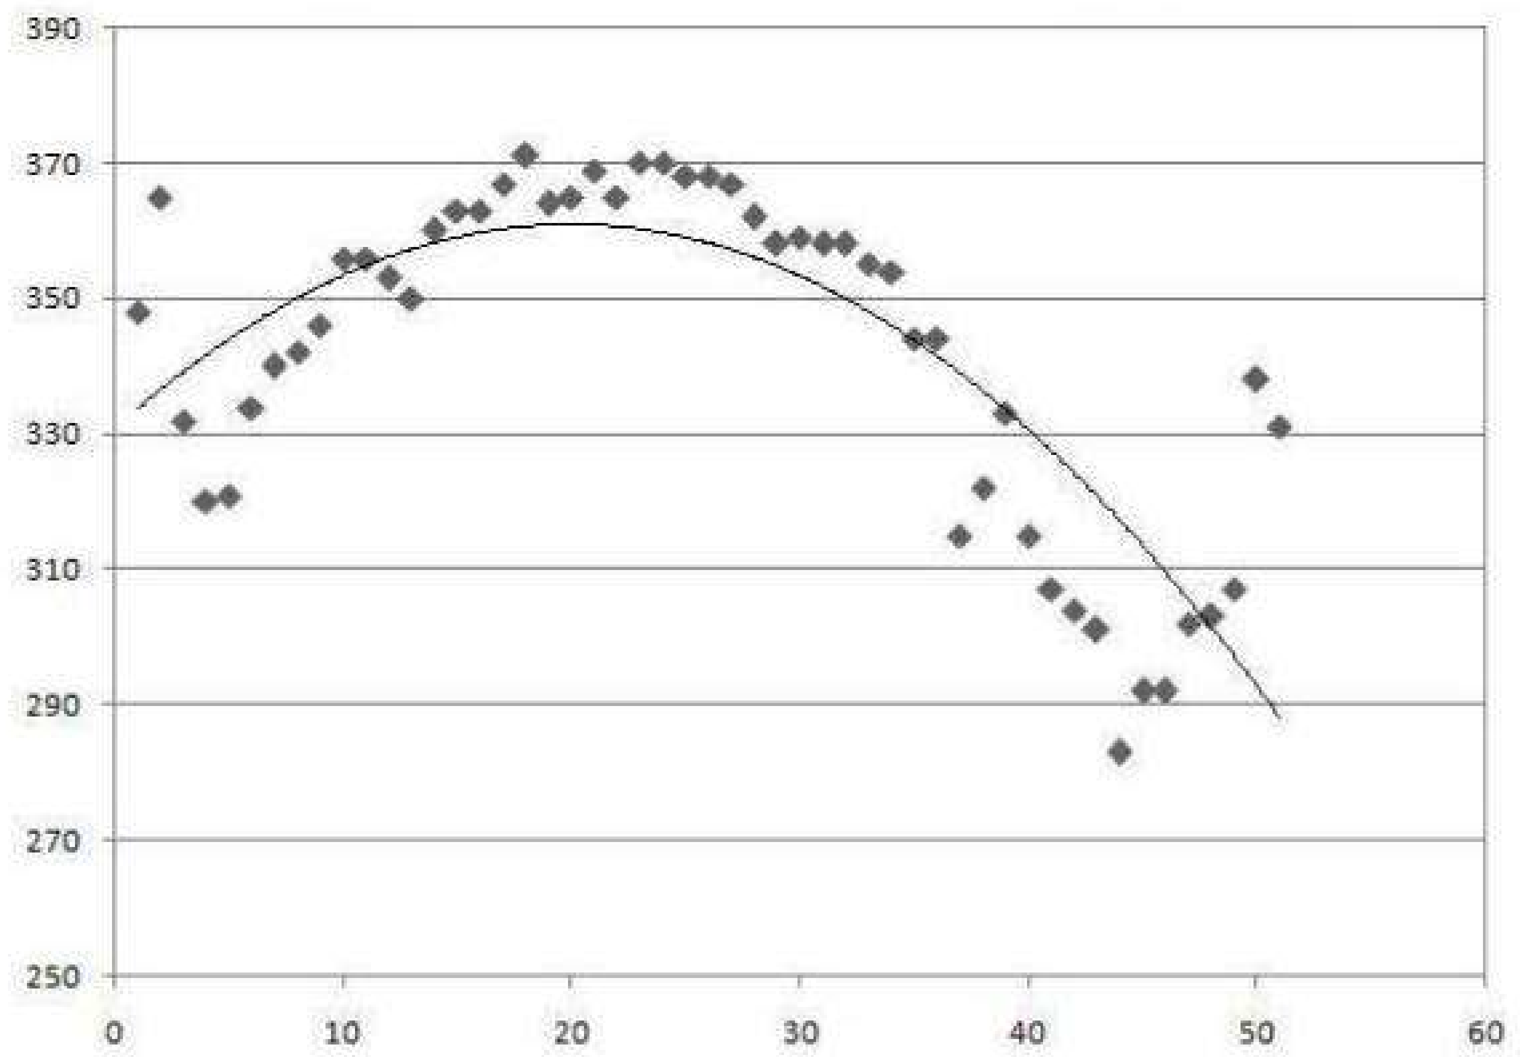


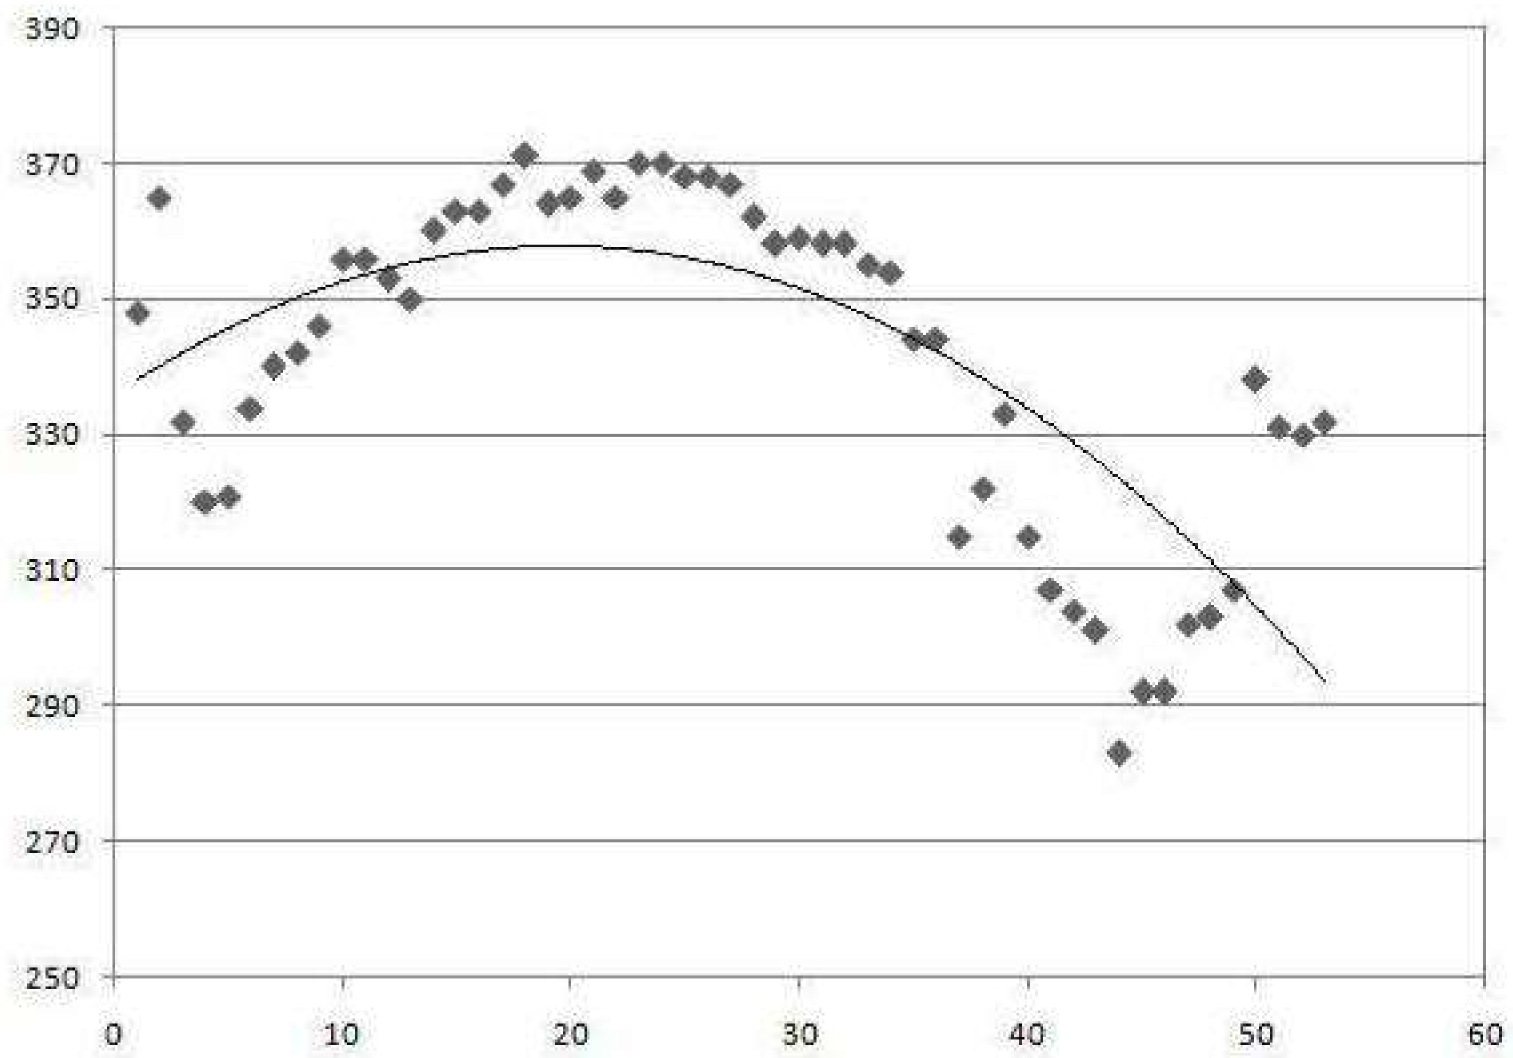


***Goodness of fit***

So far we have used two distinct types of trend lines – straight lines and quadratic lines. The simple way in which to tell if these are reasonable representations of the trend in the data is to apply a statistical calculation. This is a huge topic and those interested may of course search for more information on the internet. A useful website can be found at: www.curvefit.com.

One method is to use a quantity known as $R^2$. This quantity lies between 0 and 1.0. The higher the value, the better is the fit of the line to the data from which it is derived. A value of 1.0 can only be obtained if all of the data points lie on the line – this will never happen with stock market data. A value of 0 means that there is no fit, with a horizontal line placed at the average Y value for the data being just as meaningful.

The quantity $R^2$ is calculated from the squares of the vertical distances of the data point from the line. Squares are used so that negative distances will be equally as important as positive distances.

This book is not intended to delve too much into the realm of this type of statistics. I am simply pointing out that computing trend lines is not a haphazard process, but that there are mathematical processes underlying the calculations to give confidence that the approach is correct.

When using a spreadsheet such as Excel to produce a trend line the value for $R^2$ is calculated for you and can be displayed if required. Thus for the quadratic trend lines shown in Figures 7.5 to 7.7, the values for $R^2$ are 0.699, 0.631 and 0.575 respectively. Clearly in this particular instance the quadratic is not a good trend line. It was only used here as an example to show the effect of the arrival of new points on the overall calculation.

## MOVING AVERAGES

Another form of trend line is that produced by using moving averages. The variable in this case is the period used for the average. There are also various ways of plotting such averages. They can be plotted with no lag, or with a lag which is usually half of the period used. In this case the averages are termed centred. This type of average will be used exclusively in this book except in this current section, where un-lagged averages are used to illustrate this same issue of change of trend line as each new data point arrives.

Figure 7.8 shows a five-point moving average of the same set of data from the London Stock Exchange that was used for the previous examples. Note how much more responsive the average is to the position of the last few data points compared to the quadratic trend line.

***Figure 7.8 – A five-point moving average of 50 points of the London Stock Exchange data.***

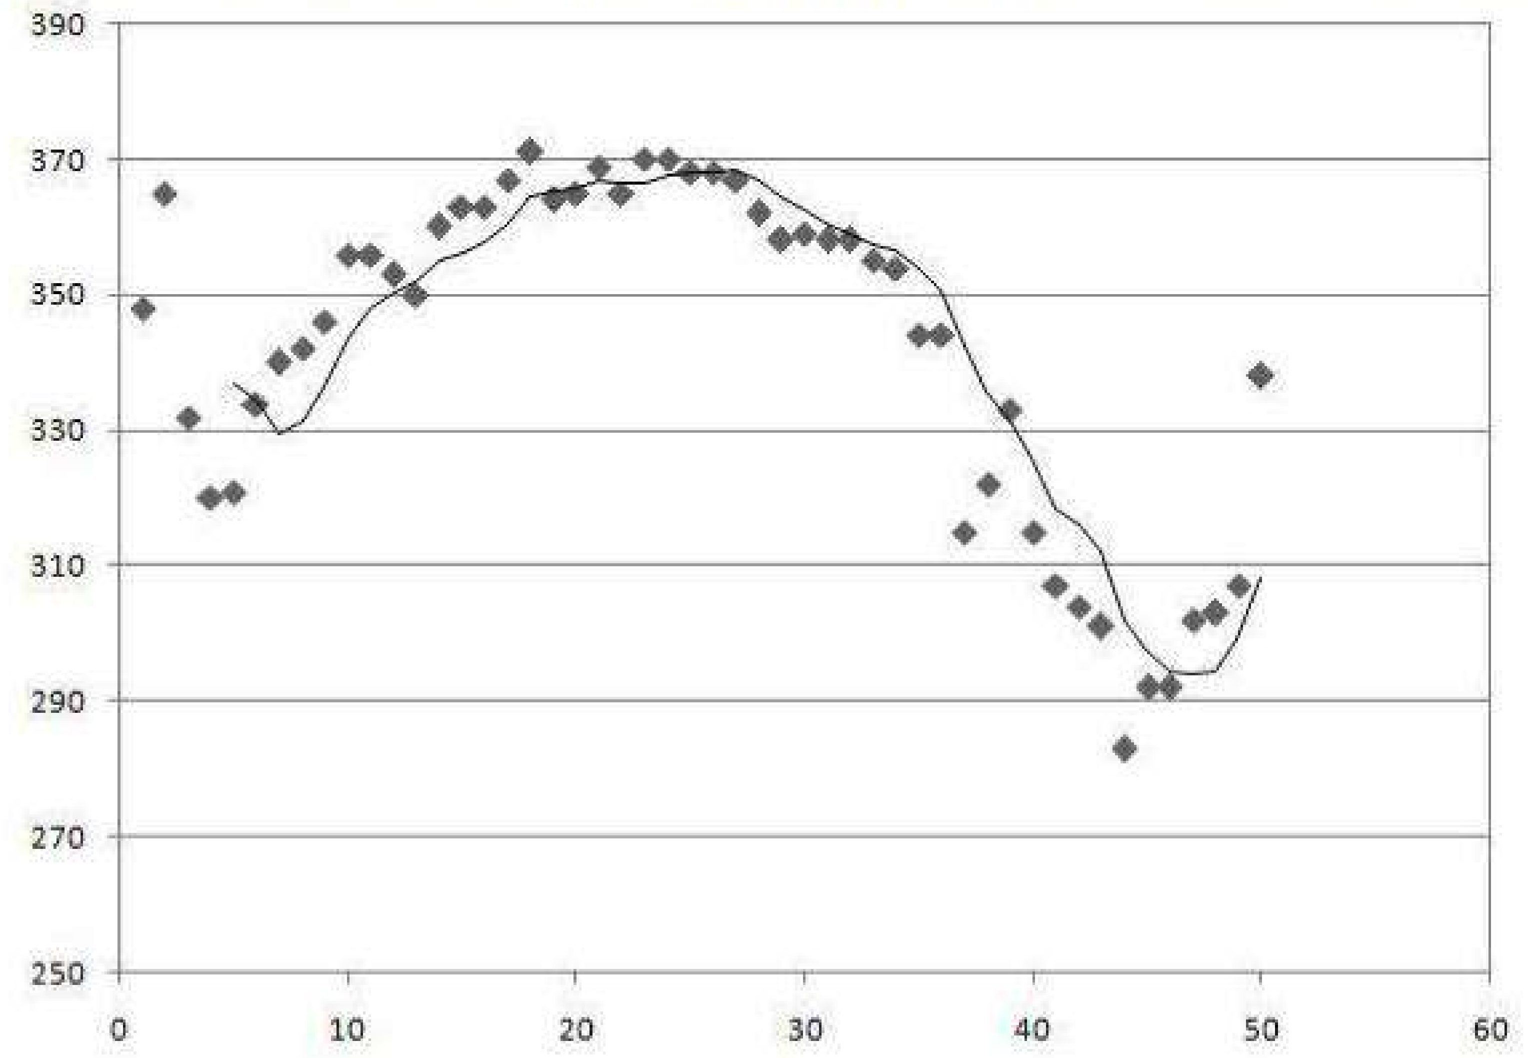


Shown in Figures 7.9 to 7.11 is the change in the plot of the average as new points arrive. Although only slight, there is a fall-off in the rate at which the average is rising.

