9. Averages as Proxies for Trends

In this chapter we will build upon the discussion in Chapters 7 and 8. We will do this by applying various centred moving averages to market data and seeing the effect of using different values for the period. When the moving average is plotted as a centred average, as in Figure 9.1, one feature will be obvious. That is the oscillation of the data around the average.

Figure 9.1 – A centred 15-day moving average applied to the Motorola stock price. The price oscillates around the average but remains within certain limits.

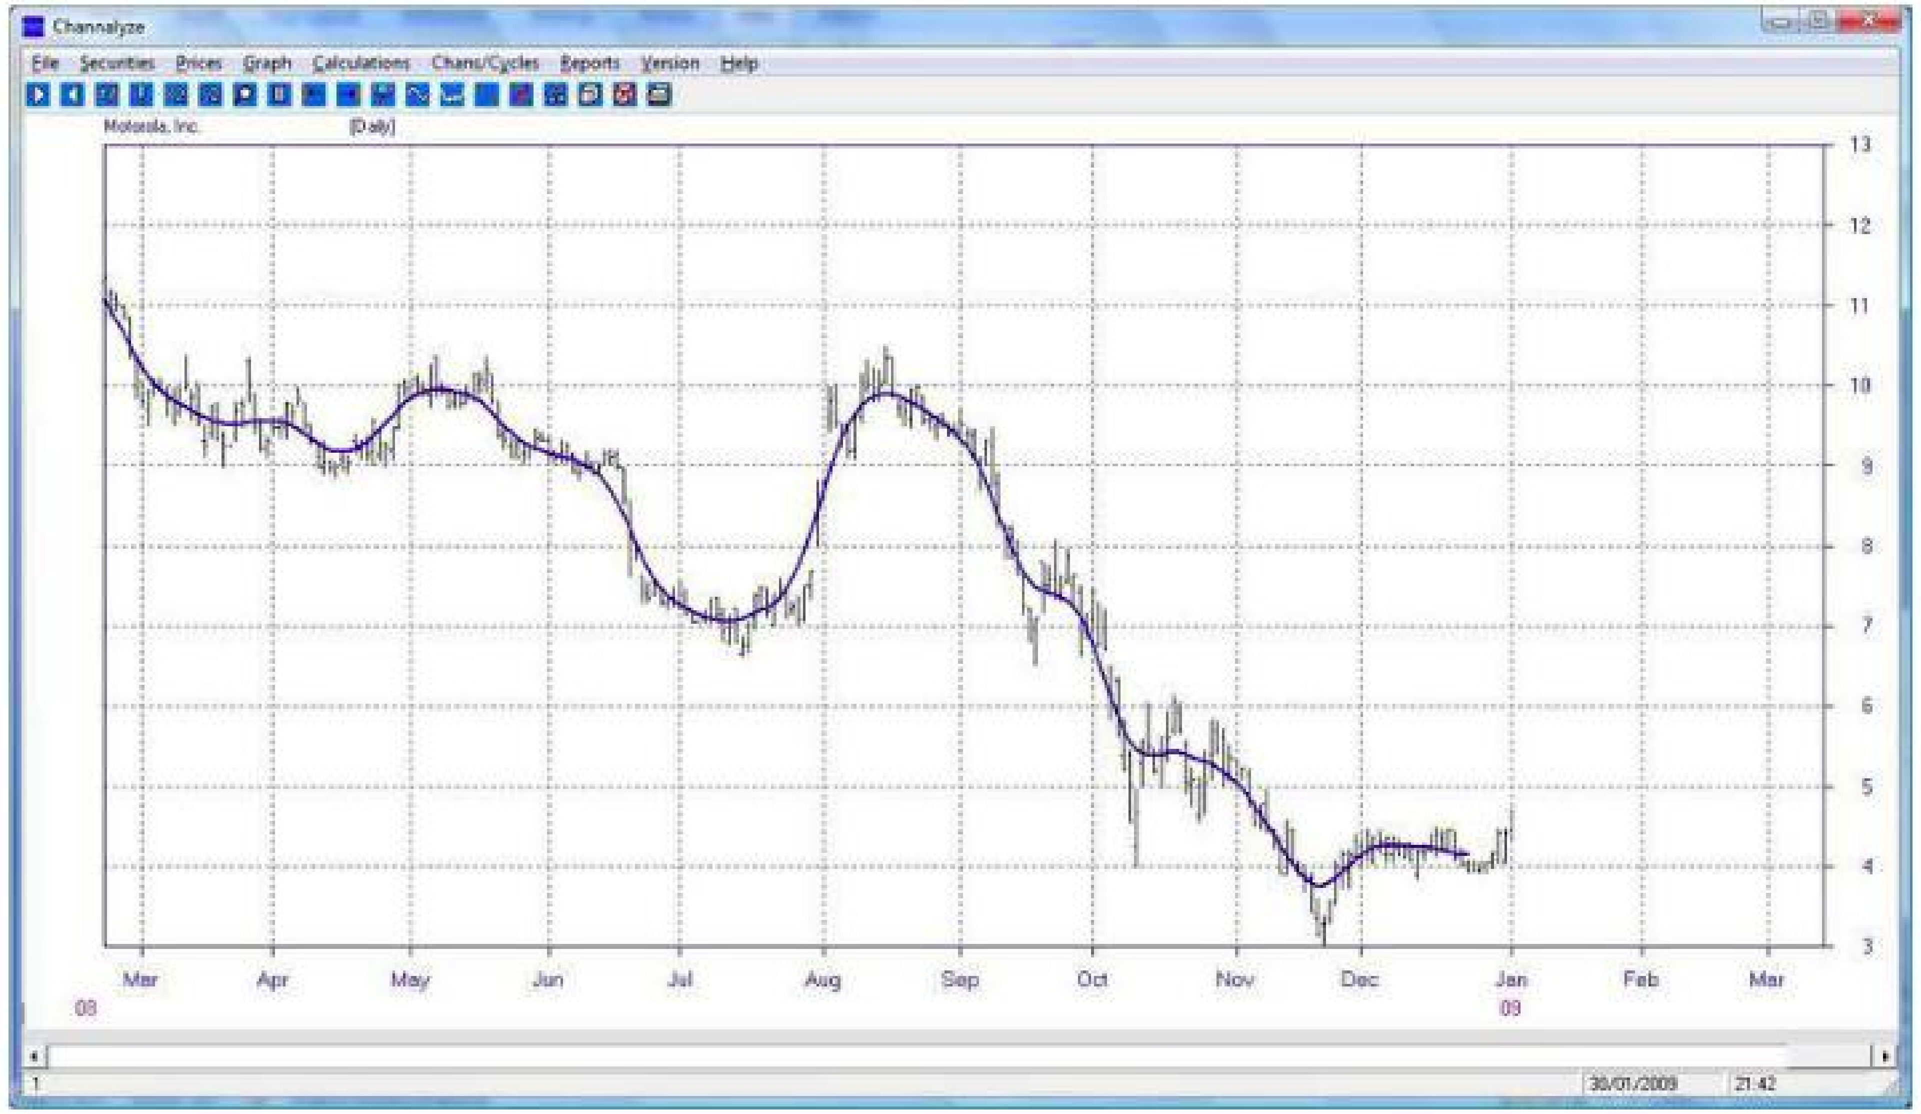

In the terms we used in Chapter 7 the centred average represents a trend. From our earlier definition (Chapter 7), whether this is a short-, medium- or long-term trend depends upon the period used for the average. Thus a short-term trend is caused by cycles of period 10 to 60 days, a medium-term trend by cycles of period 60 to 400 days and a long-term trend by cycles of period over 400 days.