***Figures 7.9 to 7.11 – Note the small change in the latest position of the trend line.***

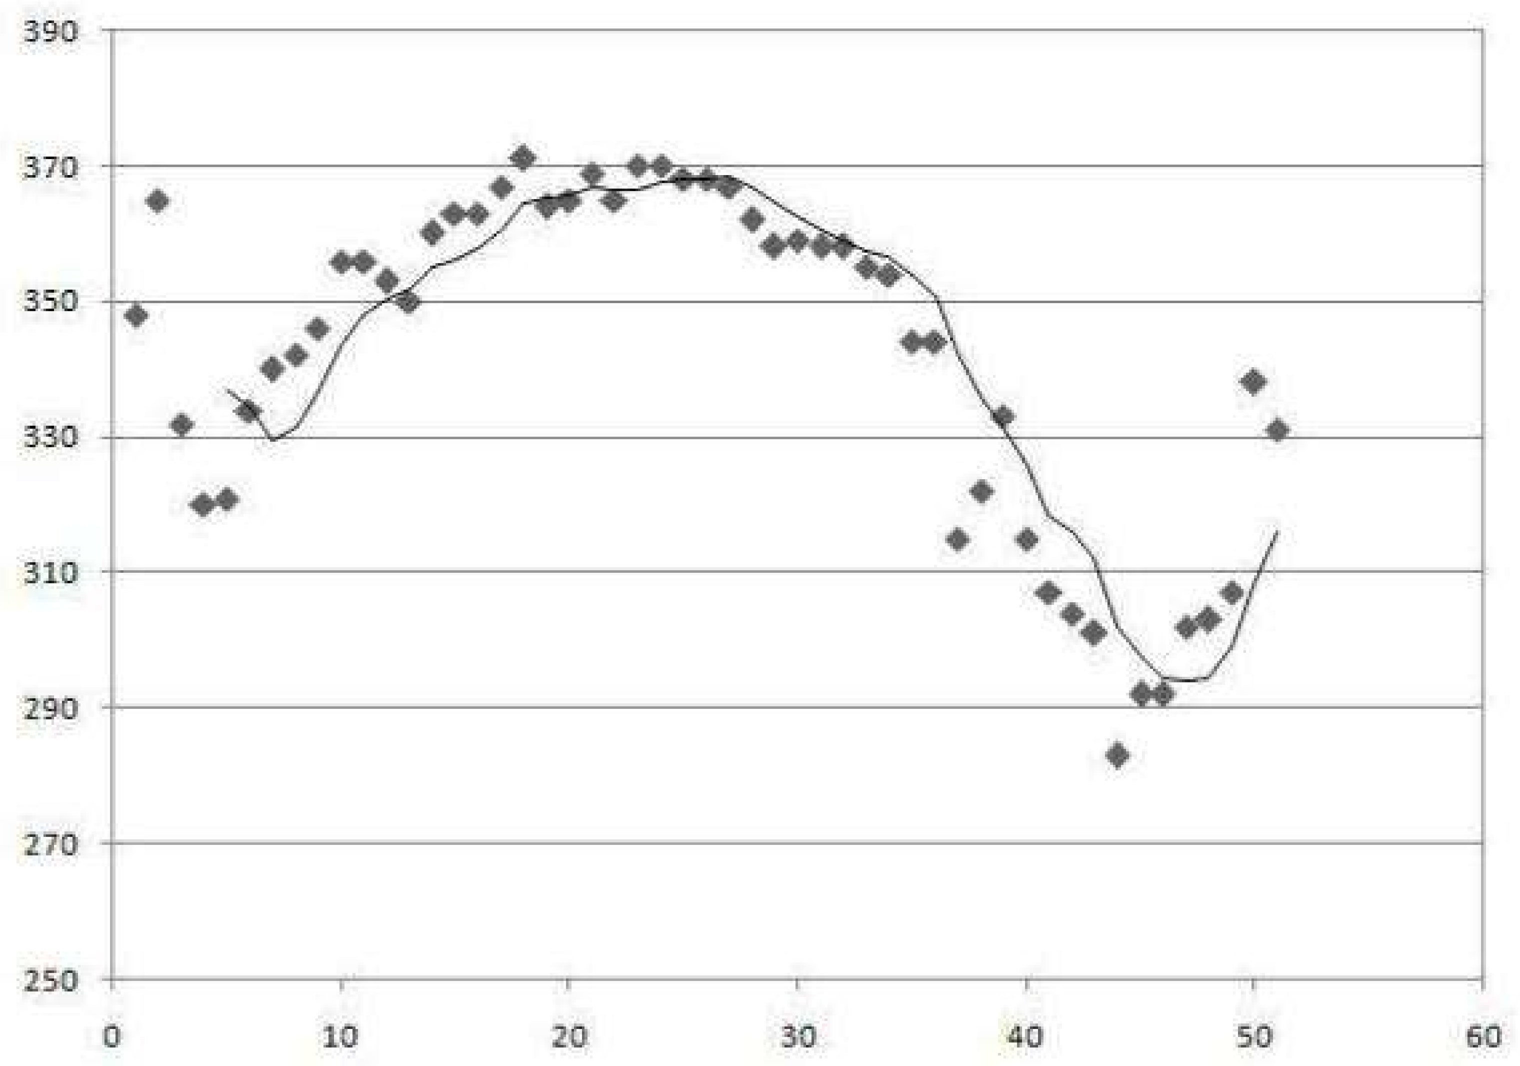
390
370
350
330
310
290
270
250
0
10
20
30
40
50
60

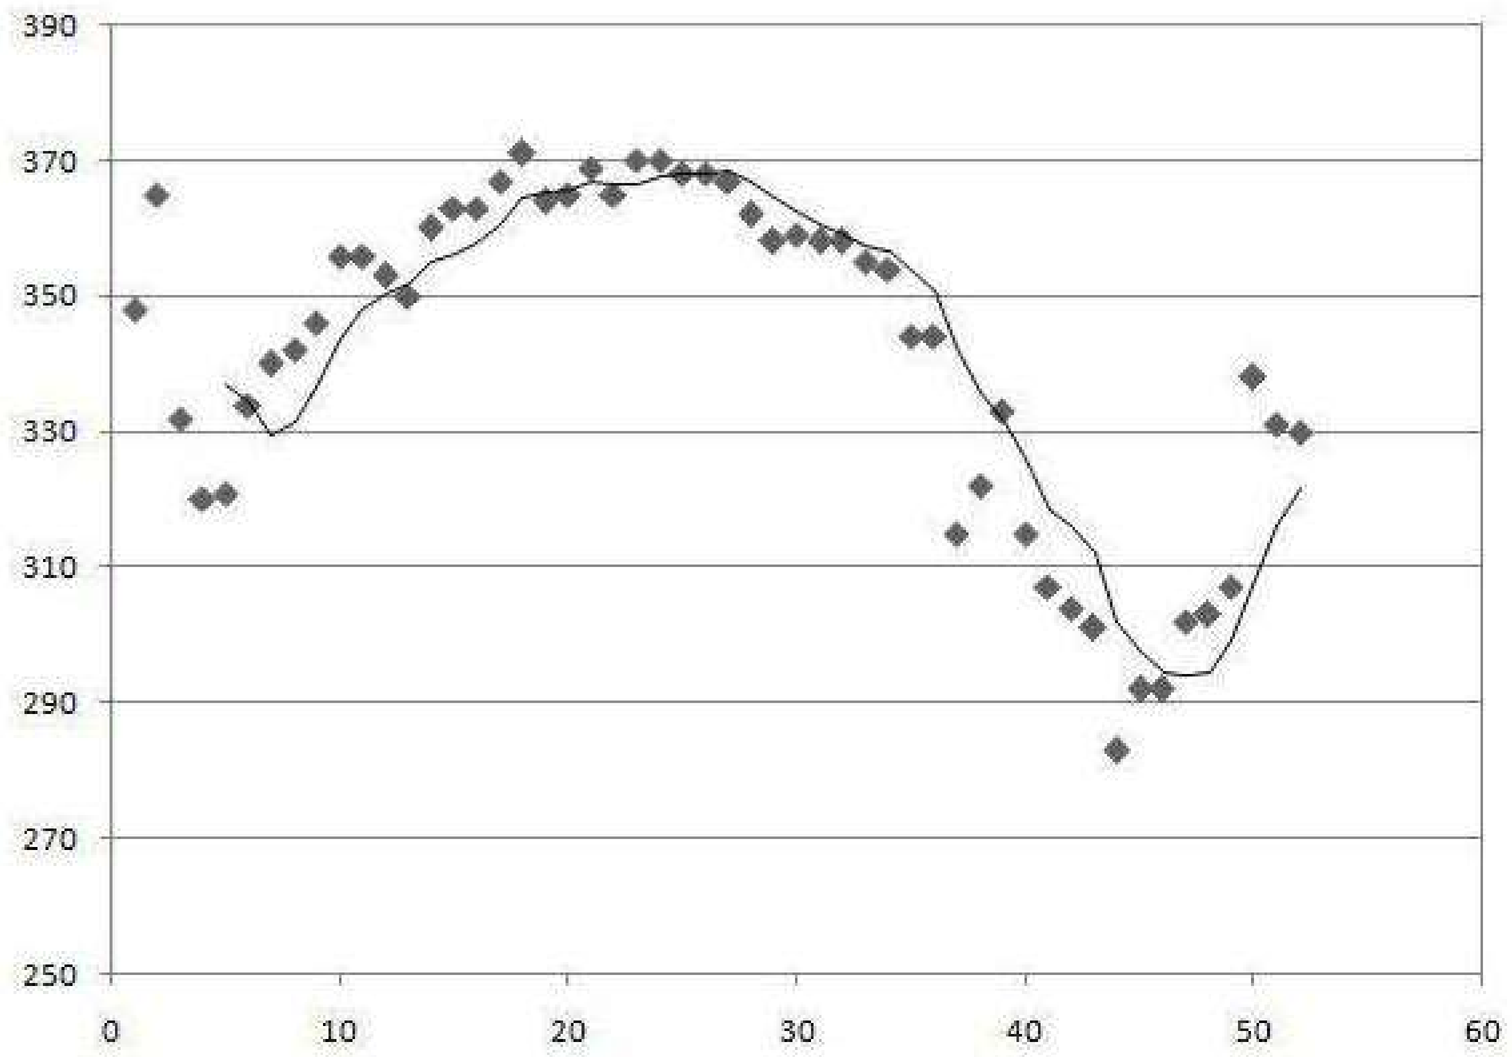
390
370
350
330
310
290
270
250
0
10
20
30
40
50
60

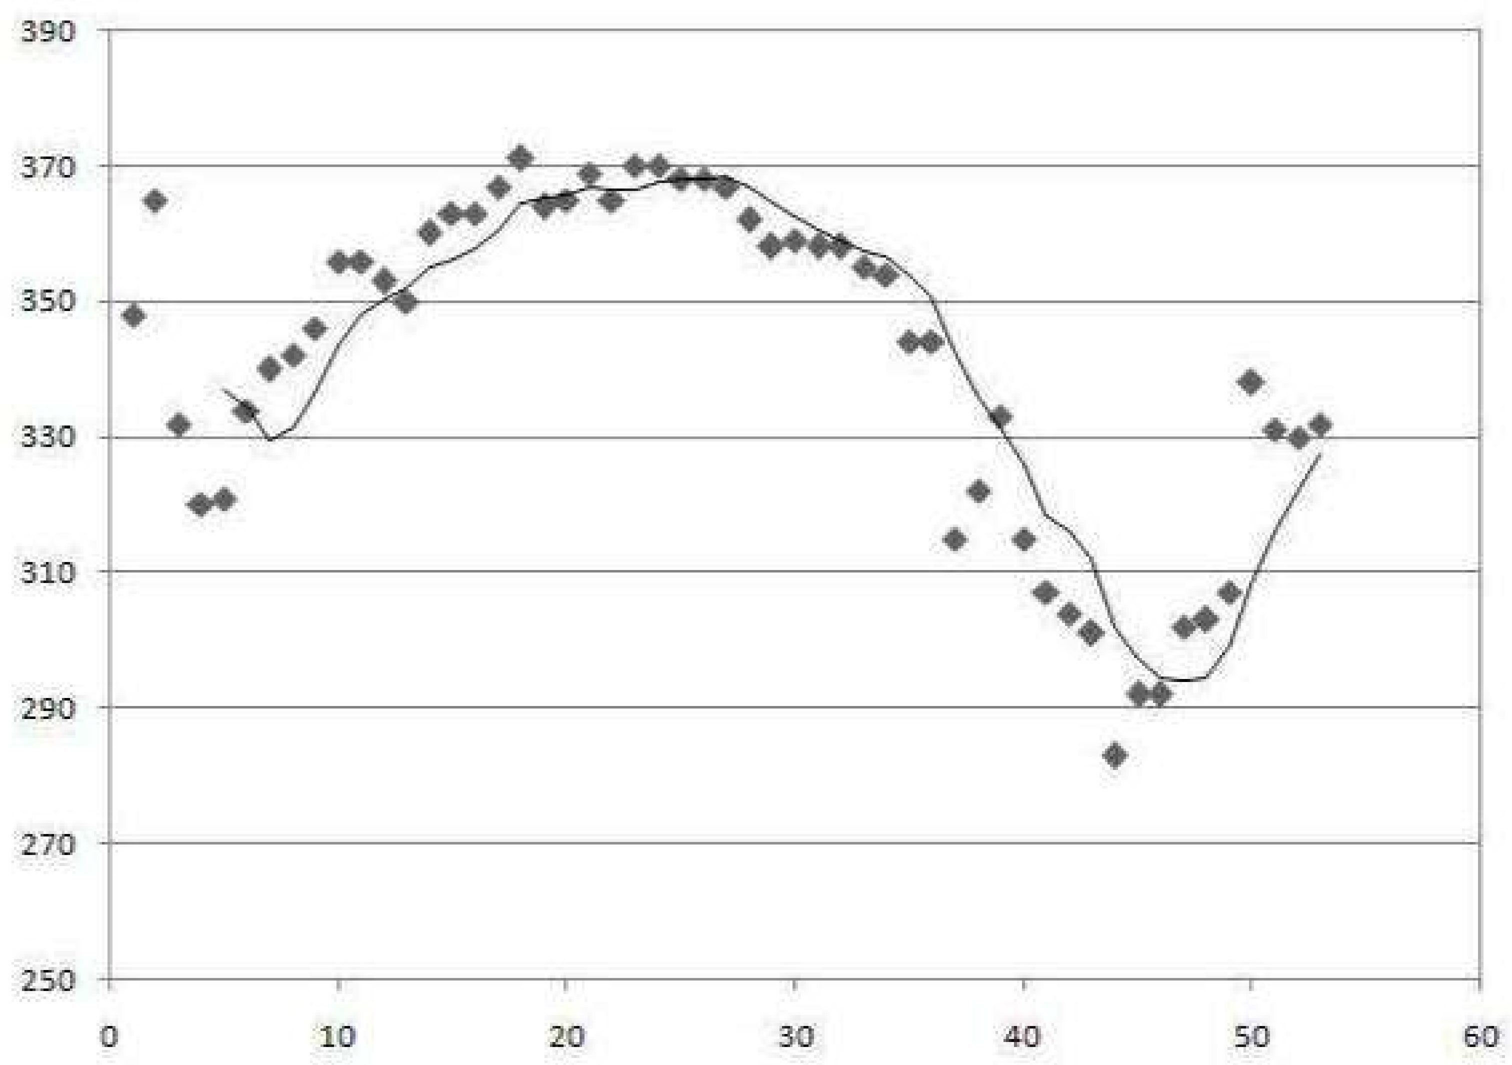


***What moving means***

There is a major difference between the way in which straight trend lines and quadratic trend lines are calculated compared with moving averages. In the case of the former, all of the data points were used in the calculation. Thus as new points arrive they have an influence on the whole of the plotted trend line.

On the other hand, moving averages use only the last n points in the calculation (where n is the period used for the average).

Starting from the beginning of the data, the first n points are used to calculate the first value. This sum is then divided by n to bring the resulting value to a similar order as the original data.

The calculation then shifts one point forwards and the first n points from this position are used to calculate the next value for the average. Thus the calculation moves forward through the data one point at a time, which is the reason for the term '*moving*'.

Because only the last n points are used to calculate the final value there will be no change in the position or shape of the average over the points prior to these final n points. In other words a new point only has influence on the n points of which it is the final point.

***Figure 7.12 – A plot where a 25-point moving average has been used.***

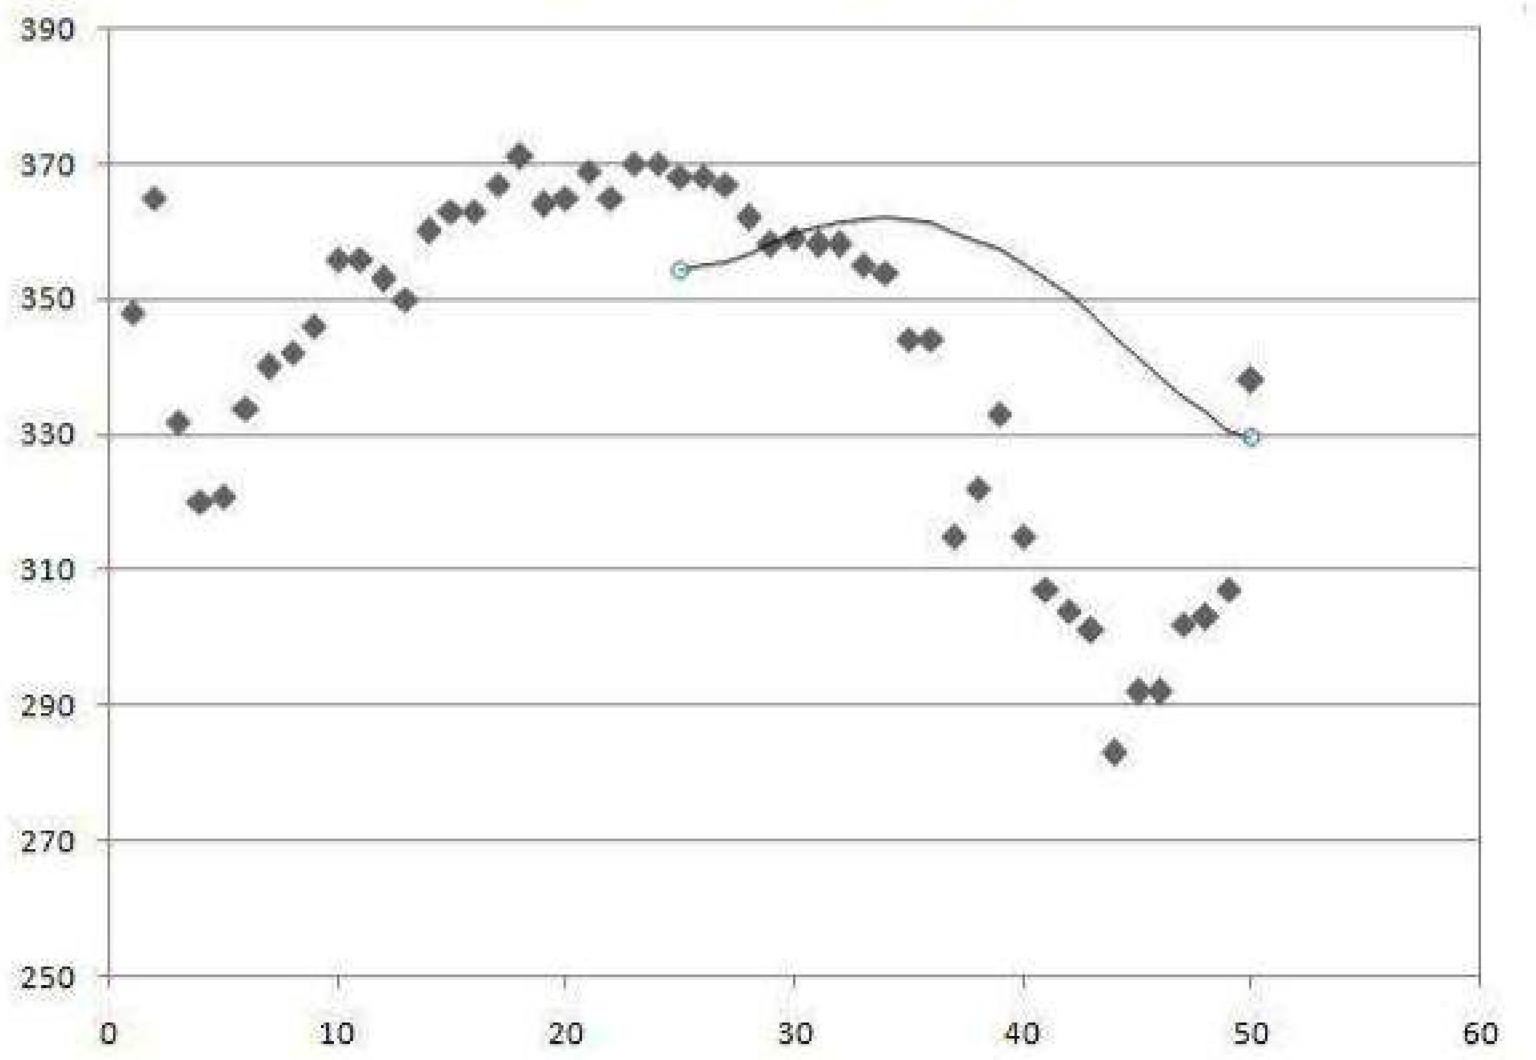


Naturally, the larger we make n, the further back the influence of the new point extends. However, on the other hand, its influence gets smaller because the sum of the n points is then divided by n. These issues will be discussed in more detail in the next chapter.