15-DAY AVERAGE

An average of period 15 days will remove cycles with wavelengths less than or equal to 15 days and remove to a lesser extent cycles which are longer in wavelength. The attenuation of these cycles decreases as their wavelengths increase. In this context therefore this average represents all of the trends – short, medium and long, except for very short-term trends which have wavelengths equal to or less than 15.

101-DAY AVERAGE

In Figure 9.2 we show the same data with a centred 101-day moving average applied.

Figure 9.2 – A centred 101-day moving average applied to the Motorola stock price. The price oscillates around the average but remains within certain limits.

An average of period 101 days will remove cycles with wavelengths less than or equal to 101 days and remove to a lesser extent cycles which have a longer wavelength. The attenuation of these cycles decreases as their wavelengths increase. In the context of the definition of timescales, this average represents two of the trends – medium and long.

501-DAY AVERAGE

Shown in Figure 9.3 is the result of applying a centred average of period 501 days. Since this is over 400 days, this average represents the long-term trend.

Figure 9.3 – A centred 501-day moving average applied to the Motorola stock price. The price oscillates around the average but remains within certain limits.

As we see in Figure 9.3, and which was discussed briefly in Chapter 8, if we take the centred average to represent the trend, then because it terminates half a span of the average back in time, we have no idea at first glance as to the current state of the trend. Is it still headed downwards or has it changed direction during this gap of half of its period? In the case of the 501-day average in Figure 9.3, this gap amounts to 250 days, which is about one year of business days. As we have seen during the past year, huge moves in the market can and do occur over such a long period.

OSCILLATION OF DATA AROUND AN AVERAGE

The feature common to these three averages is of course that the data oscillates around them. A closer inspection shows that there is a limit to these oscillations in the sense that the excursions of the price data from the centred average reach a certain maximum before reversing direction. These excursions increase in extent as we move from the short-term average via the medium-term average to the long-term average. It is these excursions that we can turn into values that will produce probabilities for us. Of course, since these are centred averages, each one will terminate half a span back in time from the latest data point.

Probabilities From Centred Averages

We saw in the last chapter that the differences between a centred average and the data at each point in time were called, quite logically, average differences. It is worth restating what centred averages and centred average differences represent.

1. The *centred average* represents essentially the sum of all those cycles with wavelength greater 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 on this sum.
2. 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 on this sum.

HOW DATA IS DISTRIBUTED AROUND THE CENTRED AVERAGE

Imagine stretching out the average so that it becomes a straight line with the data still dispersed about it in the same way as the original. Now imagine looking down the whole length of this line. The effect is shown in Figure 9.4. The position of the average is shown by the vertical line in the centre of the rectangle. The -ve and +ve represent the extremes of movement below and above the position of the average. Quite clearly the data spends more time near the position of the average than it does at the extremes. It is from this distribution around the average that we can deduce the probabilities, just as was done for the daily changes in Chapter 4.

Figure 9.4 – How the data is distributed around the position of the average. The -ve and +ve extremes denote the furthest that the data moves from the position of the average.

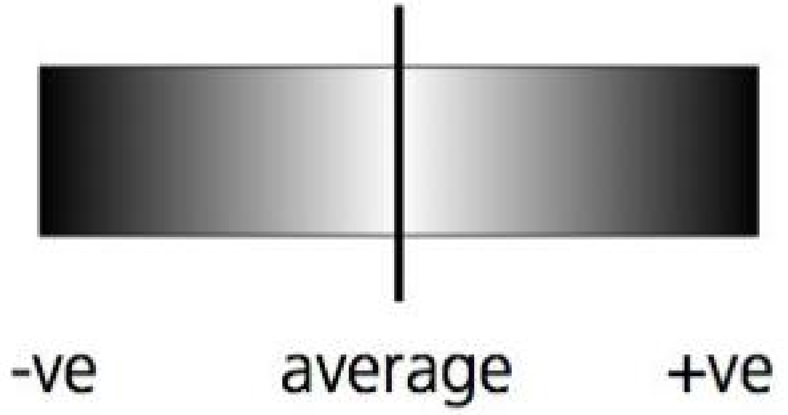

Obviously since the data spends very little time at the extremes, these represent **low probability positions**, whereas **the centre portion represents a high probability area**. Quite clearly, if the data is in a low probability area, it is unlikely to stay there as more data points arrive and the probability is high (but not 100%!) that it will move back towards the centre.

NUMERICAL DATA

Distribution around the 15-day average

Figure 9.4 indicates that the price spends most of the time in the middle third of the whole range of differences. This is confirmed by Figure 9.5. This shows the distribution of Motorola closing prices about the 15-day centred average which was shown in Figure 9.1.

Figure 9.5 – The distribution of Motorola data around the centred 15-day moving average. The data has been sorted into 200 bins (slots) and the vertical axis is the number of values that fall into each bin.