## Extrapolating a trend line

***Linear and polynomial***

Although moving averages are the key to prediction of future trends, as we shall see later in this book, there is a major difficulty in extrapolating them into the future. Linear trend lines and polynomial trend lines are characterised by an algebraic equation. Putting on one side the question of whether or not the trend line will be valid into the future, it is this equation which can be used to determine its future position. As we have discussed, although each new datum has the potential to change the trend, at any given time there is a mathematical basis for assuming that the trend will continue into the future and as such can be used as a predictor.

***Moving averages***

Unfortunately there is no underlying equation associated with a moving average line. Thus it cannot be projected accurately into the future by simple mathematics. We can use the moving average's current rate of curvature to project it into the future, but as we will see later this is subject to huge error at those points where a trend is changing direction.

We will show in later chapters that we can extrapolate a moving average trend line by three methods – using probable future movement derived from distributions, using channel analysis, and, the most fundamental of these three methods, using studies of stable cycles and their sums.

## TIMESCALE OF TRENDS

It is obvious that the moving average shown in Figure 7.8 is quite different from the moving average shown in Figure 7.12. Thus, the fact of using two different periods for the averages has resulted in the isolation of two different trends from the same data.

This concept, that there can be a number of trends co-existing within the same set of market data, is a vital one, since it will lead to an understanding of how stock prices move. Once this has been grasped, the process by which trends can be isolated and then predicted becomes clear.

As far as the stock market is concerned, for convenience we can divide trends into short term, medium term and long term.

According to *Wikipedia*, trends can be divided into primary, secondary and secular trends. However, it is clear that *Wikipedia* is applying these terms to an entire market. The secular trend is very long term and lasts from five to twenty-five years. The primary trends can last for a year or more. The secondary trend lasts for between a few weeks and a few months. According to *Wikipedia* the secondary trend acts in the opposite direction to the primary trend. I do not take this view, since the methods we will be using to isolate trends will make no distinction between rising trends and falling trends, but only between their timescales. It is perfectly possible to have two trends (of different timescales) both rising or both falling at the same time.

In this book we shall take long-term trends to be those which last for more than nine months, medium term between six weeks and nine months and short term one to six weeks. Expressed in days, therefore, long term is over 200 days, medium-term 30 to 200 days and short-term five to 30 days.

Since we will consider an uptrend to be the rising part of a cycle and a downtrend the falling part of a cycle, then obviously these categories of trend timescales are caused by cycles with wavelengths twice the time taken for the trend to rise or fall. Thus a short-term trend is caused by cycles of period 10 to 60 days, a medium-term trend by cycles of period 60 to 400 days and a long-term trend by cycles of period over 400 days. Trends which fall into these categories are shown in Figure 7.13.

***Figure 7.13 – How trends can be categorised into long-term (A), medium-term (B) and short-term (C) trends.***

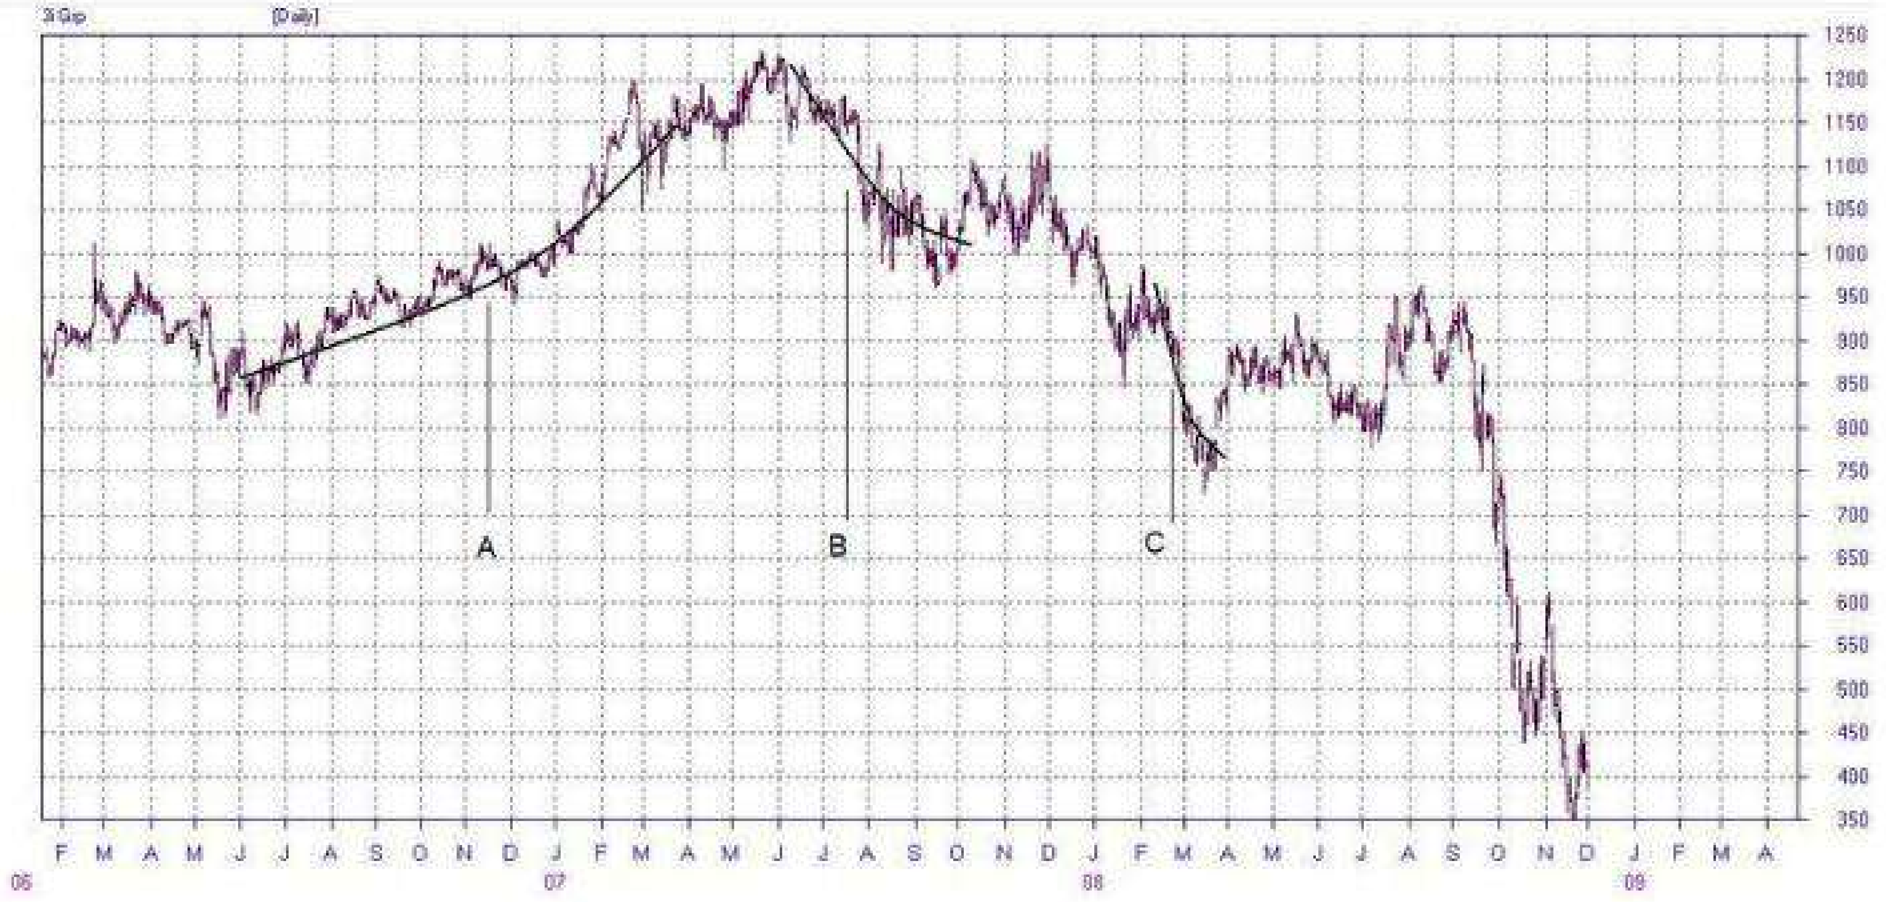


These are of course arbitrary distinctions, but they will serve us well when we come to decide how to isolate cycles or bands of cycles corresponding to these categories. Note the use of the term *bands*. Only rarely is a trend caused by just one cycle. There is usually a range or band of cycles which are responsible for a particular short-, medium- or long-term trend.

This will become obvious as we delve further into the relationship between cycles and trends in the following chapters.

# 8. Properties of Moving Averages

In this chapter we will be investigating the properties of moving averages and how they can be used to represent a trend in market data. We will also see the advantage in plotting centred averages rather than unlagged averages, which is the common way of presenting them on charts.

## Calculation of Moving Averages

Since, as mentioned in the previous chapter, all technical analysis packages can calculate averages, it might appear to be unnecessary to show how they are calculated. However an appreciation of the process will serve us very well in investigating the properties of moving averages and will also show their value when used to investigate cycles.

We will take as an example a set of market data of which the first ten values are designated by $V_1$, $V_2$, $V_3$, etc up to $V_{10}$. Suppose that we also need to calculate a five-point average. We can designate the values for these averages as $A_5$, $A_6$, $A_7$, etc up to $A_{10}$. Note that we do not start with an $A_1$, because the first average point is calculated by using the first five data points. Thus $A_5$ is the value we calculate once we have reached point $V_5$.

Once we have arrived at a total for the first five data points, we have to divide this total by five to arrive at a value for the average $A_5$.

Thus:

$A_5 = (V_1 + V_2 + V_3 + V_4 + V_5)/5$

$A_6 = (V_2 + V_3 + V_4 + V_5 + V_6)/5$

$A_7 = (V_3 + V_4 + V_5 + V_6 + V_7)/5$

and so on up to the final point $V_{10}$.

In general, therefore, to calculate an n-point average we would sum the first n points to arrive at a total and then divide this total by n. This gives us the first value for our average and using the nomenclature shown for the five-point average this first calculated point would be *An*.

They are arranged in this way to help to draw attention to an important aspect of the calculation. This is that, as we move the calculation onwards through the data, the last point of the five is dropped and the next point added in. Thus, whether the value for the average rises or falls when the next point is taken into consideration depends only upon the difference between the new point and the dropped point. If the new point has a higher value than the dropped point, the sum of the five points increases and so the new value of the average increases. Conversely, if the new point has a lower value than the dropped point, the new value of the average decreases.

This can be seen from Figure 8.1, where a nine-point average has been plotted over the data. The points from 248 to 256 have been added and divided by nine to produce the average value at point 256 (embedded in the smooth average line). The next calculation will drop the point at 248 (which has a high value) and add in the point at 257 (which has a low value). Thus the average calculated at point 257 will be lower than the value at point 256.

***Figure 8.1 – This shows why the next calculated point of the nine-point average will be lower than the previous value. The new point has a lower value than the dropped point.***

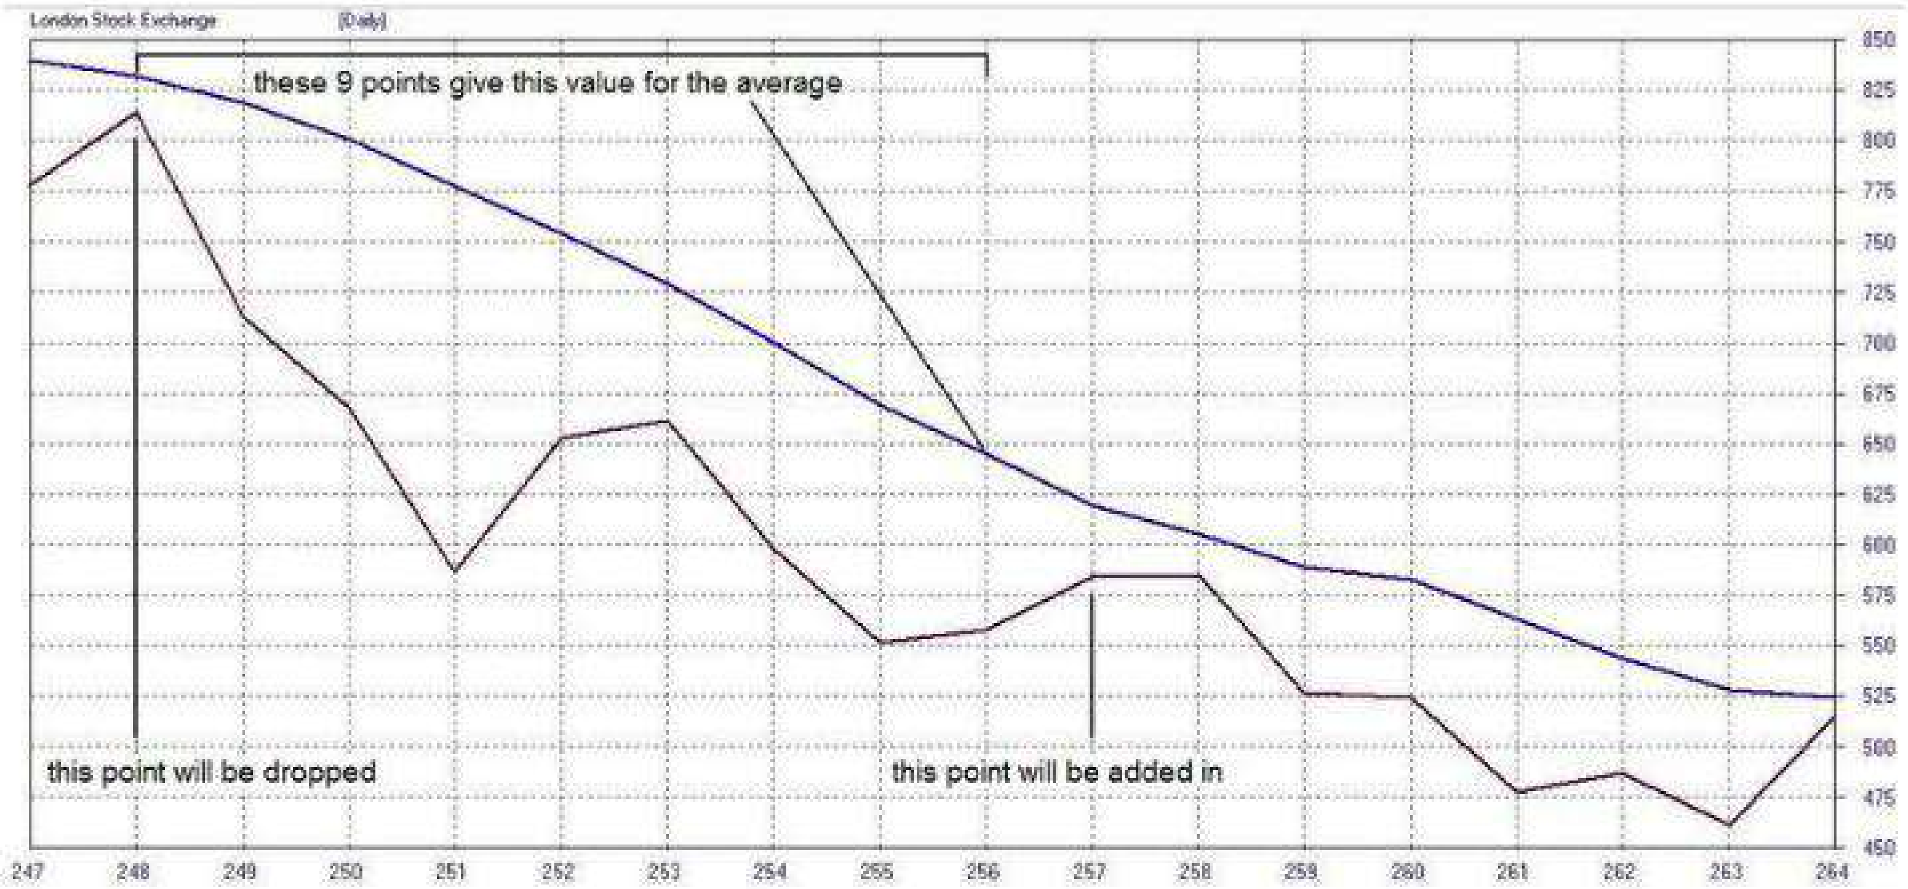


Of course, the average plotted in Figure 8.1 is plotted with no lag. Thus the average calculated using points 248 to 256 is plotted at the same position in time as point 256 and the next calculation of the average using point 257 is plotted at the same point in time as point 257.

From this the reader can see the fallacy in the way in which moving averages are used by many technical analysts. They use rules which are based on the general view that when the price rises above the average it is time to buy and when it falls below the average it is time to sell. However, as we have seen from the discussion of the dropped point, it is not whether the value of the data is higher than the current value of the average that causes the average to rise, but whether the value of the data is higher than the value of the dropped point.

## WHY DOES AN AVERAGE SMOOTH THE DATA?

The answer is because of the way it interacts with cyclic data. A moving average will totally or partially remove cycles from the data. We can divide this discussion into three parts:

1. there is the effect on cycles of the same wavelength as the period used for the average,
2. there is the effect on cycles of wavelength less than n, and
3. there is the effect on wavelengths greater than n.

### *Cycles with wavelength equal to period of the average*

What happens here is that an average of period n will remove cycles with wavelength n. The reason for this is easy to explain in terms of the dropped point. This is shown in Figure 8.2.