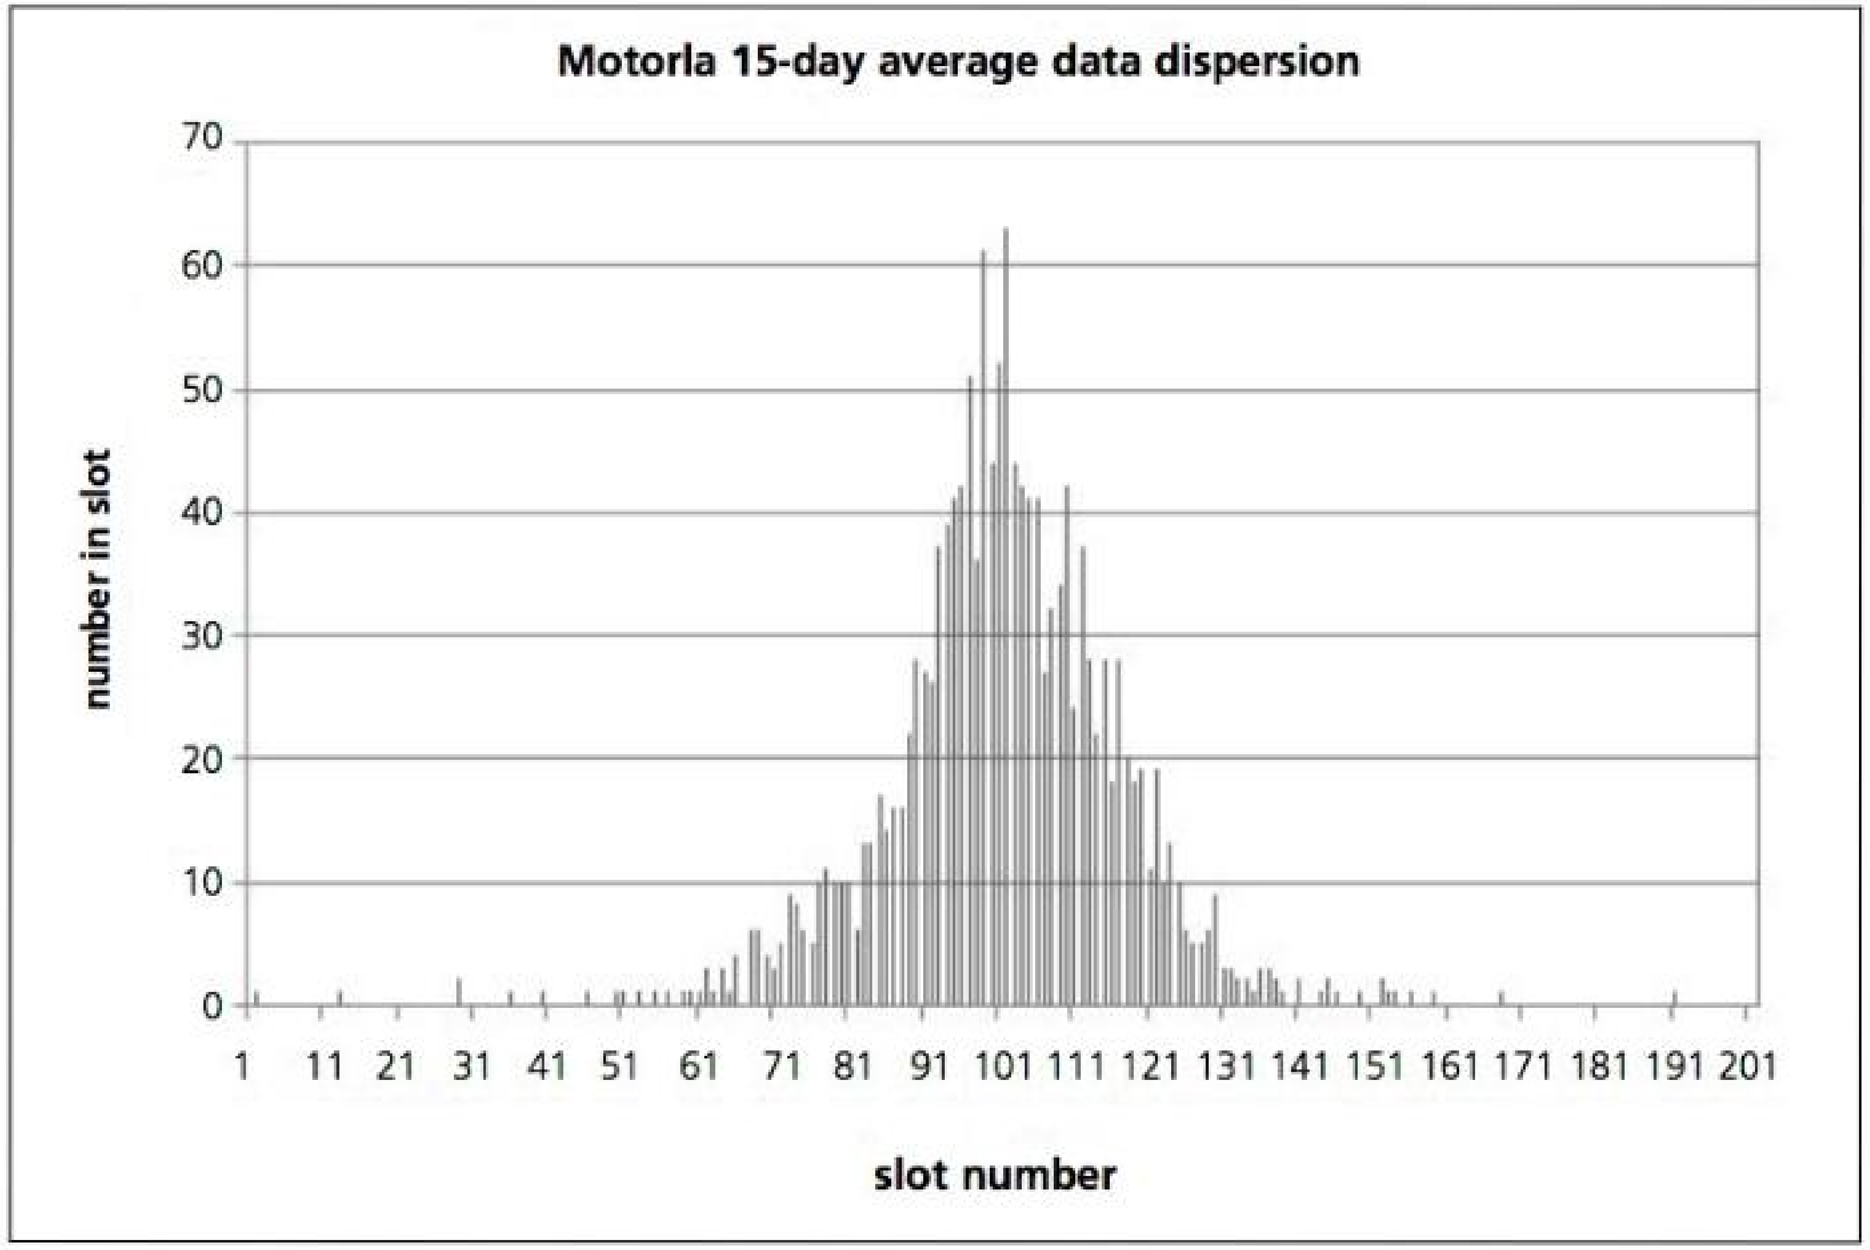

The data was allocated into one or other of 201 bins. The centre bin, bin number 101, corresponds to the position of the moving average itself. Bin 1 corresponds to the largest negative value, i.e. when the price is below the average line and at an extreme value. Bin 201 corresponds to the largest positive value, i.e. when the price is above the average and at an extreme value.

This type of distribution is of course familiar from Chapter 4, in which the distribution of daily price changes was investigated. Quite clearly, the price data spends most of its time around the centre position of the moving average (bin 101) and very little time at the extremes. Since these data are a set of numerical values, we can, just as in Chapter 4, use the concept of standard deviation to determine probabilities if we make the assumption that these data are normally distributed.

As was stated in Chapter 4, in a normal distribution:

- 68.27% of the points lie within one standard deviation of the mean
- 95.45% of the points lie within two standard deviations of the mean
- 99.73% of the points lie within two standard deviations of the mean.

If the data are normally distributed then we can produce probabilities for the location of any future particular data point relative to a known position of the centred moving average. However, as we saw previously, calculating exact probabilities is not a straightforward process and approximate values will serve us well enough for the purposes of this chapter. All we need in our approach is to grasp the fact that there is a low probability that a data point should be at the extreme left or extreme right of the plot shown in Figure 9.5. If, therefore, a data point is already in this position, then it is highly unlikely to stay there, and the next values should move towards the area of high probability, i.e. the central position of the average.

The standard deviation of the values displayed in Figure 9.5 was $13.3859. Two standard deviations amount to a value of $27 (rounded off). Therefore we can say that 95% of the points lie between bins 74 and 128.

Distribution around 101-day average

In Figure 9.6 we show the dispersion of Motorola data around the 101-day centred average. The distribution in this case is much flatter than was the case for the 15-day average. The standard deviation in this case was $7.76. Therefore, in round numbers, two standard deviations are $15. Thus we can say that 95% of the points lie between bins 86 and 116.

Quite clearly therefore the spread of bins that contains 95% of the values is much higher in the case of the 15-day average than it is in the case of the 101-day average.