***Figure 8.2 – A cycle of wavelength 25. If the moving average has a period of 25, then each dropped point has the same value as the new point, as shown by the arrowed line. Thus the average remains at its first calculated value.***

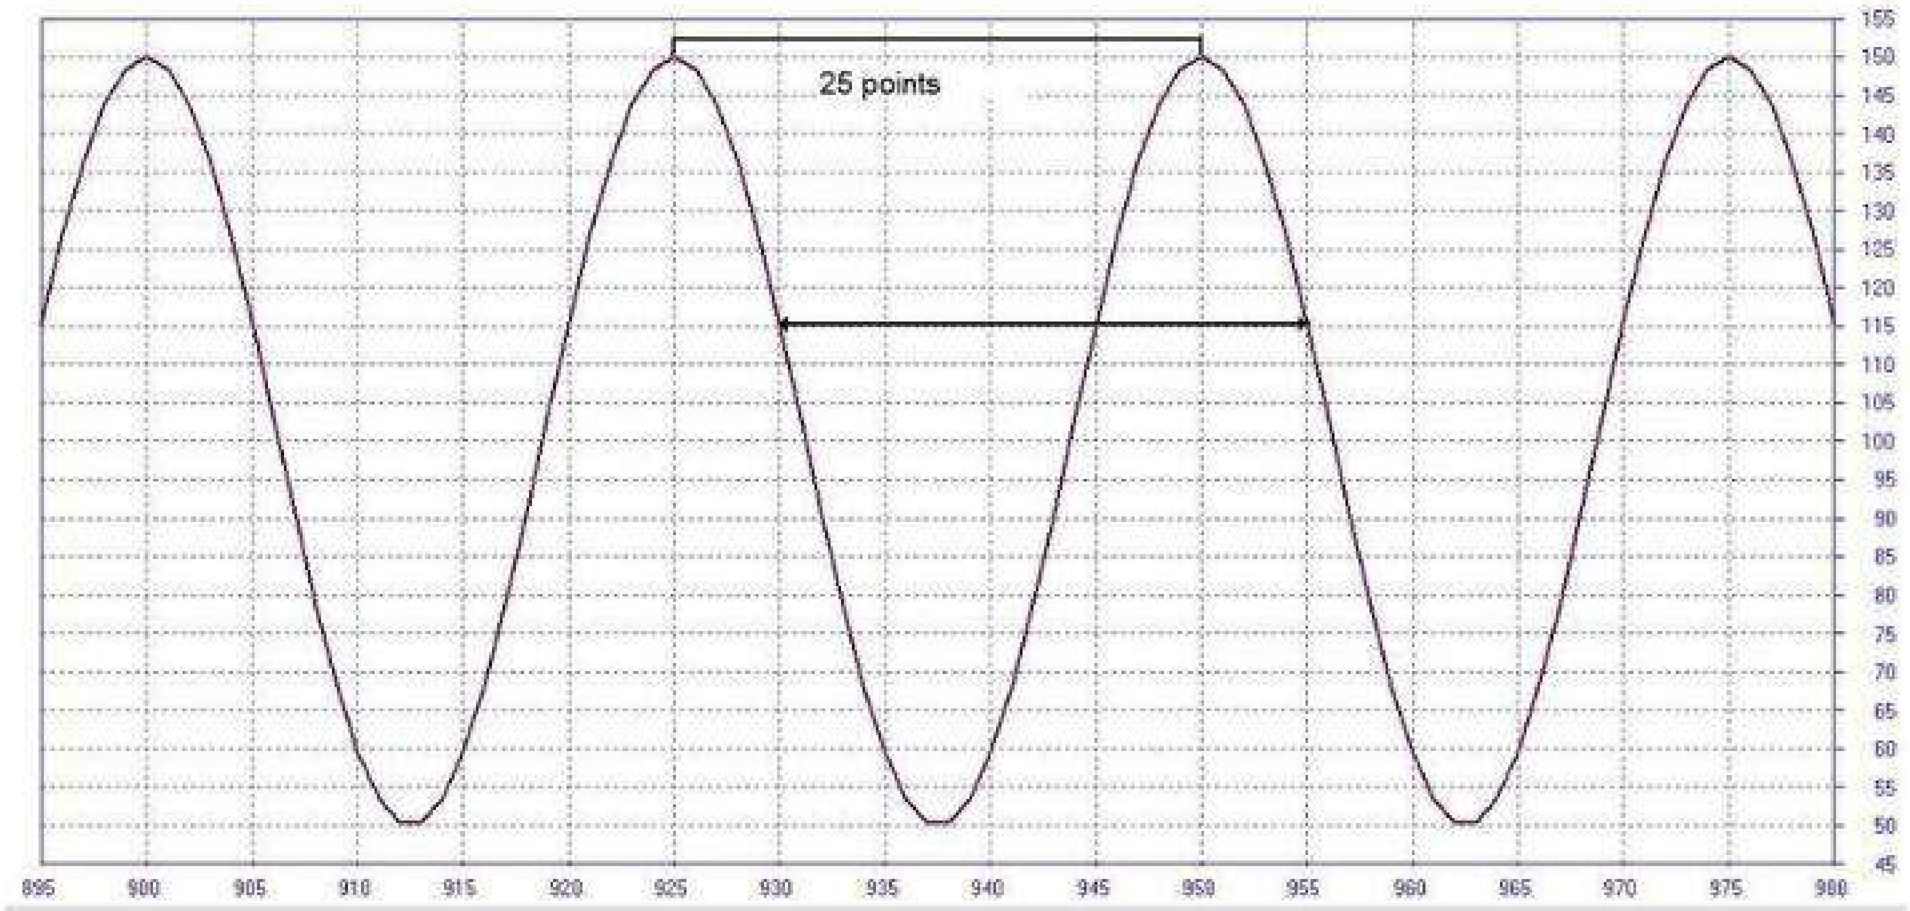


All points on the sine wave are identical in value to those points which are separated by one wavelength of 25. Because of this, every point which is dropped is replaced by a point of equal value. Thus the average remains at its first calculated value, producing a straight line, as shown in Figure 8.3.

***Figure 8.3 – The result of applying a 25-point moving average to a cycle of wavelength 25 days. The output is a straight line for the reason shown in Figure 8.2.***

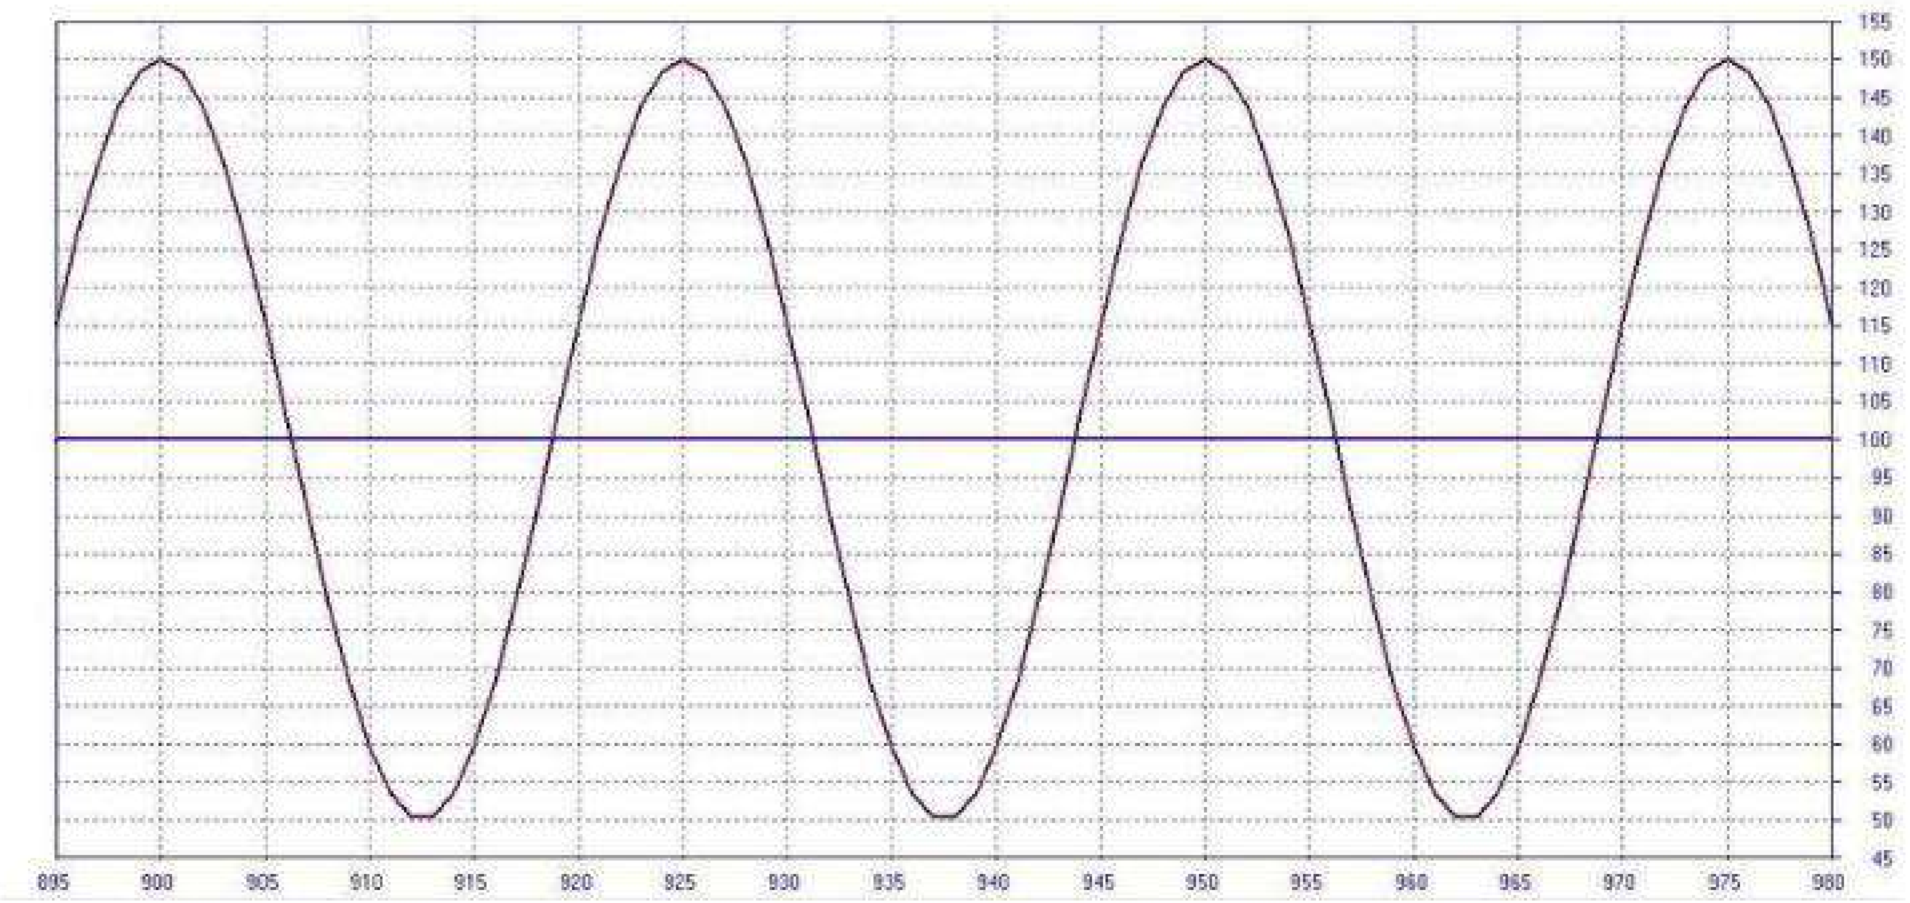

A moment's thought will lead you to the conclusion that cycles with wavelengths which are a multiple of the period used for the average will also be removed, since again every point that is dropped is replaced by a point of equal value.

***Wavelengths higher than the average period***

While cycles with wavelength equal to the period of the average are removed completely, cycles with wavelength higher than the period of the average are partially removed, i.e. attenuated. This attenuation decreases as the wavelength increases.

This can be seen in Figure 8.4. Here a cycle of wavelength 41 days has been produced. The output from the averages when applied to this cycle decreases in amplitude as we increase the period of the average, i.e. as the wavelength becomes higher relative to the period.

***Figure 8.4 – The effect of applying nine-, 15- and 21-point averages to a cycle of wavelength 41***

***days. The closer the period of the average to the wavelength of the cycle, the greater is the attenuation.***

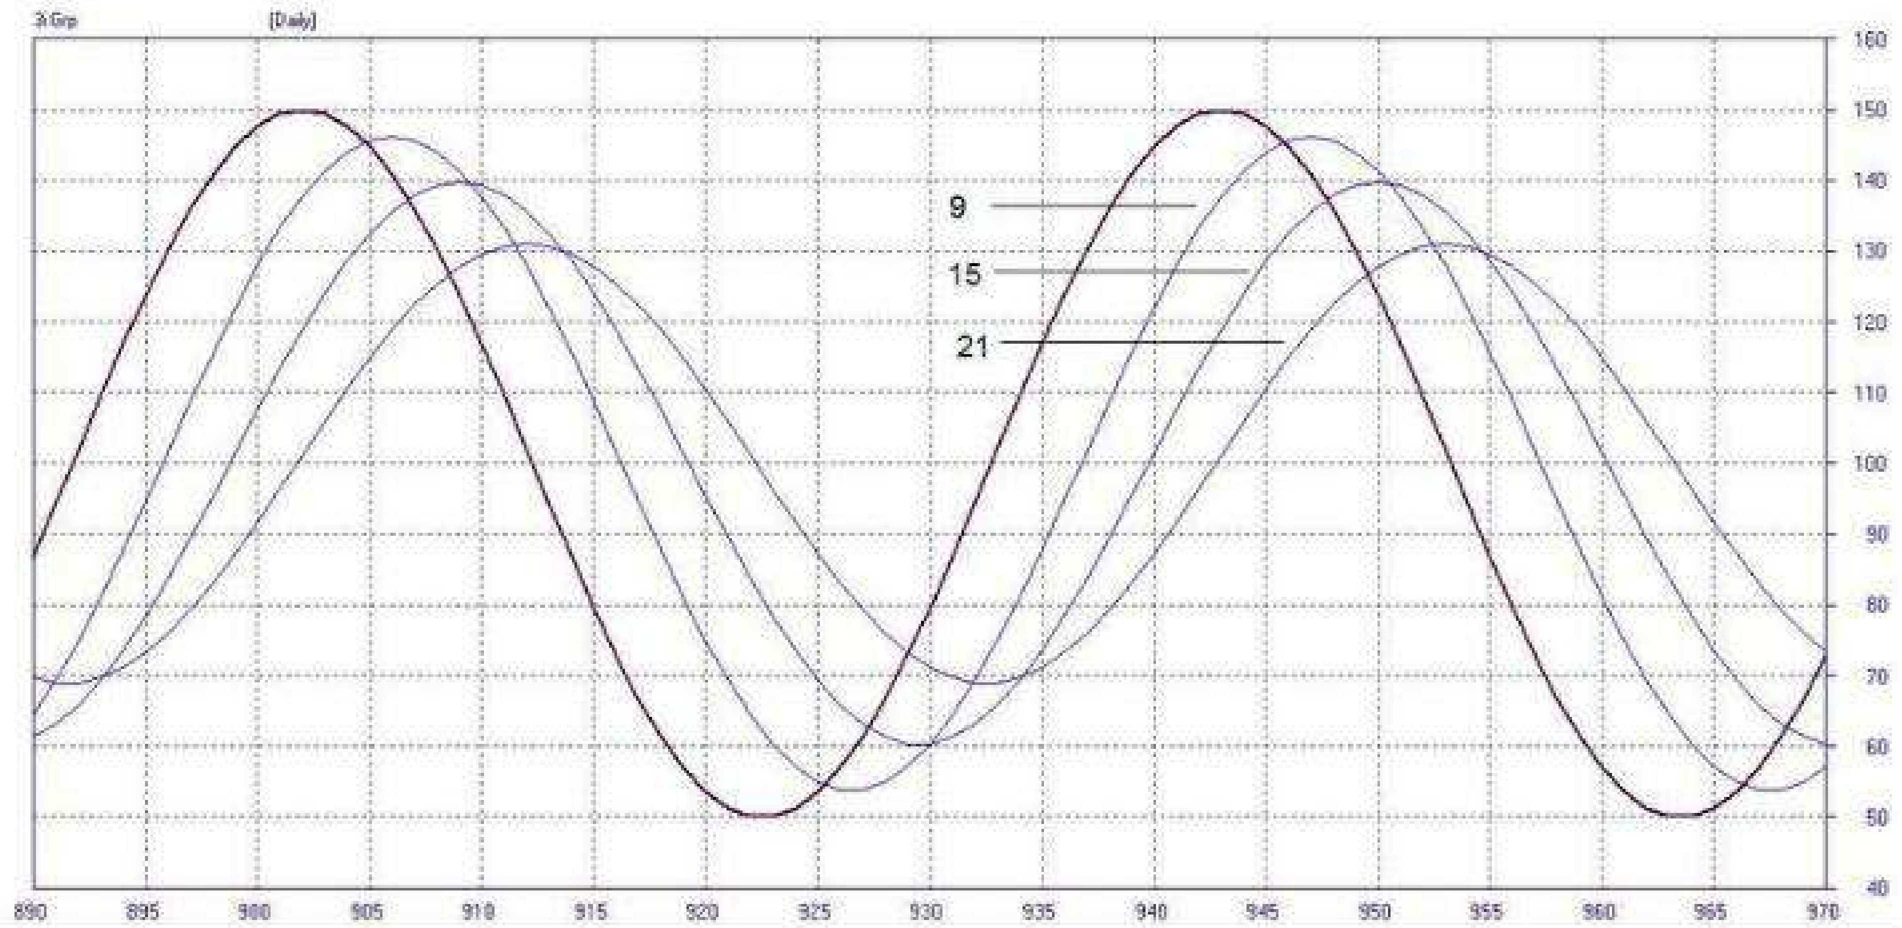


***Wavelengths lower than the average period***

While cycles with wavelength equal to the period of the average are removed completely, cycles with wavelength lower than the period of the average are partially removed, i.e. attenuated. This attenuation increases as the wavelength decreases.

***Figure 8.5 – The effect of applying 31-, 39- and 41-point moving averages to a cycle of wavelength 41 days. The original cycle is the one with the maximum amplitude.***

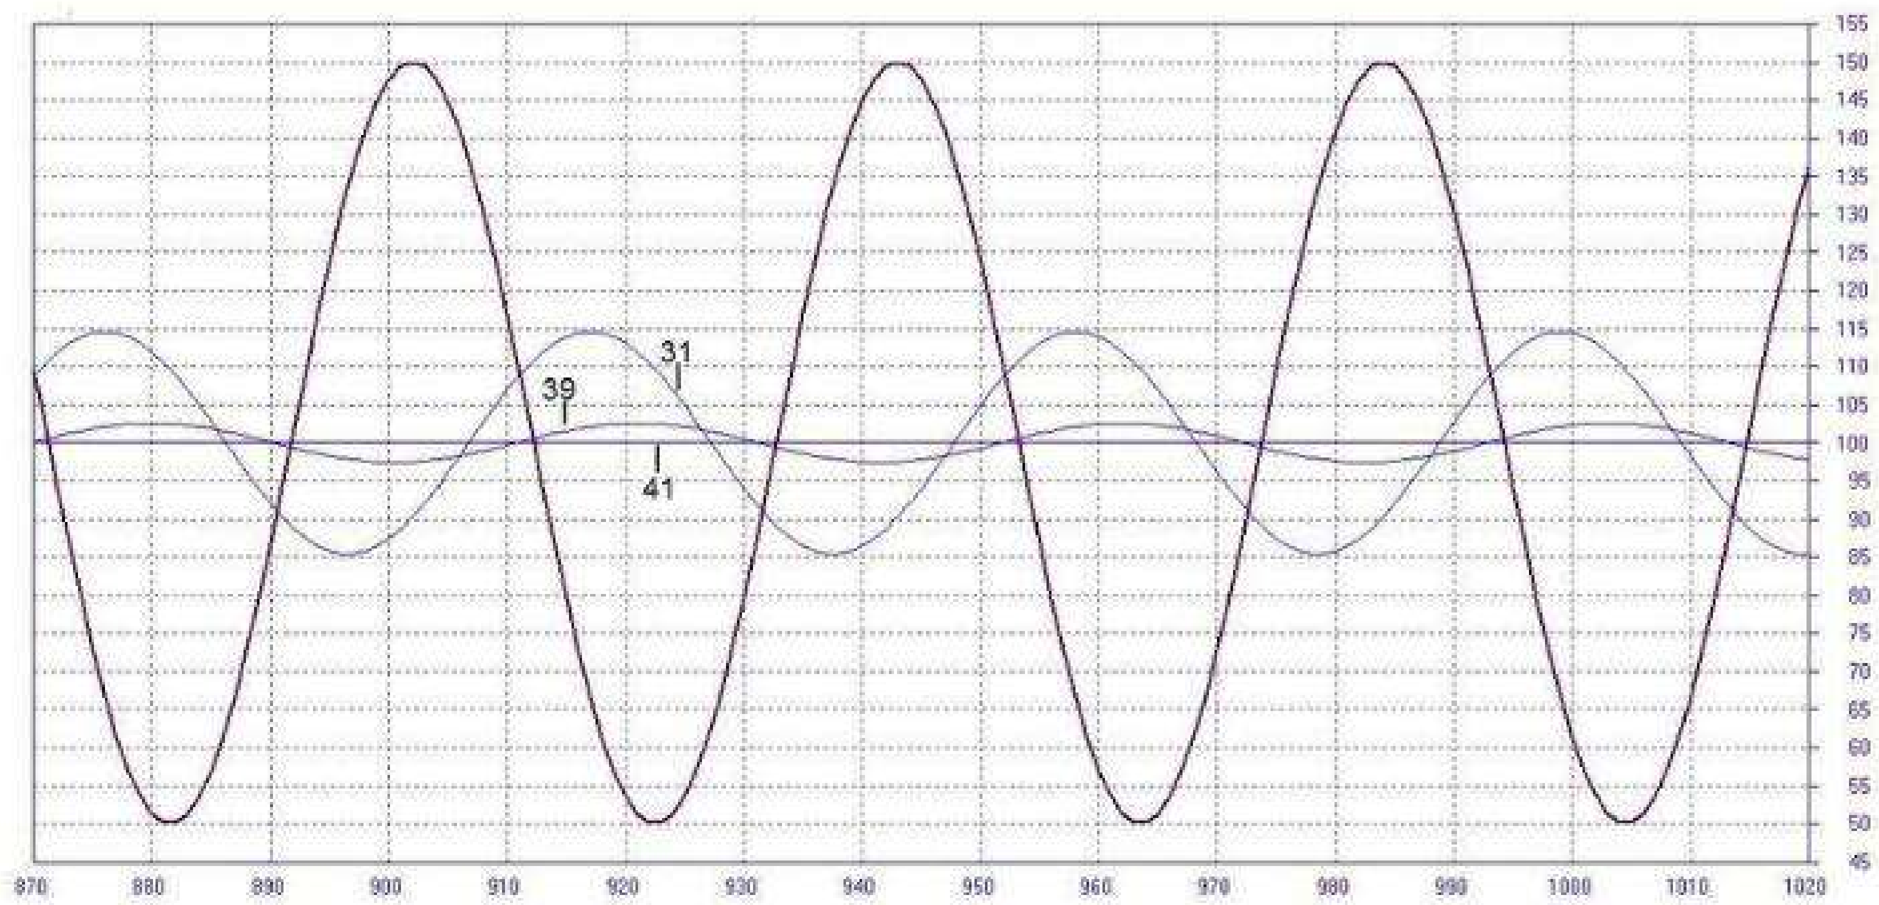


## CENTRED AVERAGES

You will note in Figures 8.4 and 8.5 that the peaks and troughs in the output from applying moving averages to cycles do not align with the peaks and troughs in the original data. This is because the output has been plotted with no lag and therefore this means that the last calculated value for the average is plotted at the same point in time as the last data point. This is of course the traditional way used by technical analysts to plot averages.

We see also that the peaks and troughs in the output from each average do not line up with the individual peaks and troughs in the original data. There are many disadvantages in plotting averages in this way. The analysis of market data for cycles cannot be carried out and neither can any probabilities be derived.

To plot a centred average, it is necessary to understand that the latest value for the average must be plotted back in time from the latest data point. This offset is such as to plot the average value at the same position in time as that of the mid-point of the data points used to calculate that value of the average. In order that any point of the average coincides with a position in time of a data point, it is necessary to use an odd number for the periods. This will become clear if we now rewrite the formula used earlier.

Now we change this so that $A5$ becomes $A3$, the centre point of the five data points used to calculate it. In other words, when we come to plot this average, the first calculated point, now designated $A3$, is to be plotted at the same position as data point three. It is important to understand that this does not affect the value of any of the points, simply how they are to be plotted.

Thus:

$A3 = (V1 + V2 + V3 + V4 + V5)/5$

$A4 = (V2 + V3 + V4 + V5 + V6)/5$

$A5 = (V3 + V4 + V5 + V6 + V7)/5$

and so on up to the final point $V10$.

When plotted in this way, the average of course terminates at a point in the past. The last plotted point will then be plotted at a position half of the span of the average back in time. Thus, we do not know how the average is likely to move over this gap between the position of the last calculated average point and the latest data point. This will not be known until we have moved that number of points which is equal to half a span of the average into the future. Then we will have enough future data points to be able to fill in the values across the gap.

Since in the next chapter we will be discussing using centred averages as proxies for trends, we will have to find ways of estimating how the average might move across this gap, what is its likely position at the present time, and what it is likely to do in the near future. It is this difficulty which will be explored in the next chapter.

When the averages which were plotted in Figure 8.4 are now plotted as centred averages, their peaks and troughs will now line up perfectly in time. The picture is now much more meaningful, as shown in Figure 8.6.

***Figure 8.6 – When the nine-, 15- and 21-point centred averages are applied to the cycles of wavelength 41 days the peaks and troughs now line up perfectly.***

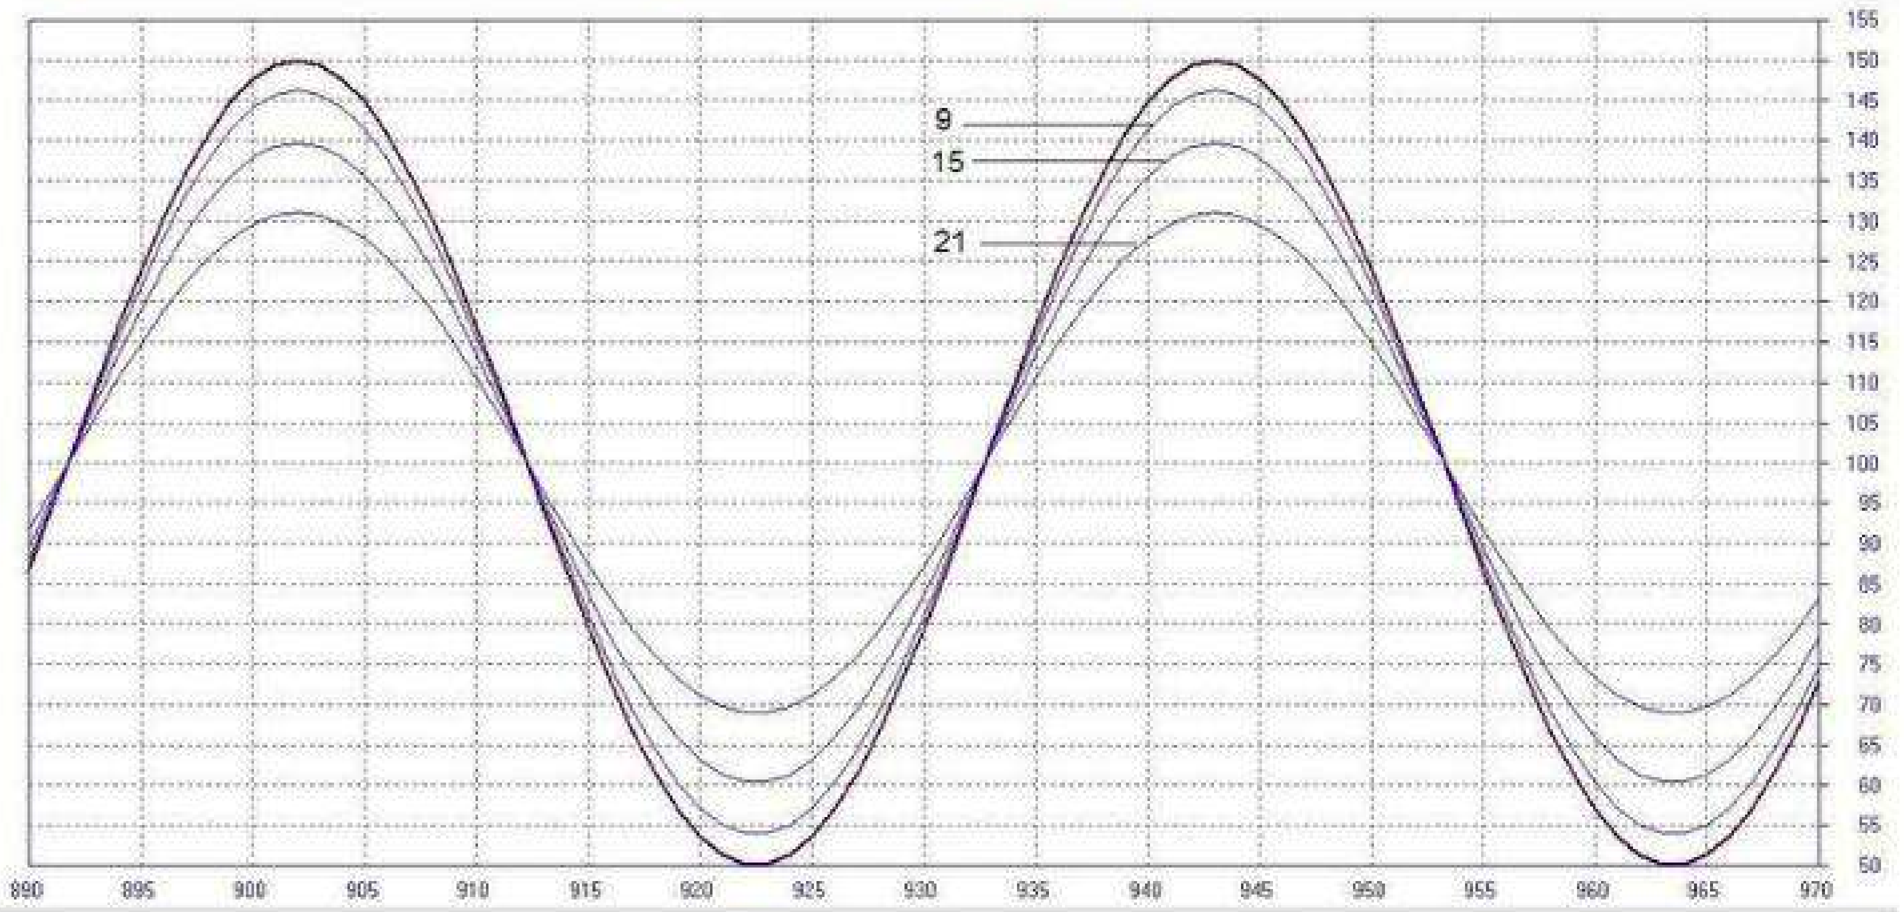


The effect of increased attenuation by averages as their period approaches the wavelength of the cycle is now much more obvious. The greater the wavelength of the cycle relative to the period of the average, the less is the attenuation. For wavelengths very much greater than the period used, the attenuation is almost negligible. This state of affairs is reached when the ratio of the wavelength to the period of the average exceeds about three. This is a very important point that we need to take into account later when isolating cycles by means of centred moving averages.

When the averages which were plotted in Figure 8.5 are now plotted as centred averages, once again the peaks and troughs line up perfectly. This is shown in Figure 8.7. Notice that, as in Figure 8.5, the effect of applying the 41-point average to the 41-day cycle is to completely remove it, leading to a straight line.

***Figure 8.7 – When the 31-, 39- and 41-point centred averages are applied to a cycle of wavelength 41 days the peaks and troughs of the first two line up perfectly. The 41-point average has totally removed the cycle in its output, leading to a straight line. The closer the period of the average to the wavelength of the cycle, the greater is the attenuation.***

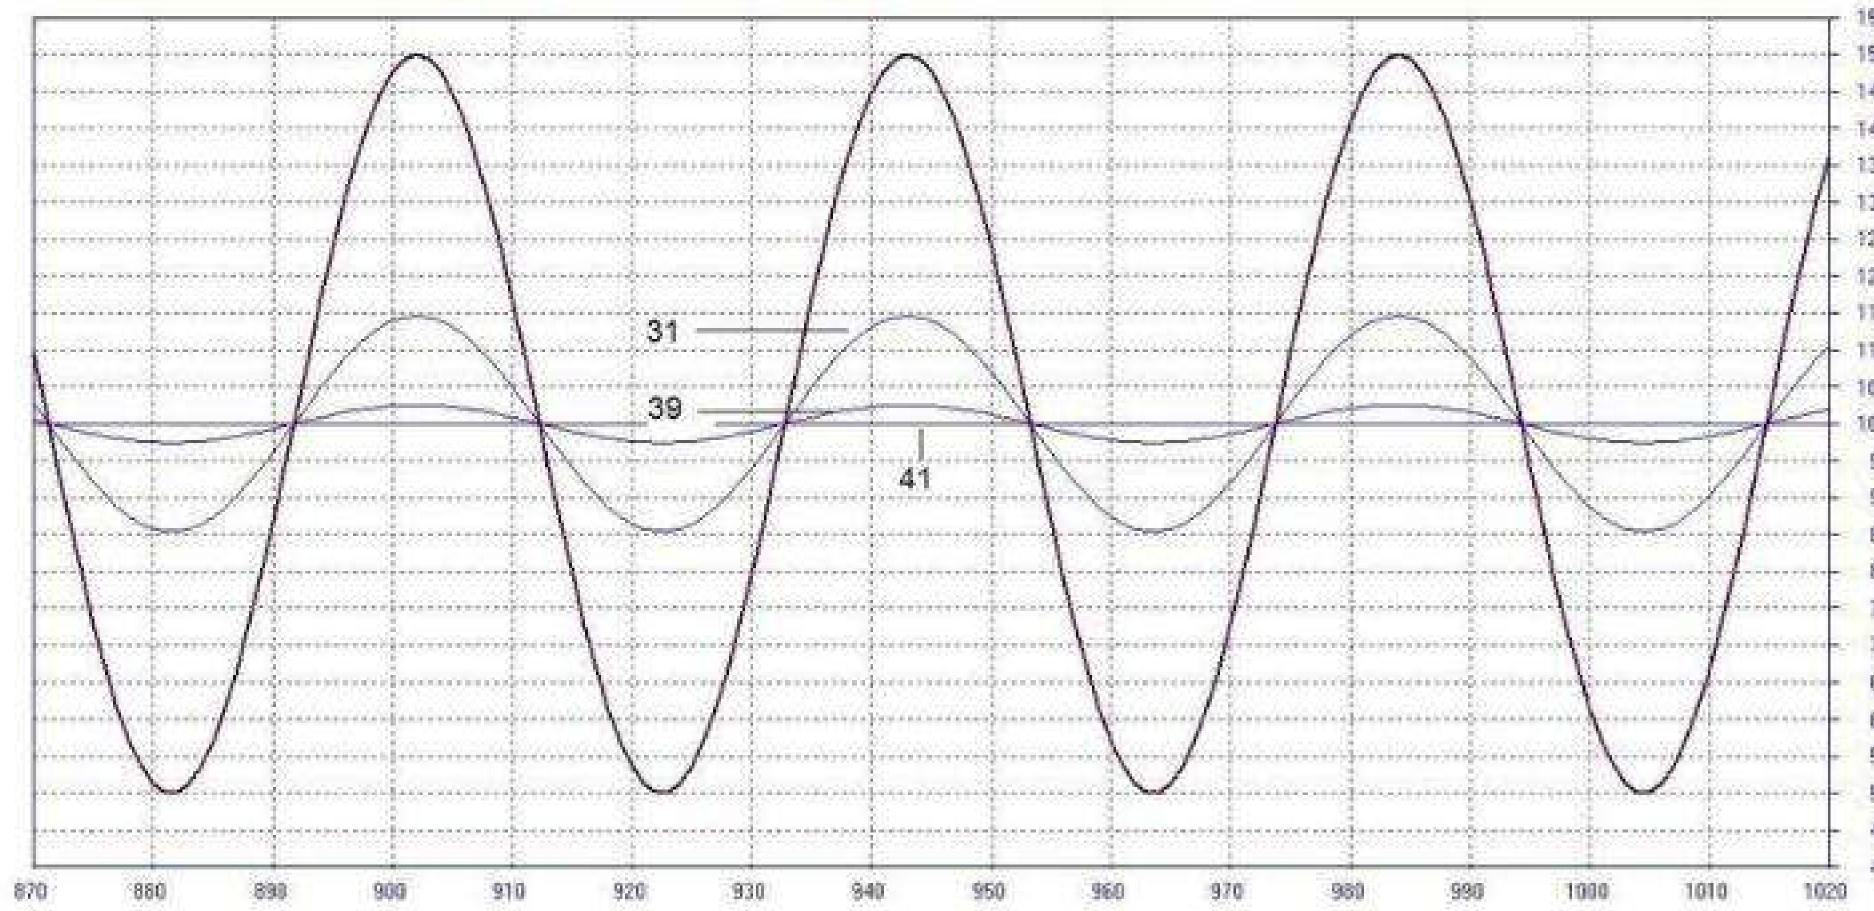


***Applying averages to a mixture of cycles***

Now we have seen the effect of applying moving averages of periods less than, equal to and greater than the wavelength of the cycle, it is important to examine how we can apply this knowledge to simple mixtures of cycles. This will enable us to arrive at the best ways of separating such cycles, so that we will be able to understand what happens when we use averages in an examination of real market data.

Three cycles of wavelength 41, 101 and 202 days with amplitudes of 25, 50 and 100 units respectively were combined with a horizontal straight line of 100 units. The plot of this complex waveform is shown in Figure 8.8.