Figure 9.6 – The dispersion of Motorola data around the centred 101-day moving average. The data has been sorted into 200 bins (slots) and the vertical axis is the number of values that fall into each bin.

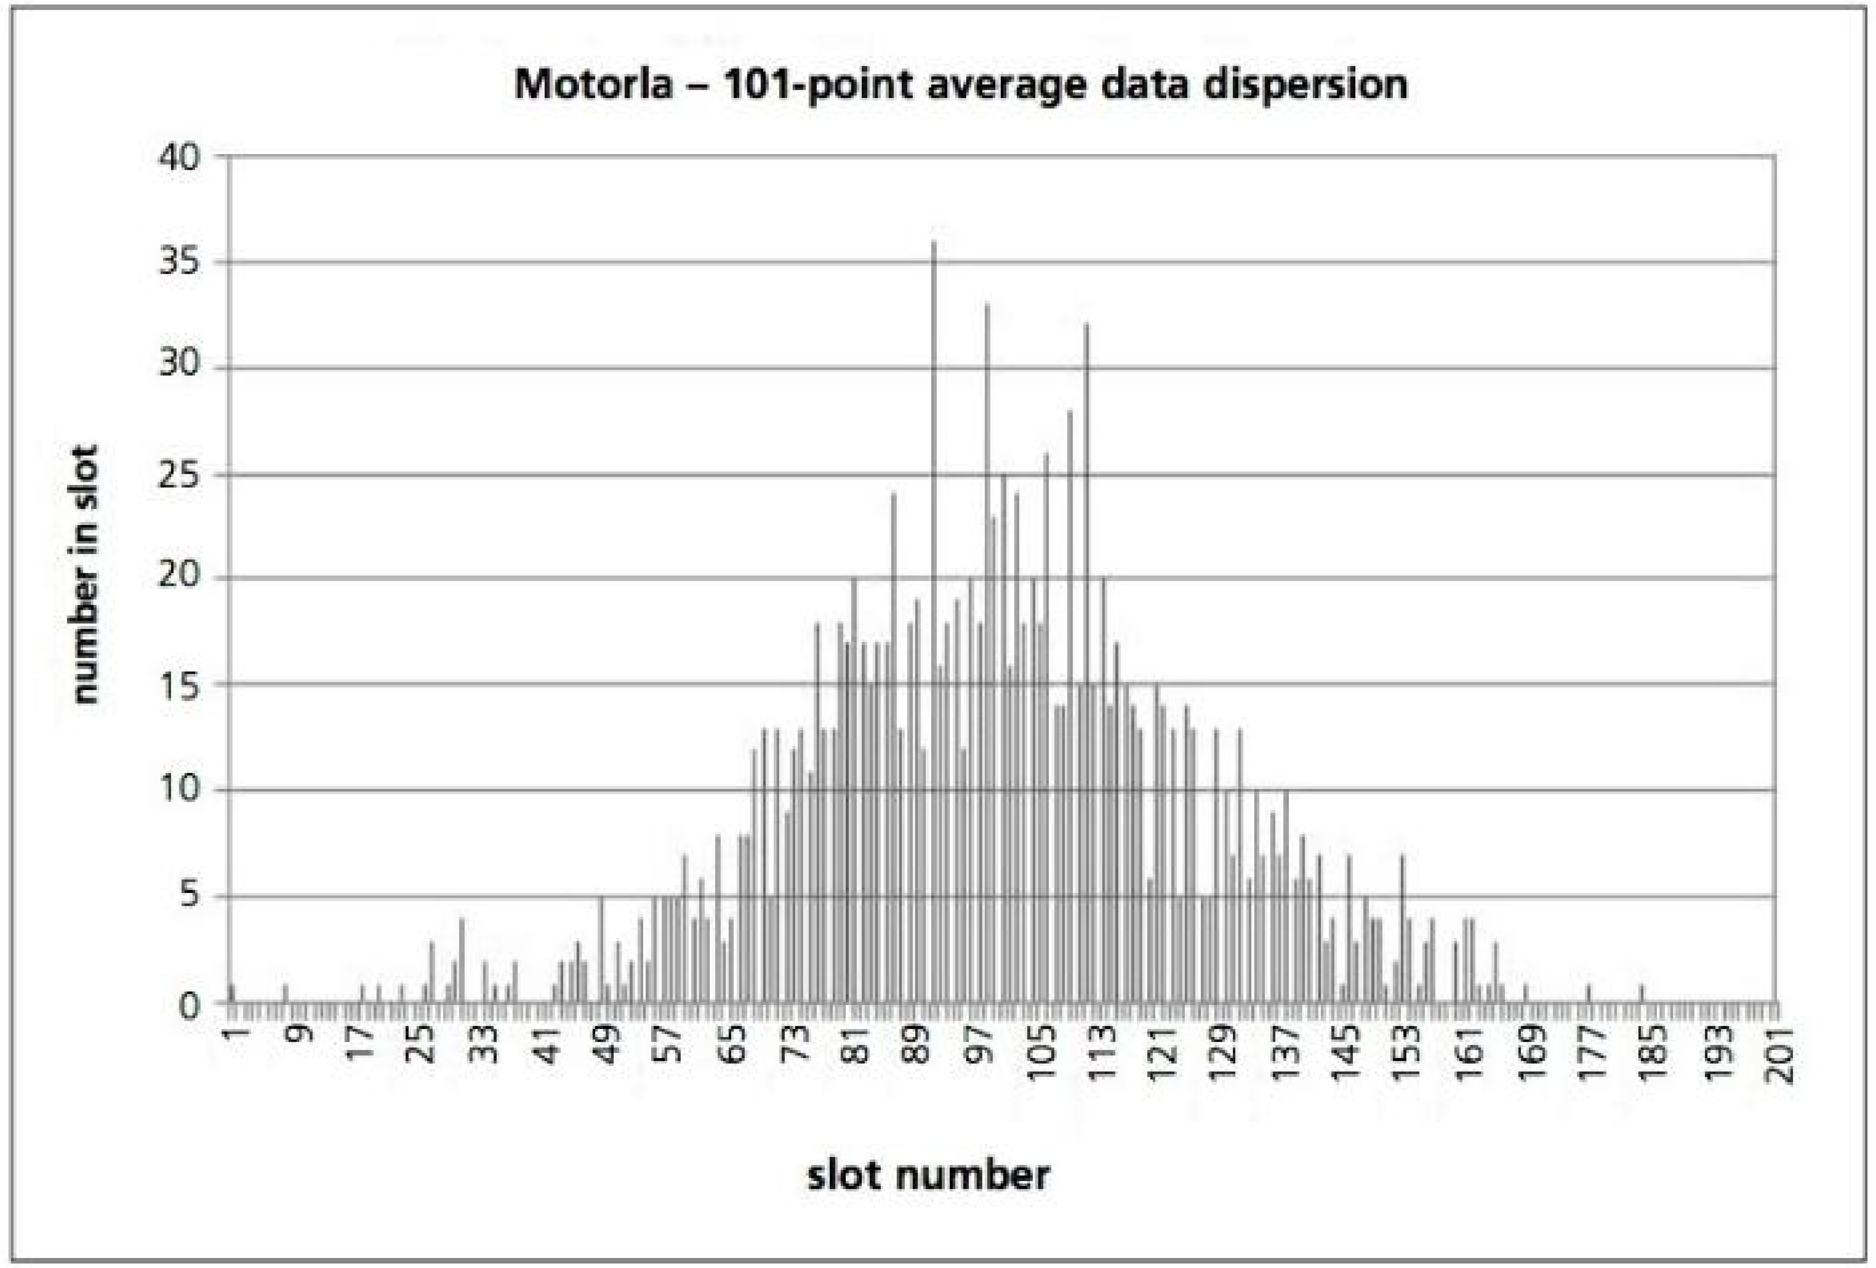

In Figure 9.7 we show the dispersion of Motorola data around the 501-day centred average. The distribution in this case is much flatter than was the case for the 15-day average. The standard deviation in this case was $4.8. In round numbers therefore, two standard deviations are $10. Thus we can say that 95% of the points lie between bins 91 and 111.

Figure 9.7 – The dispersion of Motorola data around the centred 501-day moving average. The data has been sorted into 200 bins (slots) and the vertical axis is the number of values that fall into each bin.

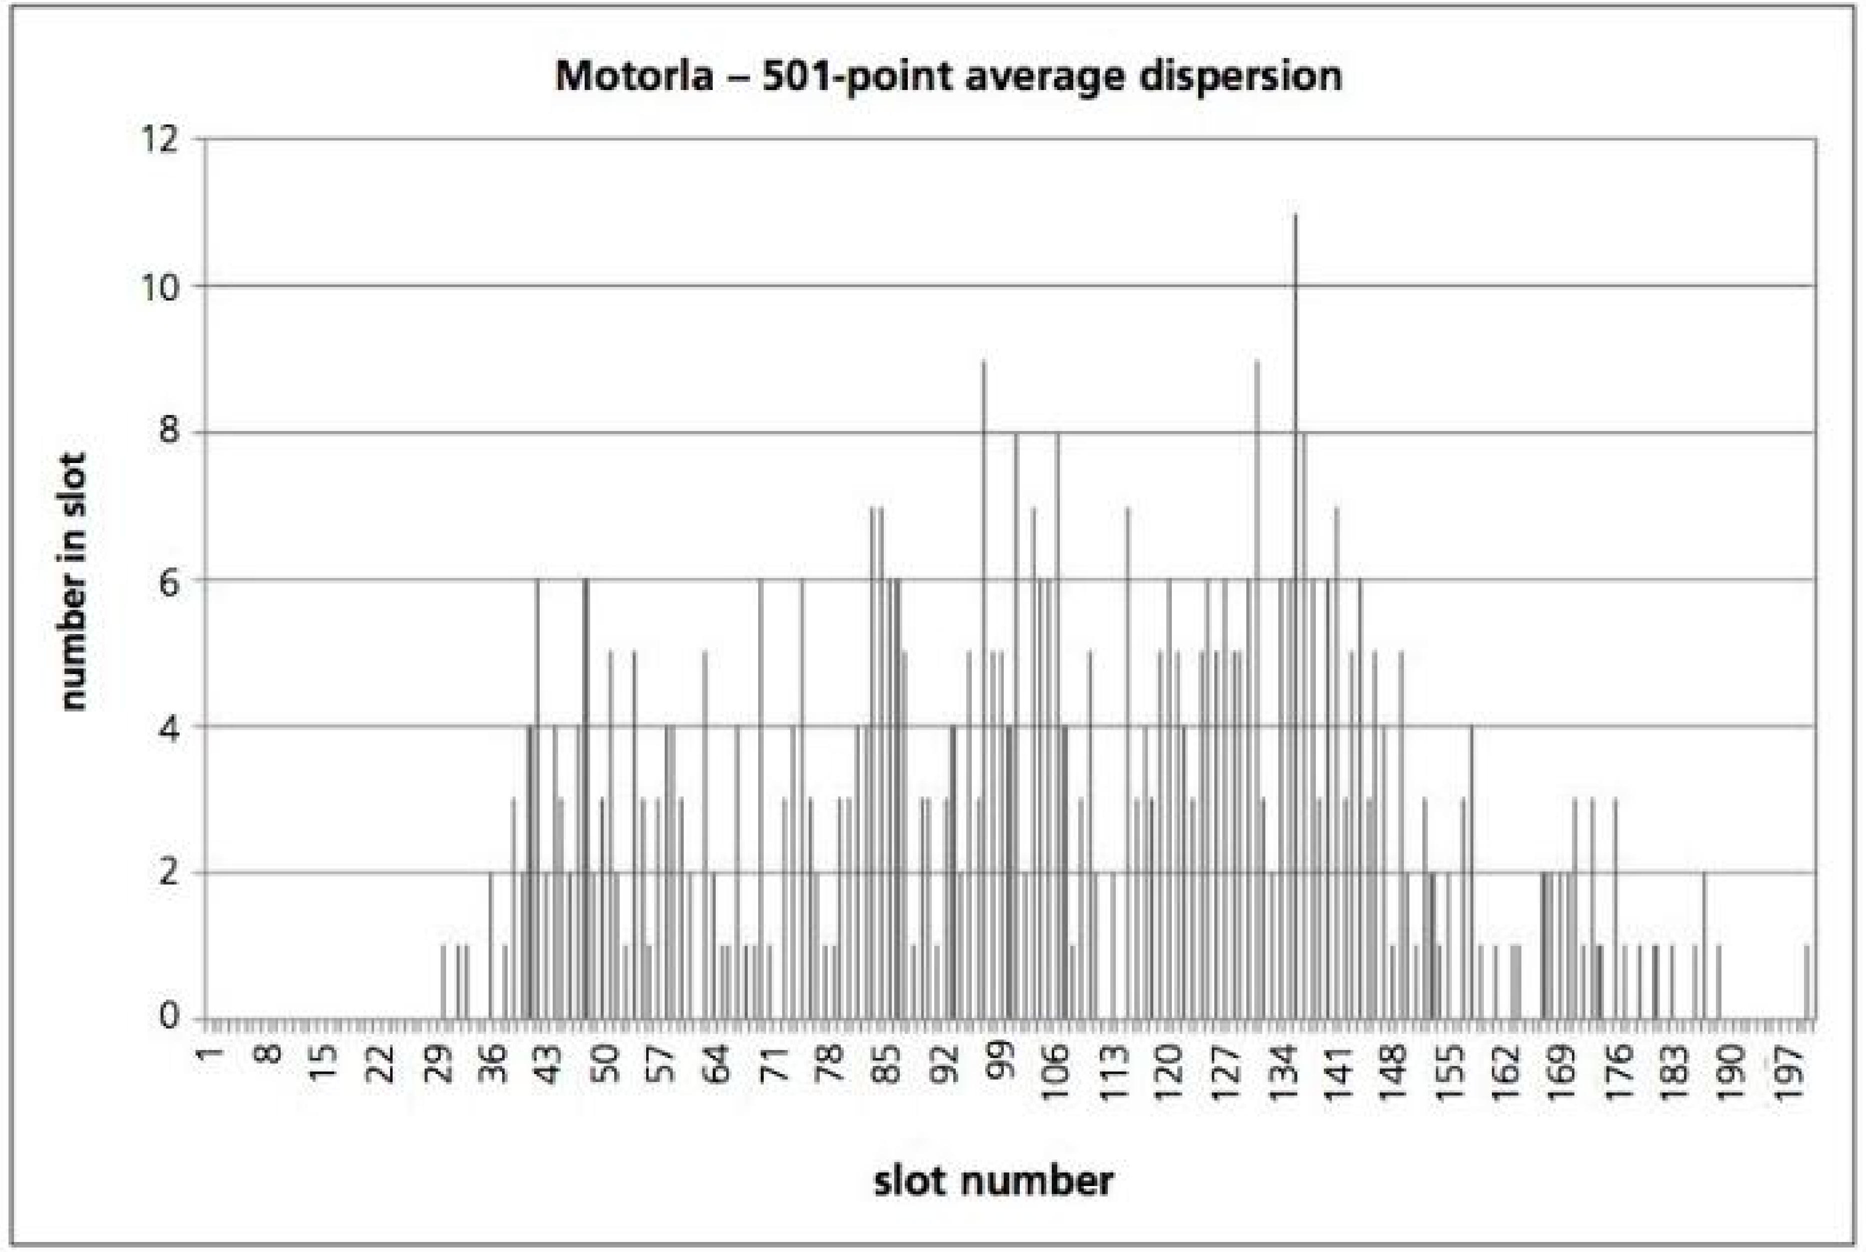

Quite clearly therefore the spread of bins that contains 95% of the values is much higher in the case of the 15-day average and 101-day average than it is in the case of the 501-day average.

This gradual erosion of the usefulness of using probability and distribution around moving averages echoes that discussed at the end of Chapter 4 for longer-term price changes. The practical limit for the period of the average in carrying out such studies is around 201.

Boundaries and Channels

The advantage of putting the differences between the centred average and the associated data point into bins is that it is easy to calculate which two bin numbers will frame, say, 95% of the values between them. As shown above, the calculation of standard deviation will enable these two bins to be defined. As we saw, for the 15-point average they were 75 and 127, for the 101-point average they were 86 and 116, and for the 501-point average they were 91 and 111.

Knowing the bin numbers, we can see what value of distance from the central average these represent and then draw lines at this constant value below and above the centred average line. This approach is shown for the 101-point centred average in Figure 9.8. Here we have drawn constant depth lines above and below the centred average. This can also be done with a pencil on a chart on which is plotted a centred 101-day average. From the bin numbers it is simple to calculate the positions of these lines so that, for example, only 5% of the data points lie outside these boundaries. These boundaries then constitute a *channel*, which is why this method is called channel analysis. Using such a visual method rather than staying with numerical distributions makes it very much easier to understand what is happening to the data as it oscillates around the central average line.

Quite clearly, as shown in Figure 9.8, there is a high probability that the price will retreat once it reaches the boundary. Since in this particular case the boundaries are set so that only 5% of the data points will lie outside, there are some points of approach to the boundary in which the boundary is then penetrated. However, this is only temporary, since there is a low probability of the next data point to arrive remaining in this location.

Figure 9.8 – The centred 101-point average now has a line at a constant distance above and below the average. The position of these lines has been chosen such that only 5% of the data points lie outside of them.

These points of penetration are shown in better detail in Figure 9.9.

Figure 9.9 – This expanded portion of Figure 9.8 shows the points of penetration of the boundaries.