***Figure 8.8 – The plot of a sum of three cycles of wavelengths 41, 101 and 201 days and amplitudes 25, 50 and 100 points, with a straight line of value 100 points also added.***

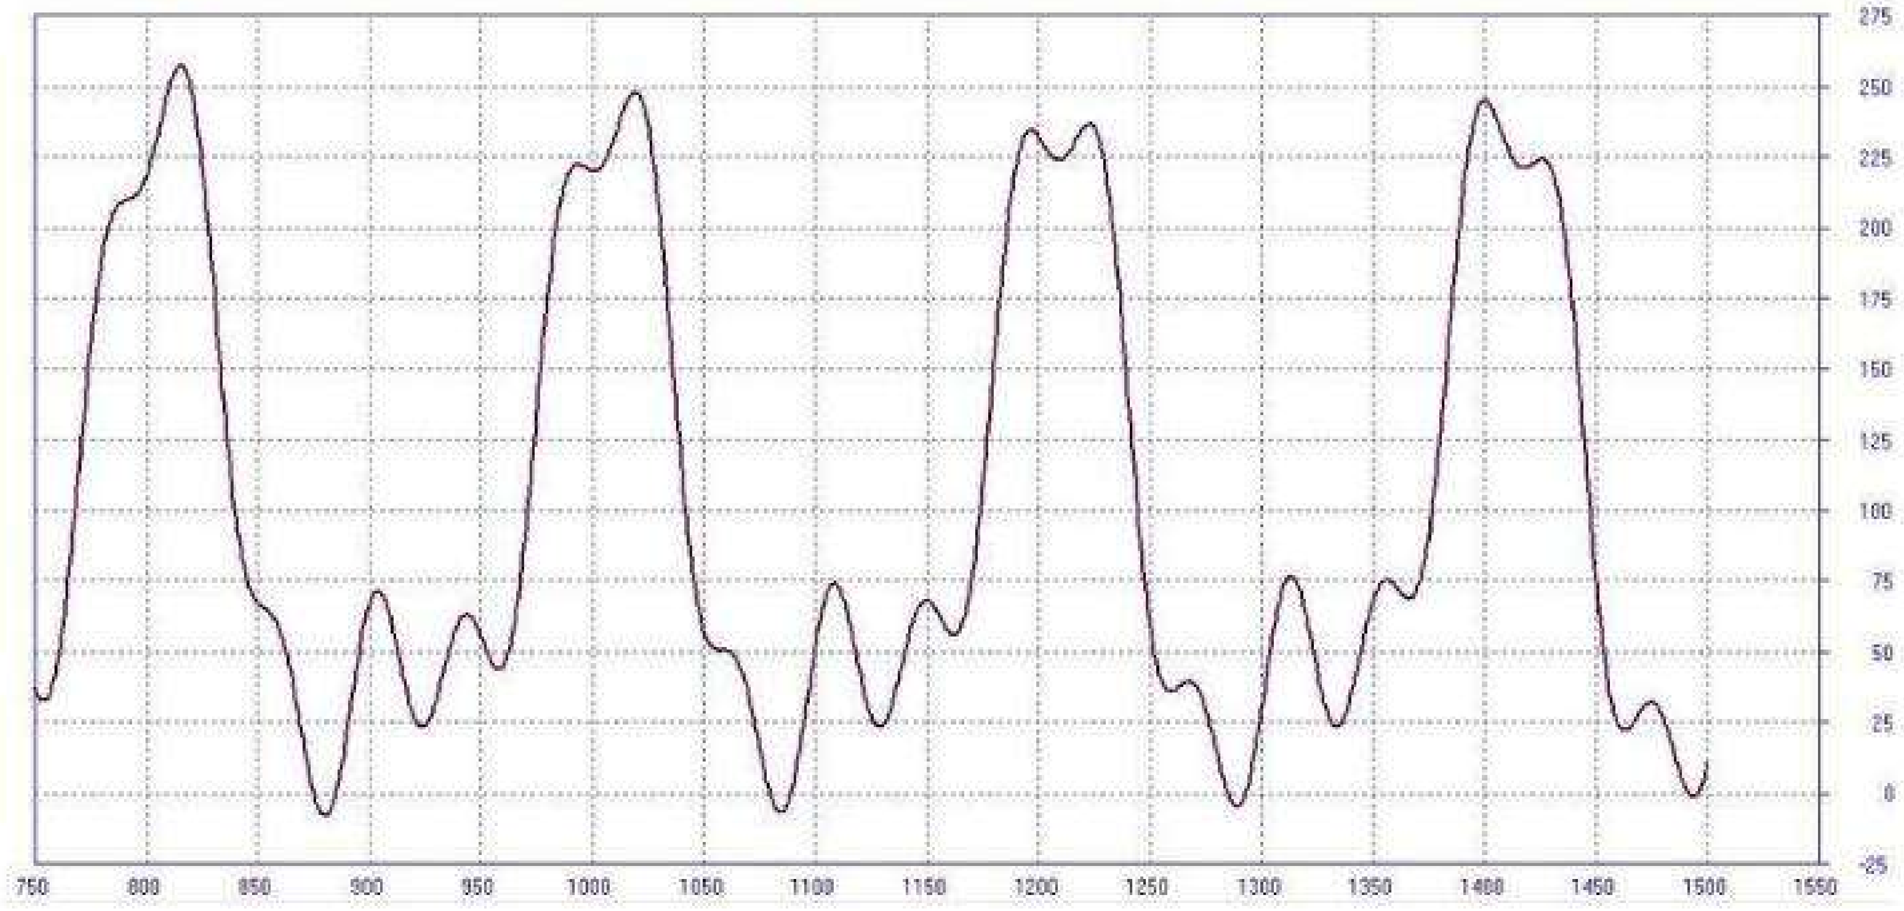

***151-day average***

If we were unaware of which cycles were present then we would have to test a few averages of different periods to form an idea of which cycles were present. A good starting point is a period of around half of the wavelength range of cycles in which we would be interested if examining the market. In this case we have applied a centred average of period 151 days. The result is shown in Figure 8.9

***Figure 8.9 – A moving average of period 151 days has been applied to the data from Figure 8.8.***

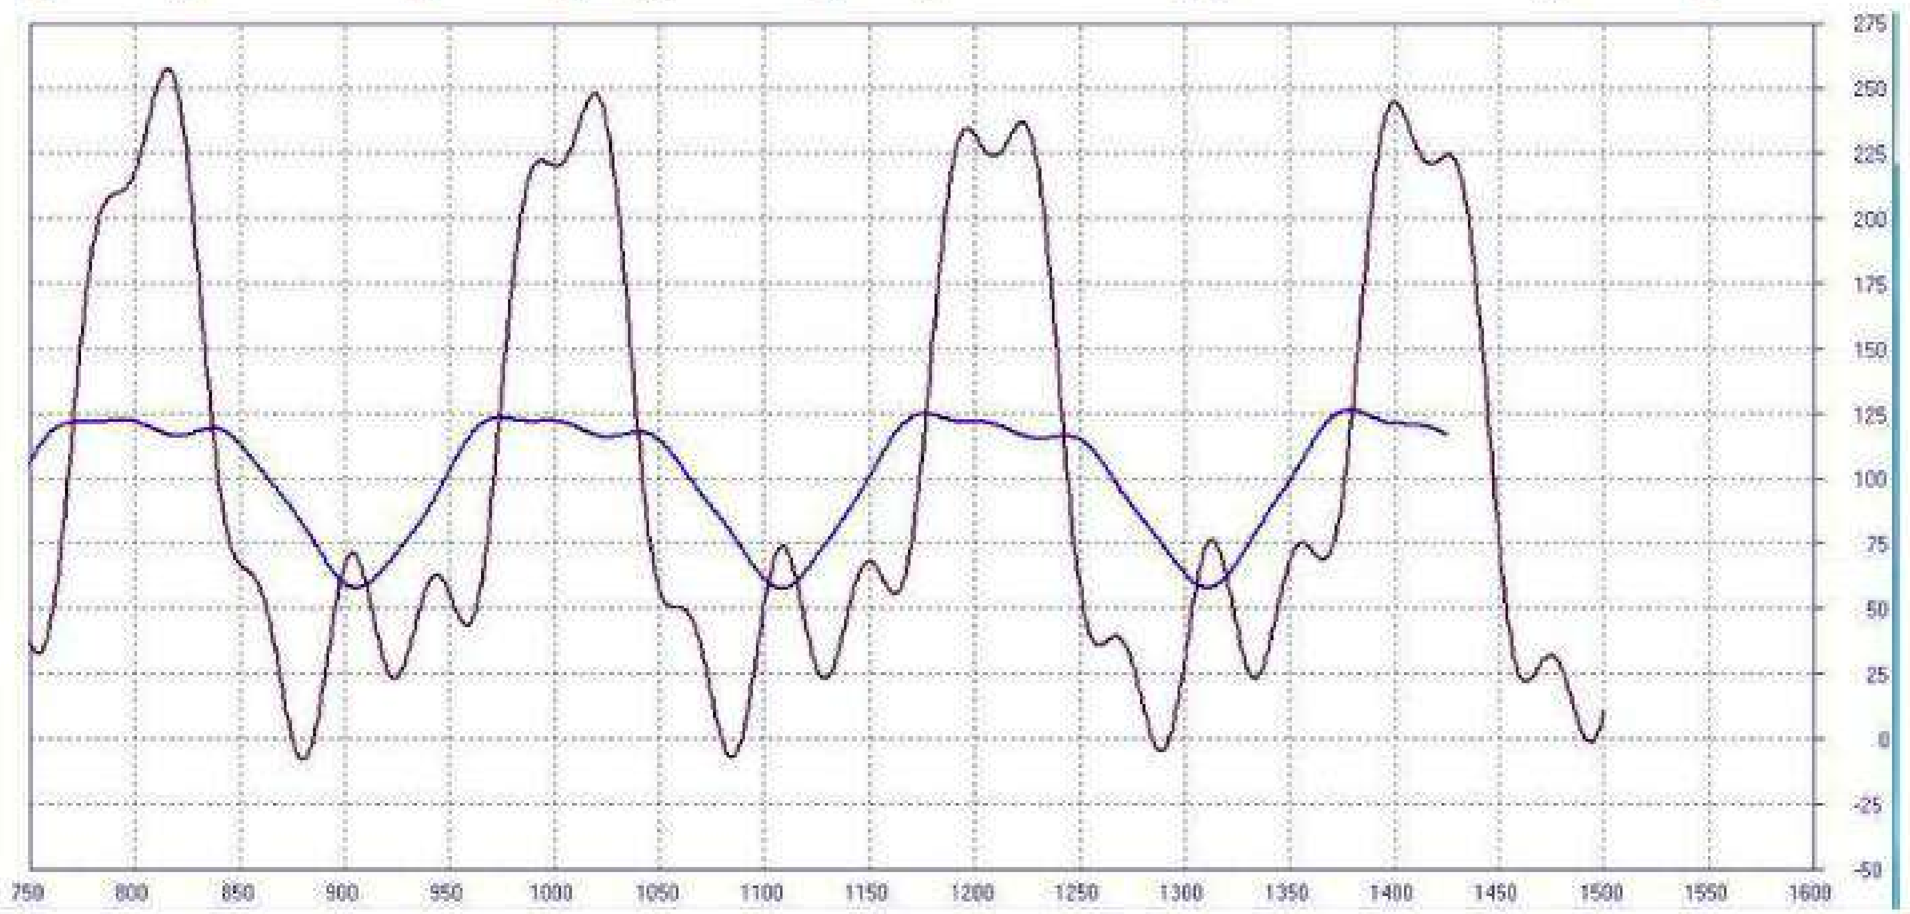

At this stage a great deal of information can be obtained. Firstly this average contains only a small amount of perturbation, so that we can take the distances between the clean troughs as being significant. These distances are 202 and 203. Thus we deduce that there is present a cycle of wavelength of around 202 or 203 days. Note that in a complex mixture these distances may be slightly distorted so that the wavelength deduced from such troughs or indeed peaks may be subject to a small error.

The other piece of information is that a line drawn through the centre of this average corresponds to a value of about 90 to 95, suggesting this is a constant amount which has been added to the cycles.

***201-day average***

Since we have discovered a wavelength of around 200, we can now apply a centred average of period 201. The result of doing this is shown in Figure 8.10. This is very informative since this has removed all the various cycles, leaving an almost straight line at a level of about 100. We can now say with confidence that if all the cycles have been removed, the resulting straight line represents the original line of value 100 points which was added into the mixture. We can also say that there are no cycles with wavelengths appreciably more than 201, since these would still be apparent in the output from the 201-point average.

***Figure 8.10 – A centred average of period 201 days has been applied to the complex waveform. The result is an almost straight line at a value of 100 points.***

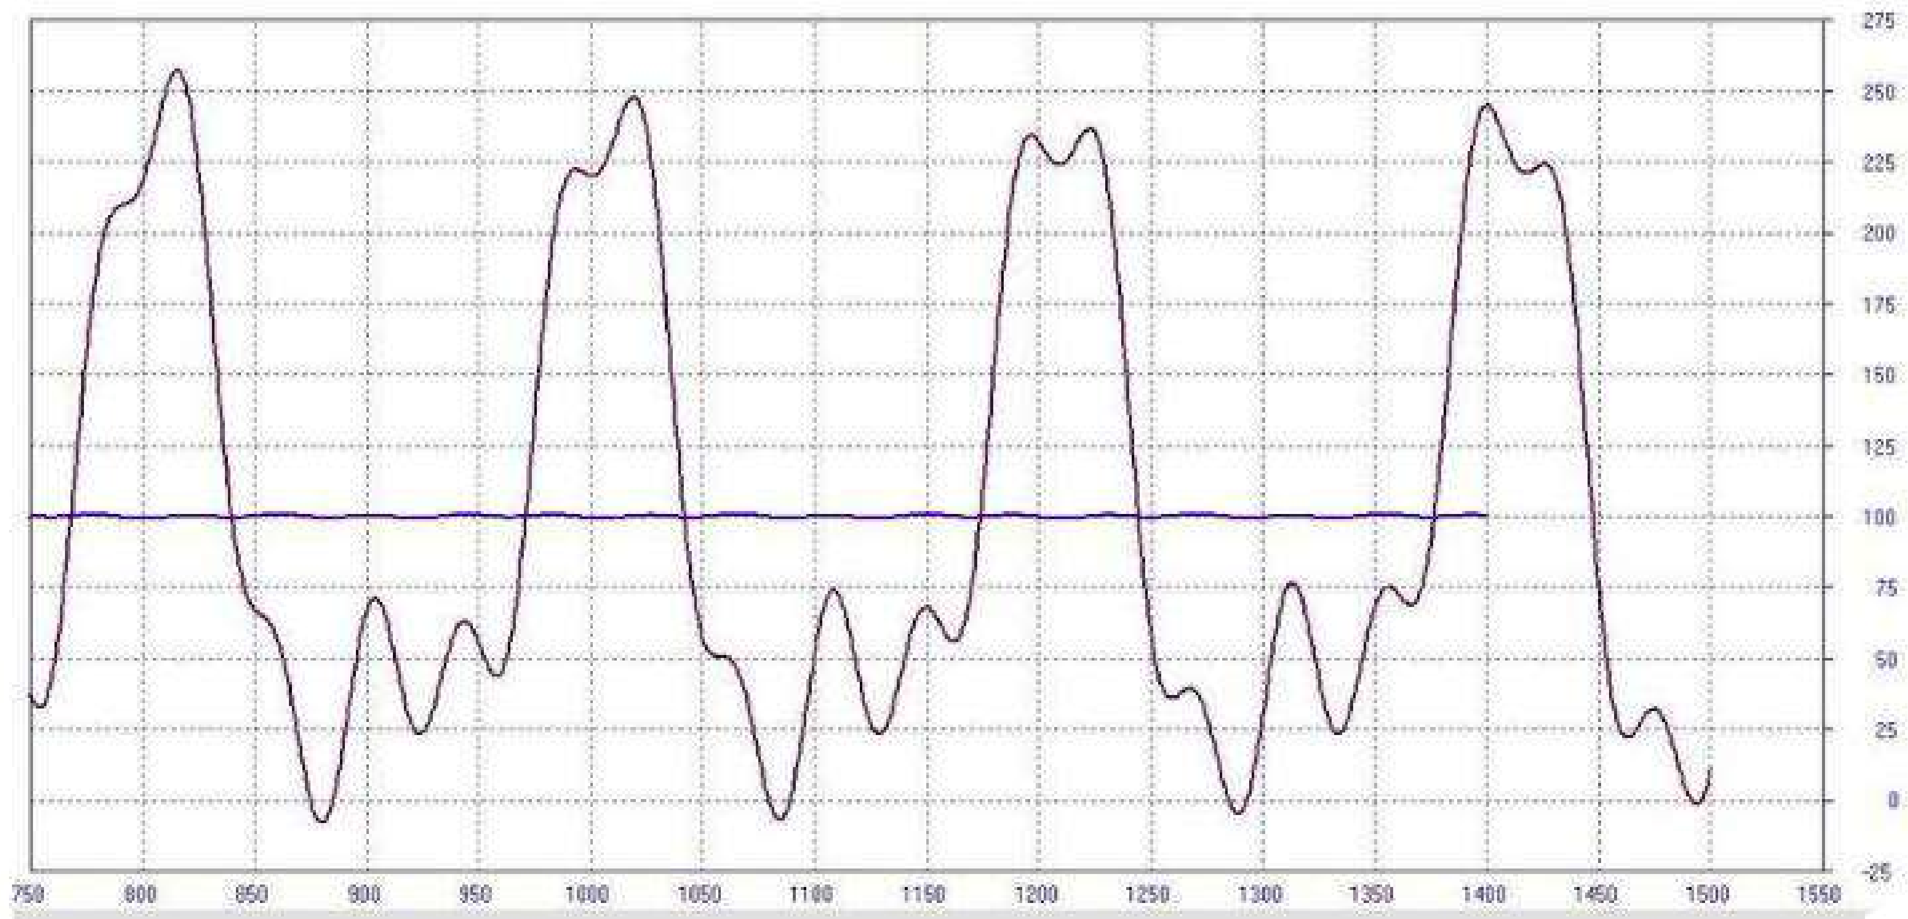


**101-day average**

We can now go through a sequence of testing a few averages with periods between 51 and 151. The simplest output is that given by a 101-day average. This is shown in Figure 8.11. This shows the cycle of period 201 days to be only slightly distorted.