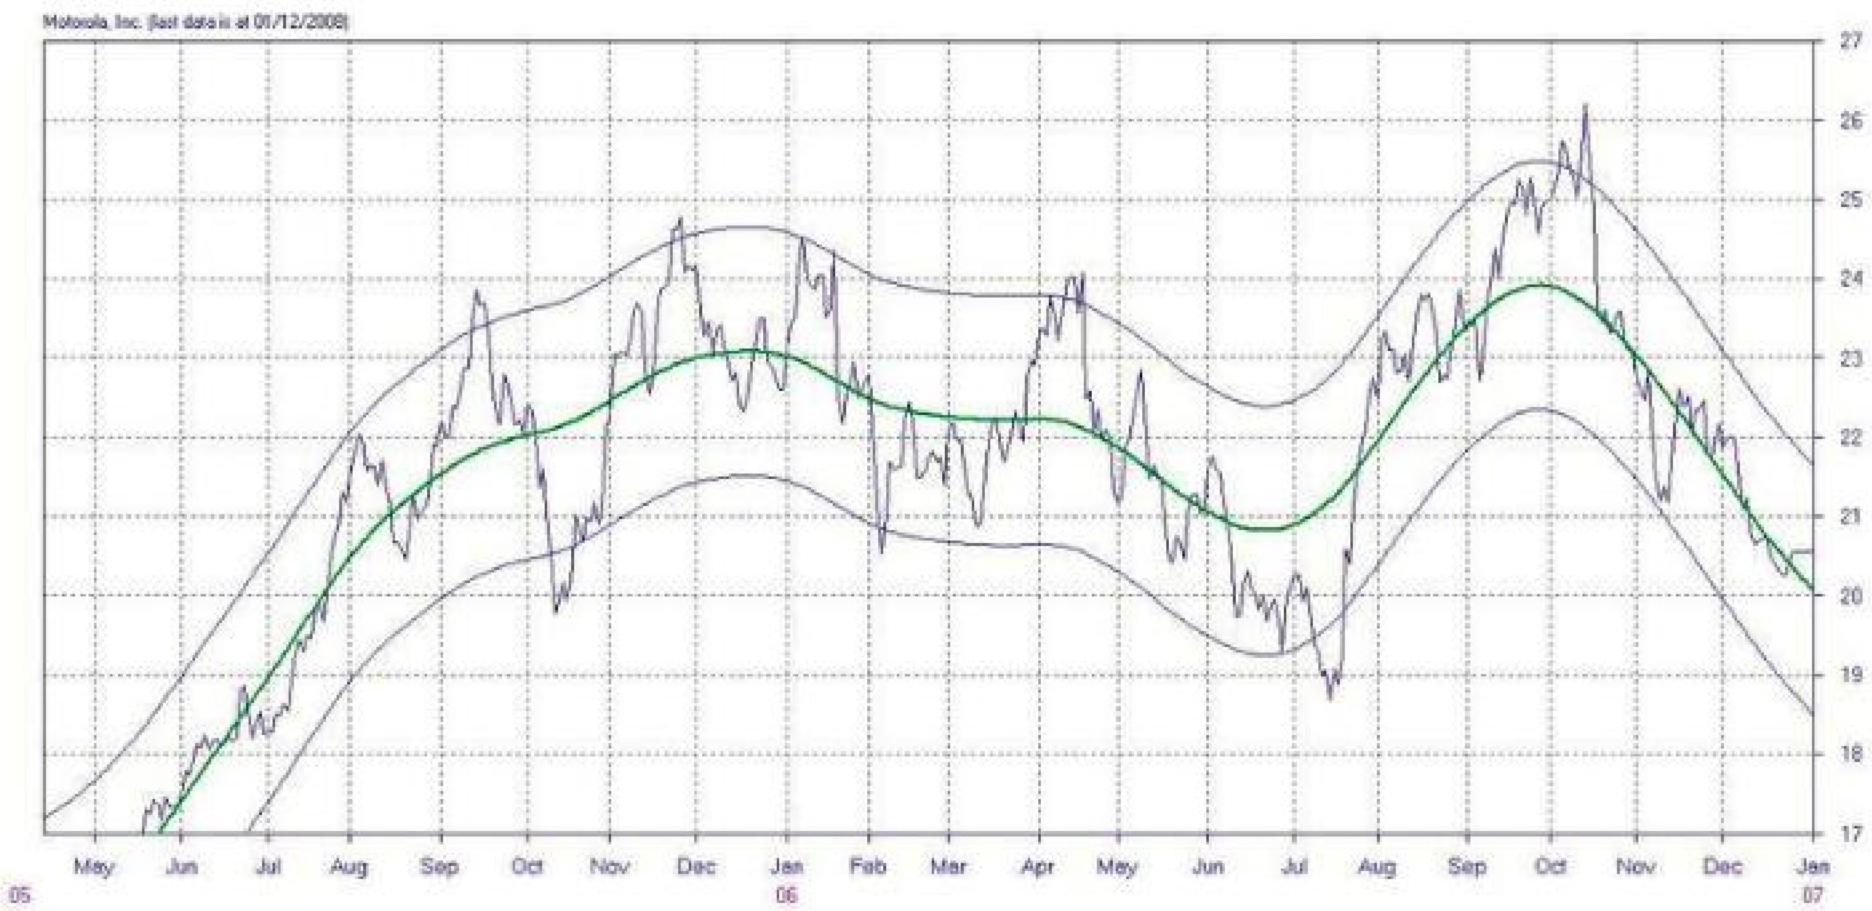

It is possible with the Channalyze software to adjust the position of the boundary so as to increase or decrease the number of data points which lie outside the boundary. In Figure 9.10 the channel boundaries have been adjusted so that no points lie outside the channel. However, this may be counter-productive. In Figure 9.10 the price movement did not reach a boundary from May 2003 until July 2006. Since the idea behind channel analysis is to provide buying and selling signals, this situation of zero points lying outside of the channel would have been ineffective over that period of time. As always in trading there is a balance to be struck between too many signals which might be false and too few so that traders will remain inactive for too long.

Figure 9.10 – Here the channel boundaries have been adjusted so that no points lie outside the channel.

WHAT CHANNELS TELL US

We saw from Chapter 8 that the centred average represents essentially the sums of all of those cycles of wavelength greater than the period used for the average.

The movement of the price data within the channel is due of course to the sum of all of those cycles which have wavelengths less than or equal to the period used for the centred average upon which the channel is based. Using a channel to analyse the price movement separates the trends into two categories, thus enabling a better focus on the overall movement of the stock:

1. The channel itself represents the longer-term trend – whether this is short-, medium- or long-term depends on the value chosen for the average.
2. The movement within the channel represents the shorter-term trend – whether

this is short-, medium- or long-term depends on the value chosen for the average.

This point about the movement of price data within the channel representing the sum of cycles with wavelengths less than the period of the average can be demonstrated in another way. This is done by plotting the 101-point average difference. This plot is shown in Figure 9.11.

Figure 9.11 – The upper panel shows the 101-day channel and the lower panel the 101-day average difference.

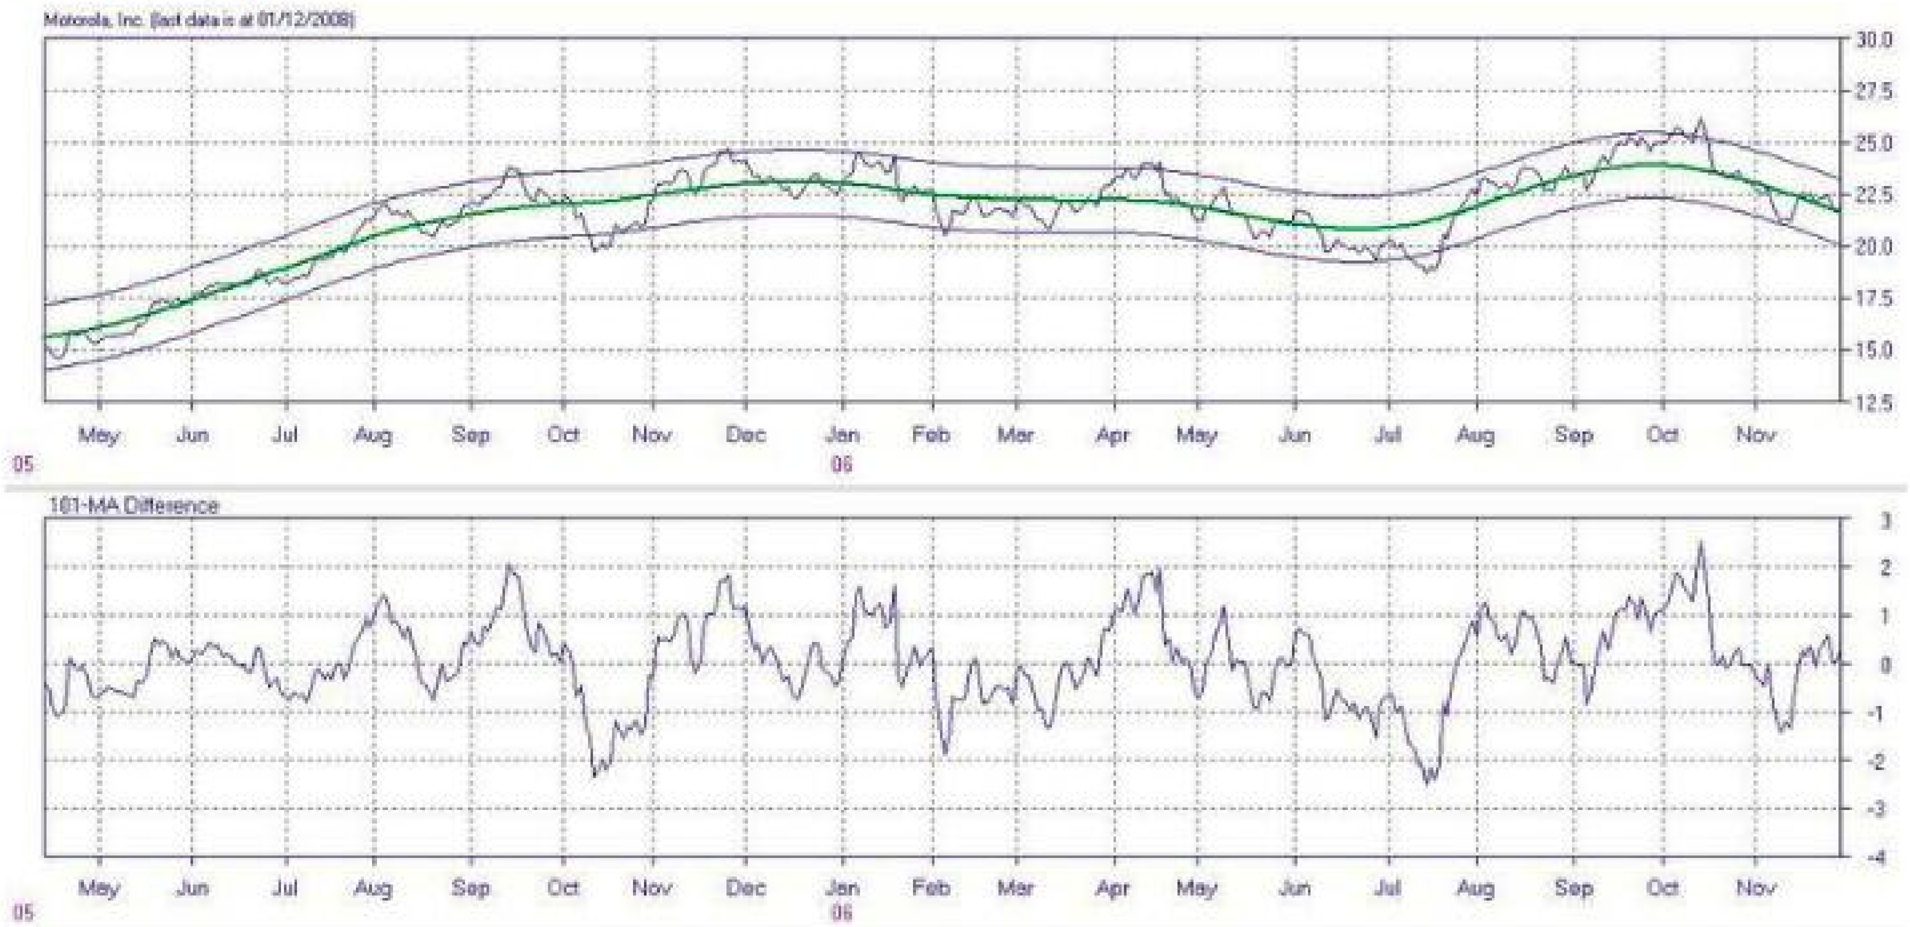

The correspondence between the movement within the channel in the upper panel and the movement in the average difference is quite clear. Points where the data reaches the channel boundaries are replicated by peaks and troughs in the lower trace.

PREDICTION OF FUTURE PRICE

What we have discussed so far enables us to predict a range in which the price will *move relative to the centred average*. Thus we are still faced with the fact that we do not know how the average has behaved over the half of its period which covers the period of time from the last calculated point to the present time and then on into the future. Unless we can make a good prediction of the path which the average (which represents the trend) has taken over this period of time, then any prediction of the price relative to this central trend will be of little value.

EXTRAPOLATING THE CENTRED AVERAGE

The dilemma facing us is exemplified by Figure 9.12. In this we show just three ways out of the many in which a particular centred average can be extrapolated into the future. This has been done by taking into account the curvature of the most recent section of the centred average and projecting it into the future. How large a section which is taken can have a profound effect on the extrapolation. Thus in the absence of any other evidence there is a large error associated with any method which simply depends on a visual estimate. However, there are several ways in which we can make progress. The first is through probability analysis, the second through channel analysis and the third is through cycle analysis. The most powerful technique for predicting future trends and price movement around these trends will be that in which all three of these are used. This approach is discussed in more detail in the final chapter of this book.

Figure 9.12 – There are many ways of extrapolating a centred average from its last calculated point to the present and into the future. Each of the three possibilities shown above can be justified.

Where channel analysis will be able to help us is in establishing a probable location for the boundaries. Since the average is in the exact middle of the channel, it follows that knowing the position of the boundaries will give us the position of the average and hence of the trend which it represents. We will be able to do this because of the fact that boundaries are places where the price movement has a high probability of reversing direction. Therefore, of necessity, these reversals take the form of a peak or trough. It is by a careful examination of the position of these peaks and troughs that we can estimate the position of the boundaries. This method of establishing trend direction is discussed in Chapters 11 and 12.

10. Trend Turning Points (I)

In the last chapter we illustrated the dilemma that we face when we use centred averages as proxies for trends. We have a gap of half of the period of the average during which we do not know what the average (trend) has been doing. Before we can begin to think about the future movement of the trend we have to examine ways in which we can make a decent estimate of its current position.

Figure 9.12 from the last chapter is presented again here as Figure 10.1 in order to focus on what is required. The average is known up to the point where it divides into three possible paths.

Figure 10.1 – This centred average could be extrapolated in many ways to bring it up to the present time and into the future. The three paths shown are just three examples, each of which can be justified.

These three are used for illustrative purposes – there will be a large number of possible paths, each of which can be justified by analysis of the price data itself; so how do we bring these possibilities down to just one? The answer of course is that we cannot end up with just one line which represents the path of the average up to the present time. It might be possible in some circumstances to be able to draw the most probable path as a widening range of values as we move forwards from the last calculated point. This is shown in Figure 10.2, where there are limits to its probable position. By this it is meant that there is a range of values (the lowest value and highest value) within which there is a certain probability (say 90%) that the price will lie.