***Figure 8.11 – A centred average of period 101 days has been applied. The output is a good representation of the original 201-day cycle. The amplitude is about 60 points.***

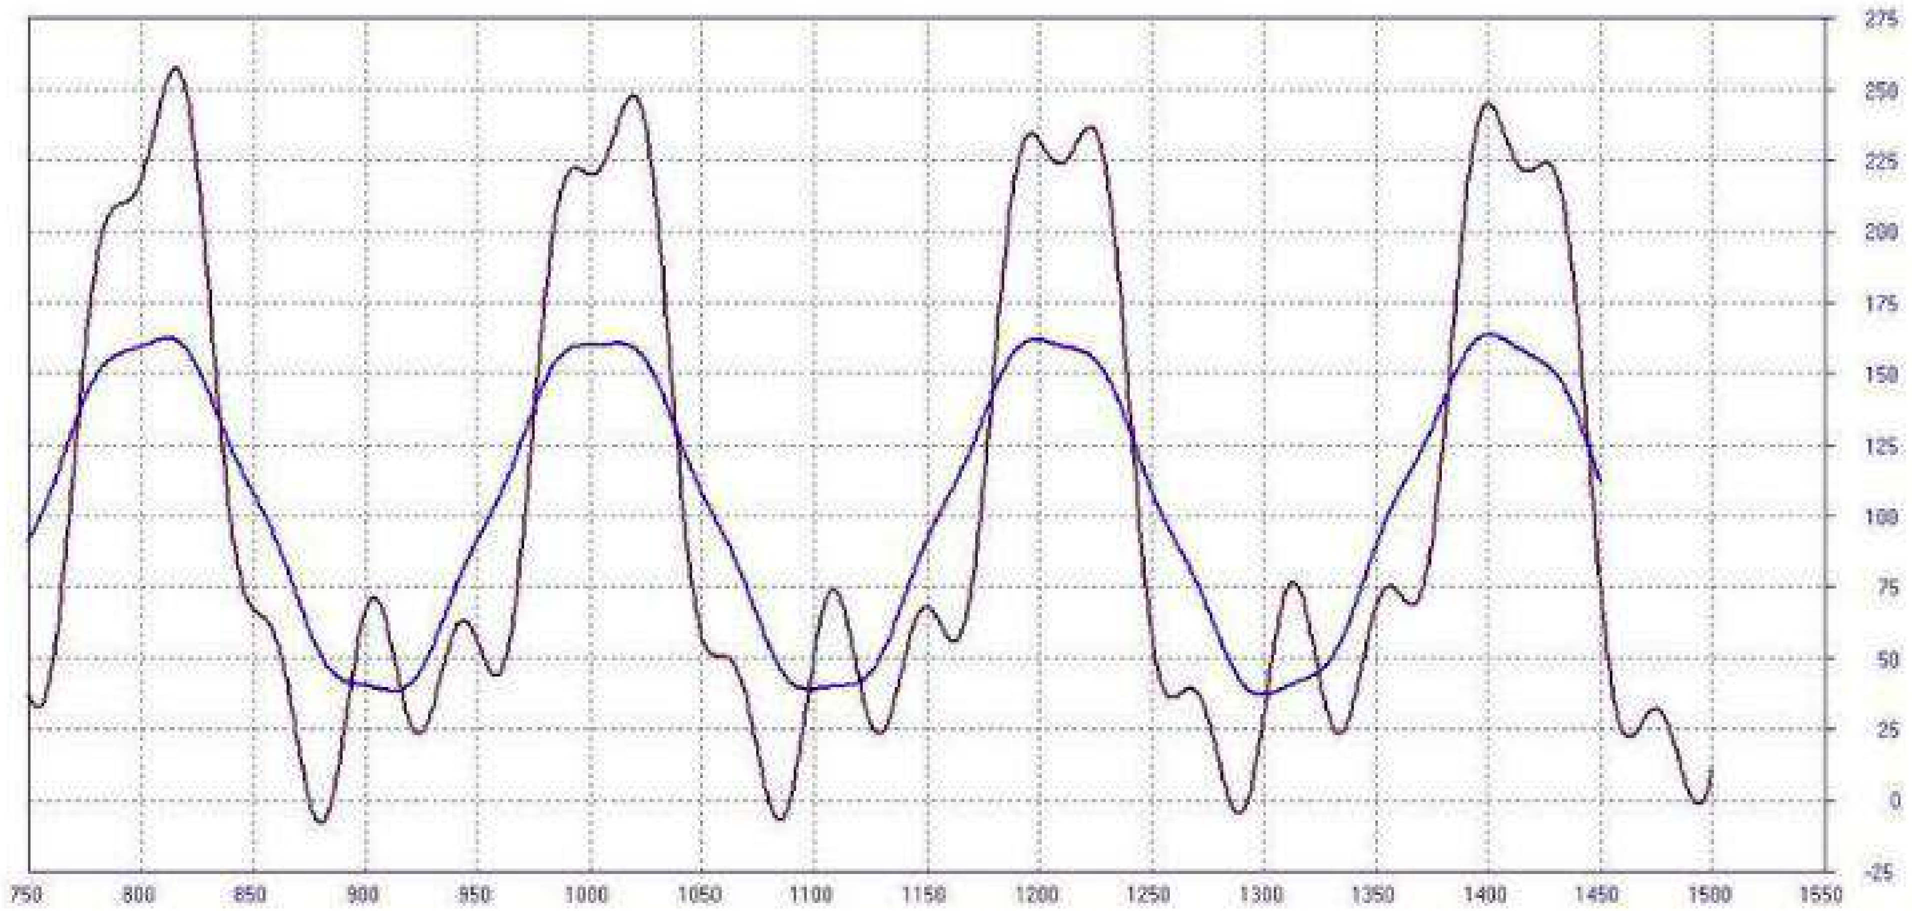


Since, of course, the application of a moving average will attenuate the cycle to an extent depending upon the relative values for wavelength and period of the average, it is to be expected that the amplitude of the 201-day cycle which has been isolated in this way will be reduced. This is true, since the amplitude of this cycle is only 62 units (half of the trough-to-peak height).

We have come as far as we can in the analysis of this complex set of three cycles, having found that there is a cycle of wavelength 201 days and amplitude greater than 62, plus a straight line with a value of 100 units.

We cannot isolate the cycles of wavelength 41 and 101 because if we use an average of period between 101 and 201 we will remove the cycles of wavelength 41 and 101 for the reasons stated earlier in the chapter. If we use an average of period between 41 and 101 we will remove the cycle of wavelength 41 but leave a mixture of the cycles of period 101 and

201. The amplitude of the cycle of wavelength 101 coming through in the output will depend upon how close the period of the average is to 101. The further towards 41 the period being used, the greater will be the amplitude of the cycle of wavelength 101 in the output. In other words, all we can hope to do is to isolate a mixture of cycles of wavelength 101 and 201.

Thus from our experiments with applying centred moving averages to mixtures of cycles we can see that the only cycle we can isolate from a mixture is the one with highest wavelength.

From this simple set of experiments we can reach a very important conclusion. For any value of period used for an average upon a mixture of cycles, the resulting output is essentially the sum of all cycles with wavelength greater than the period used for the average.

If there is a wide gap in wavelengths between the highest two wavelengths, as was the case with this example, then we can isolate the cycle with the highest wavelength when an average is used with period equal to the second highest wavelength. This removes this second highest cycle totally and greatly reduces the amplitude of those cycles of shorter wavelength.

This gives us an important understanding of what a centred moving average represents when applied to market data and is the key to predicting future trends by using moving averages.

## MOVING AVERAGE DIFFERENCES

Although it would appear that by using the methods discussed so far we can only isolate the cycle with highest wavelength, there is another property of moving averages than we can use to our advantage. We have already shown that an average with the same period as the wavelength of a cycle to which it is applied removes that cycle. The question is, where does this 'lost' data go, and can we retrieve it? The fact is, it isn't lost, and can easily be retrieved.

It has gone to form the difference between the output of the moving average and the data itself. For each value of the centred average, the difference between that value and the value of the data itself at that point will re-form the lost cycle. This is demonstrated in Figure 8.12. We can see that the vertical lines (which represent the difference) outline the original cycle. It might be more accurate to say that the values of these differences when plotted would be the original cycle. Note also that the amplitude of the reconstituted cycle is identical to that of the original.

***Figure 8.12 – The differences between the value of the centred average and the data itself will re-form the original cycle.***

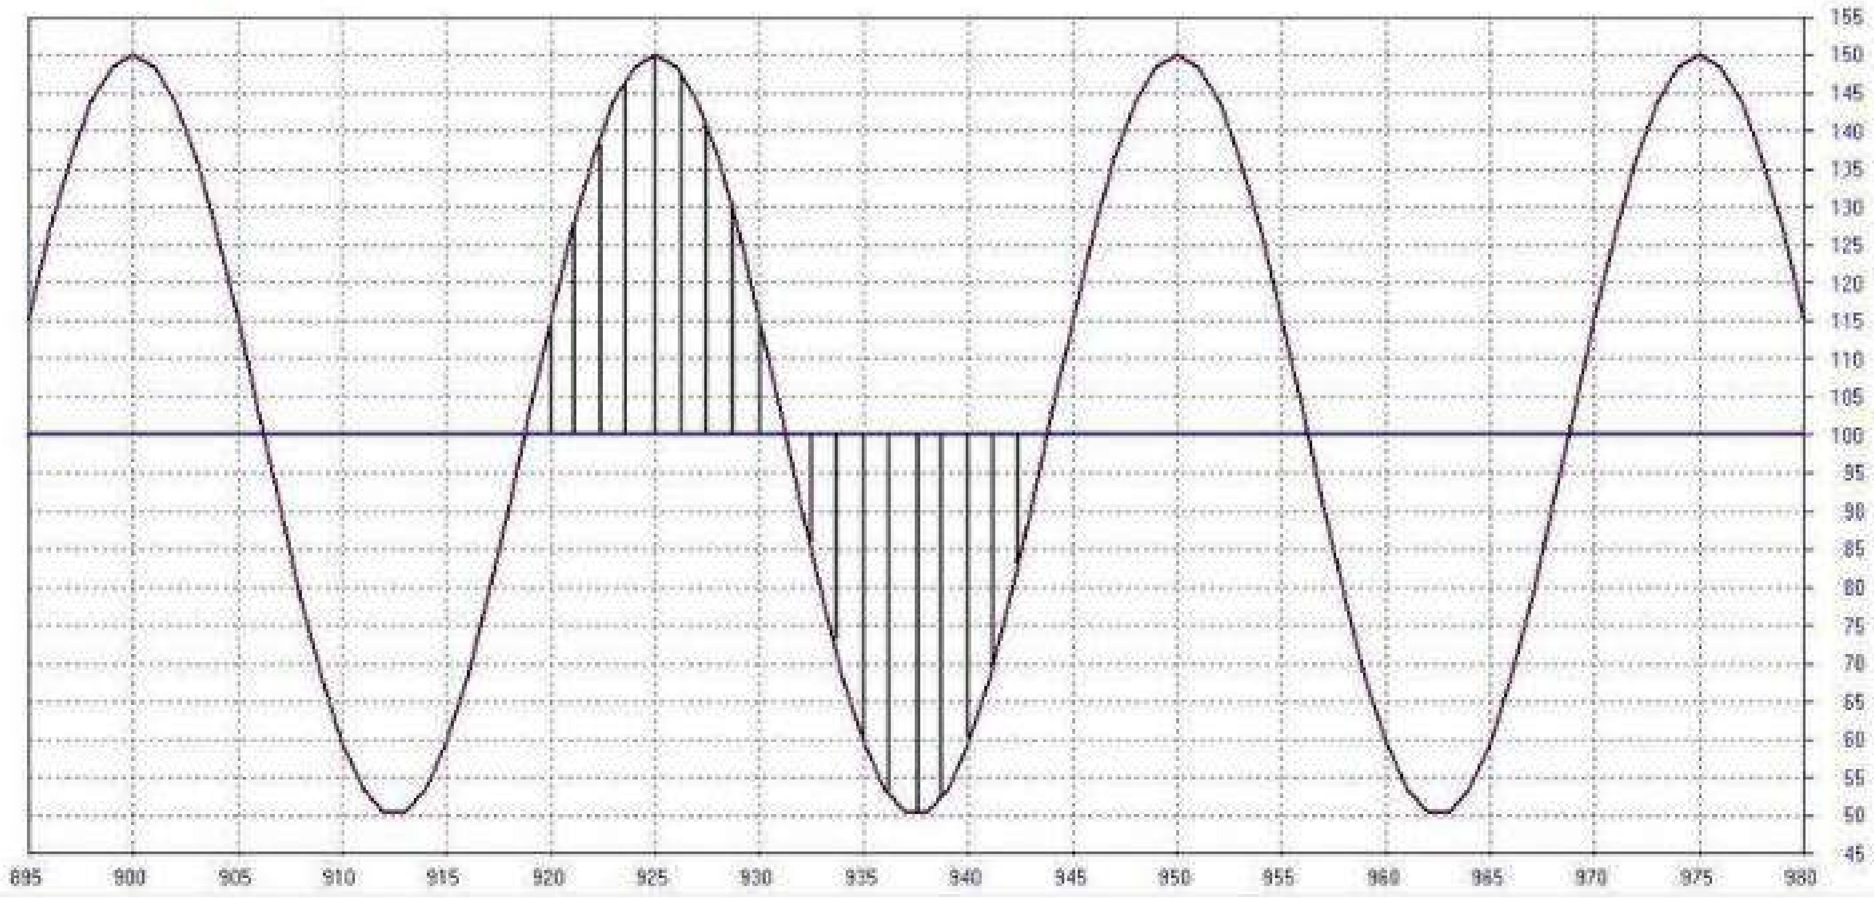

This is known as an average difference and will allow us to isolate the cycle of lowest wavelength from the complex mixture. Note that it is absolutely essential that a centred average is used, since the differences obtained from an average with no lag will be meaningless.

***41-day average difference***

If we apply a 41-point average difference to the mixture of cycles shown in Figure 8.8 we get the result shown in Figure 8.13. The cycle shown in the lower panel has a wavelength of 41 days if the peak-to-peak or trough-to-trough distances are measured.

***Figure 8.13 – The 41-point average difference of the mixture of cycles shown in Figure 8.8.***

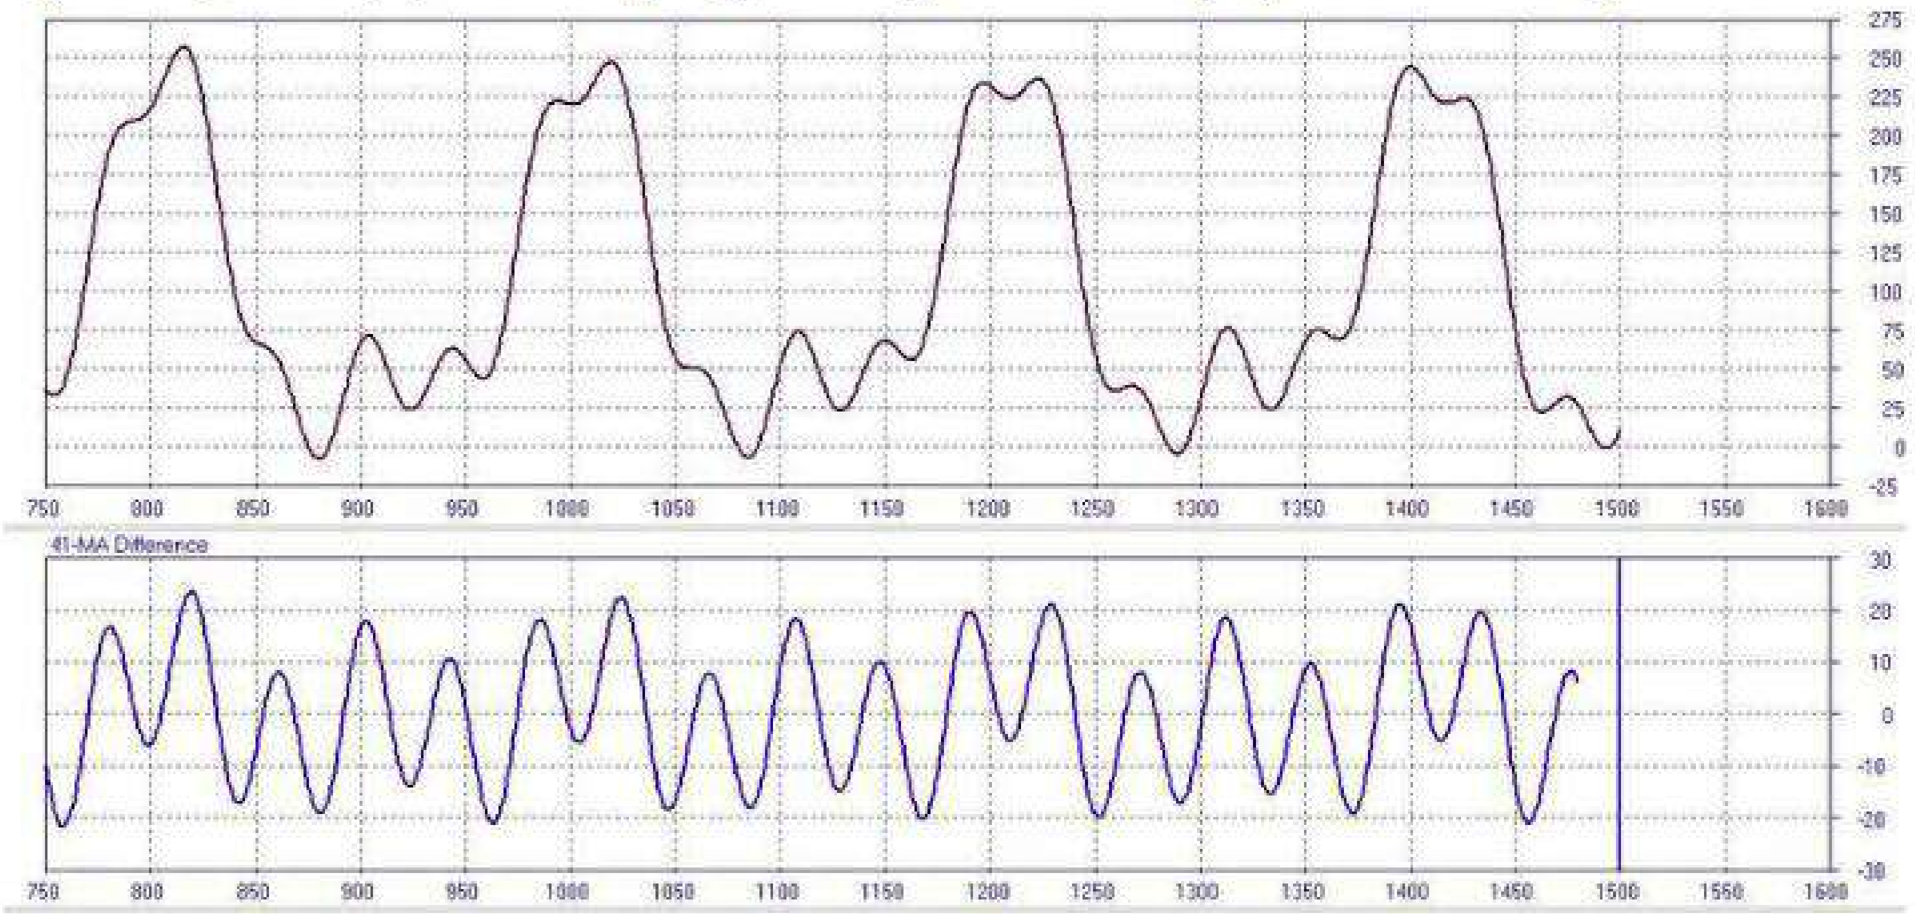


Although the isolated cycle is not absolutely clean, since there is a slight interference from the other cycles which causes the peaks and troughs not to lie in straight lines, there is no doubt at all that we have isolated and can identify the 41-day cycle. It is important also to note that the peaks and troughs of this cycle are in the same positions in time as the original 41-day cycle which was put into the mixture of cycles.

***101-day average difference***

In these two exercises we have isolated the 201-day cycle and the 41-day cycle. Neither of the methods that we have shown will isolate the 101-day cycle. To prove this point, the 101-point average difference is shown in Figure 8.14.

***Figure 8.14 – The 101-point average difference of the mixture of cycles shown in Figure 8.8 does not isolate the 101-day cycle.***

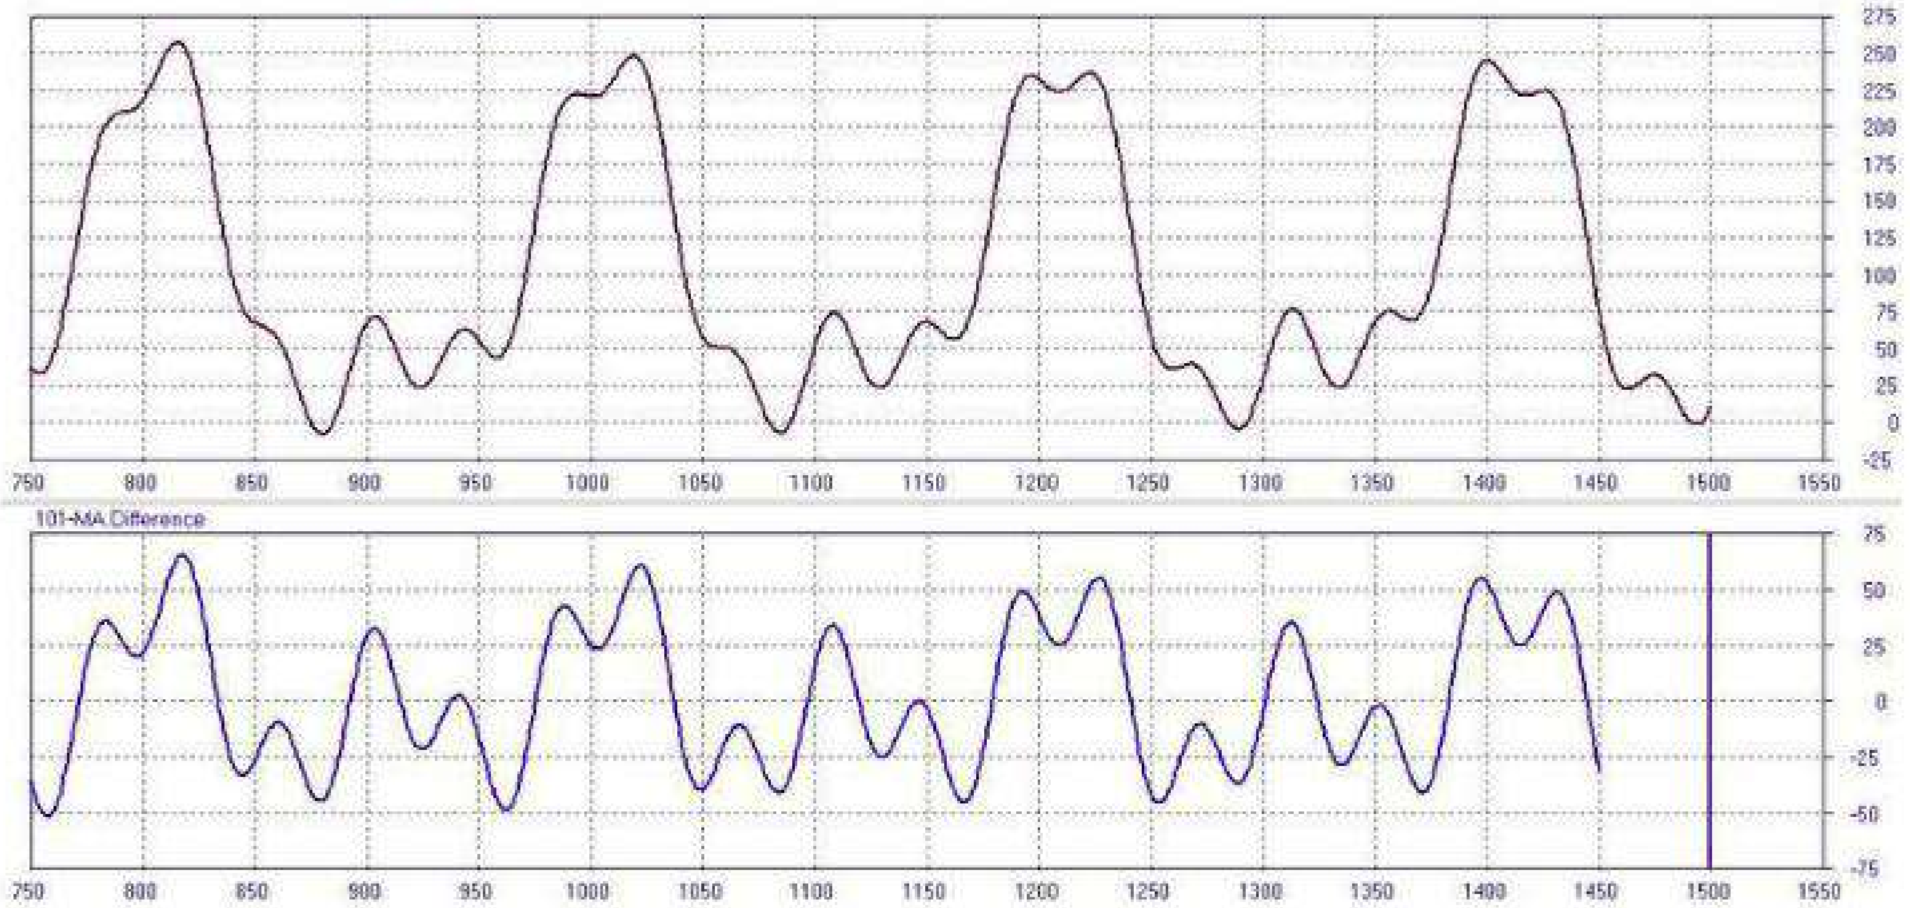


In this plot we still see in the lower panel a complex pattern which is due to the combination of the 41-day and 101-day cycles.

***Comparison of centred averages with average differences***

From the discussion so far, two very important points have been established from the properties of centred moving averages and average differences.

The **centred average** represents essentially the sum of all those cycles with wavelength greater than the period of the average. The greater the difference between the wavelength of a cycle and this period, the greater is the influence of this cycle on this sum.

The **centred average difference** represents essentially the sum of all those cycles with wavelength equal to or less than the period of the average. The less the difference between the wavelength of a cycle and this period, the greater the influence of this cycle is on this sum.

Thus we are still faced with the problem of how to isolate the 101-day cycle.

## DIFFERENCE OF TWO AVERAGES

The average difference method just described is actually a special case of a more general method. This is to calculate the differences between two centred averages. Thus the average difference discussed earlier is the case where the second average has a period of one. The two averages should be chosen so as to bracket the wavelength which it is desired to isolate. In the present case, periods of 135 and 51 were chosen, but a wide range of values can be used. The result of using these two values is shown in Figure 8.15.

***Figure 8.15 – The difference between the 51-point and 135-point centred averages is shown as the filled in portions of the plot. The original cycle data has been omitted for clarity.***

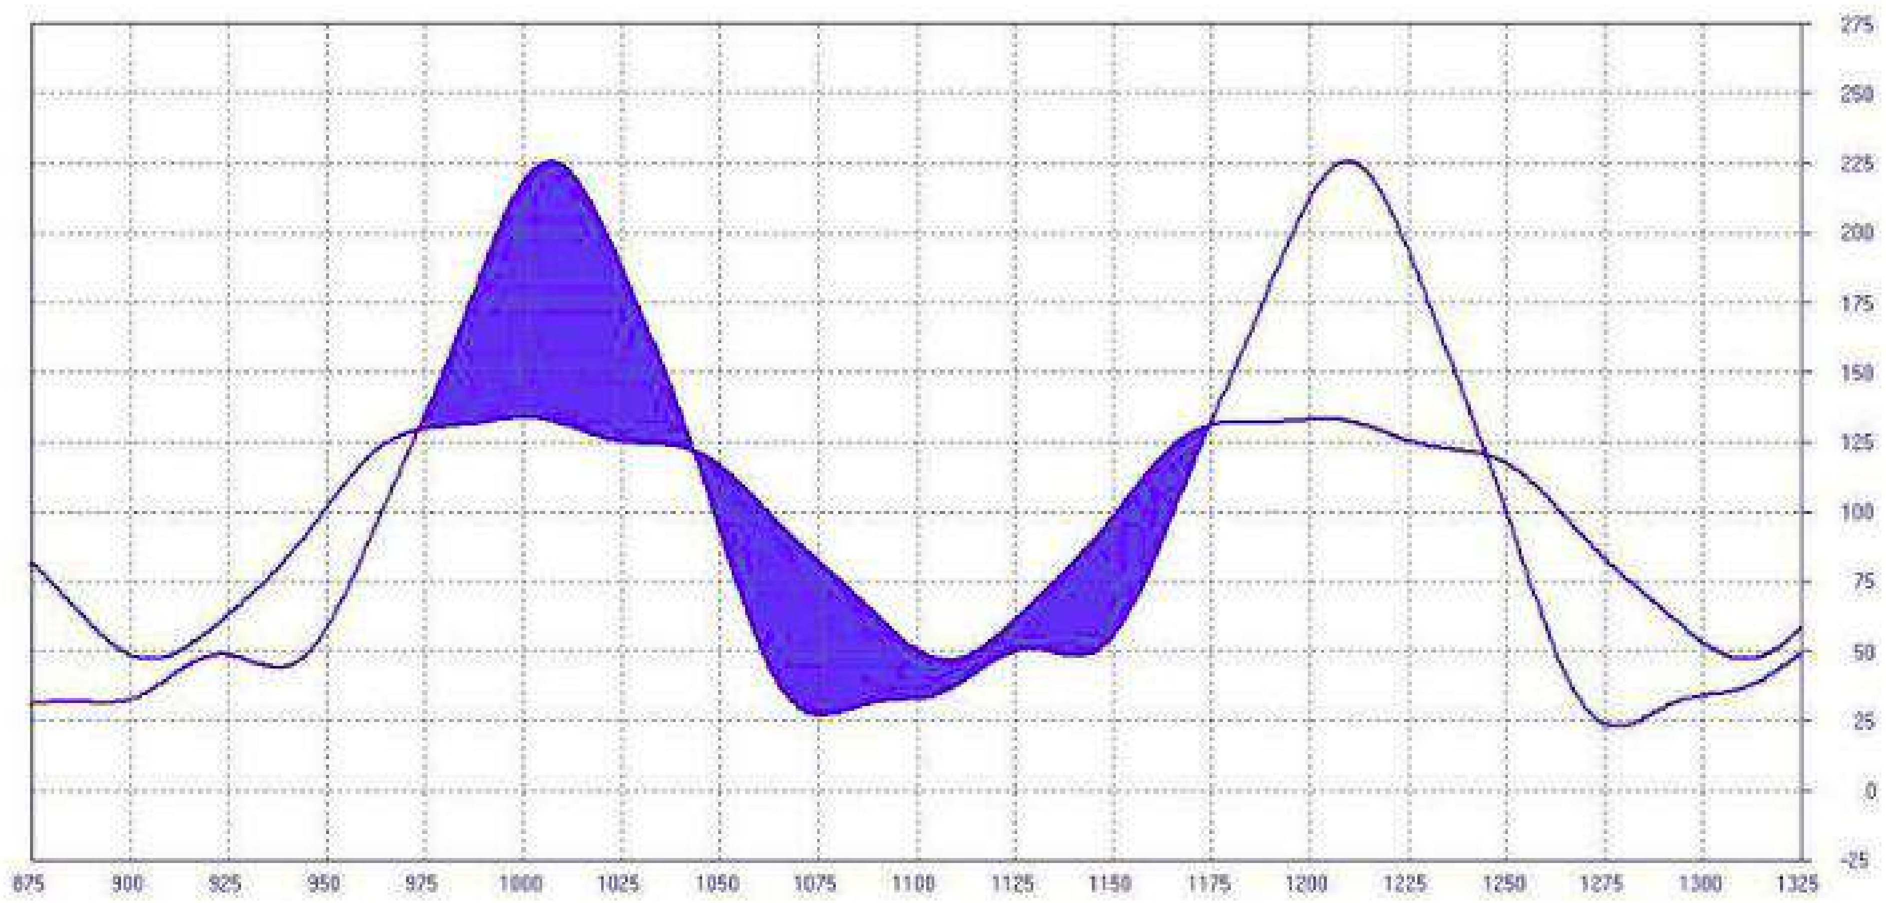


The original complex mixture of cycles has been omitted for clarity. The difference between the two averages is shown as filled in to highlight the process. When these differences are plotted, the cycle becomes obvious, as shown in the lower panel of Figure 8.16.

***Figure 8.16 – The result of plotting the differences between the two averages is shown in the lower panel.***

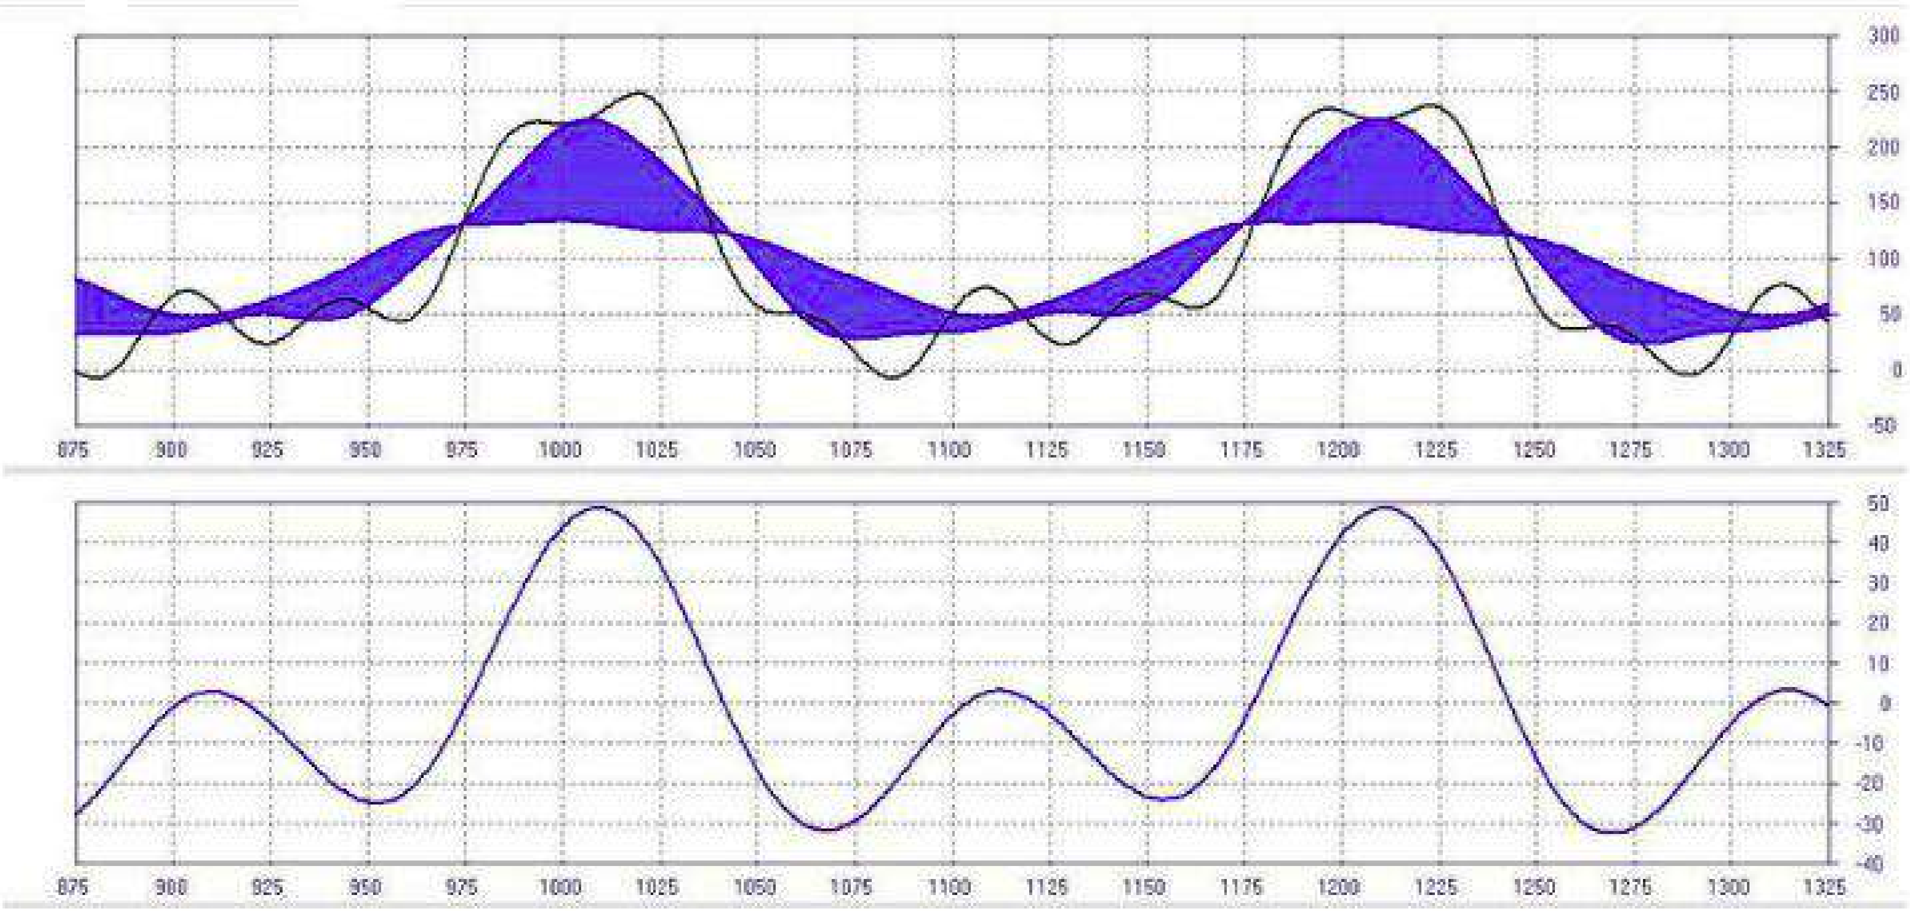


The two prominent peaks are separated by 202 days and since these represent two sweeps of the cycle, this is equivalent to a wavelength of 101 days. However, there is some distortion of the overall plot, so that the amplitude does not remain stable across the whole section. It is also worth pointing out that the difference between the positions of the prominent peaks and the next lesser peak is 104 days. In other words the minor peaks do not sit exactly halfway between the prominent peaks.

In spite of these two difficulties, we can see quite clearly that the use of the difference between two centred averages is an excellent way of isolating individual cycles from a complex mixture.