Figure 10.2 – It might be possible to extrapolate the average and put 90% limits on its probable position.

If it is not possible to predict a probable path for the extrapolated average, then the least we would hope to achieve is to decide whether it has changed direction somewhere within the gap. In this case it is not necessary to decide at which exact point it changed direction, but only that it has changed direction.

Short-term Trends

The simplest measure of determining whether a trend (average) is currently standing at a higher or lower value than its last true calculated point is to obtain a prediction of its probable value at a future point.

This issue can be addressed by reference to the drop point, which was discussed in Chapter 8. A falling average will change direction if the new point being brought into the calculation is of greater value than the point being dropped. For a rising average, it will change direction if the new point being brought into the calculation is of lesser value than the point being dropped.

How can we estimate if a new point is likely to be of greater or less value than the point being dropped? The answer lies in the probability distribution of daily changes. We need to know:

1. the value of the drop point
2. the probability that the value of the new point lies above or below a certain value, this value being that required to offset the value being lost by the drop point.

Using these values we can extend the centred average incrementally from its last calculated point half a span back in time, one point at a time to the present.

TURNING POINT IN RISING AVERAGES

Taking a nine-point centred average as an example, the initial position is shown in Figure 10.3. For ease of understanding, the x-axis has been labelled with point numbers rather than dates. The vertical line at point 1353 is the point at which the data terminates with a value of $49.95. The smoother line is the nine-point centred average. This terminates at point 1349, which is half a span (four points) back.

Figure 10.3 – The centred nine-point moving average calculated with the last data point being at point 1353 (12 May 2008). The data to the right of the vertical line shows the subsequent price movement.

The issue now is to move this average on, one point at a time, to the present. In reality, as each new data point comes in, the average needs to be recalculated so that its position will move on one point at a time. Thus, once we know the value of point 1354, we can calculate the value of the average for point 1350. Once we know the value of point 1355, we can calculate the value of the average at point 1351, and so on. The value of the average at the present time, i.e. point 1351, will only be known once we have the value of the data at point 1355, which is four days into the future. In general, therefore, to bring an n-point average up to the present time, we need to estimate the value of the data points up to half a span (n is odd, so half a span is (1 - n)/2) into the future. In the case of this nine-point average, this means that we have to estimate the next four points into the future. The only way we can do this is to use the probabilities associated with the distribution of price changes as discussed in Chapter 4. Thus we will need the distribution of one-day, two-day, three-day and four-day changes and determine the appropriate probabilities for these various price changes.

Since in Figure 10.3 the average is rising, our interest is in seeing when or if it has changed direction between point 1349 and 1353. This is where the drop point becomes essential. When we compute a new value for the average, we will bring in the estimated value at point 1354, but will drop the point 10 days back, i.e. point 1345. Thus, in order for the average to start falling, the new value must be less than that at point 1345, i.e. less than $48.42. Since the value at point 1353 is $49.95, then this can only be achieved if the one-day change from point 1353 to 1354 is a fall of more than $1.53.

The distribution of two-day changes in American Express is shown in Figure 10.4. The probability of a fall of more than $1.53 is only 4.4%. Thus it is highly improbable that the average is falling at point 1350.

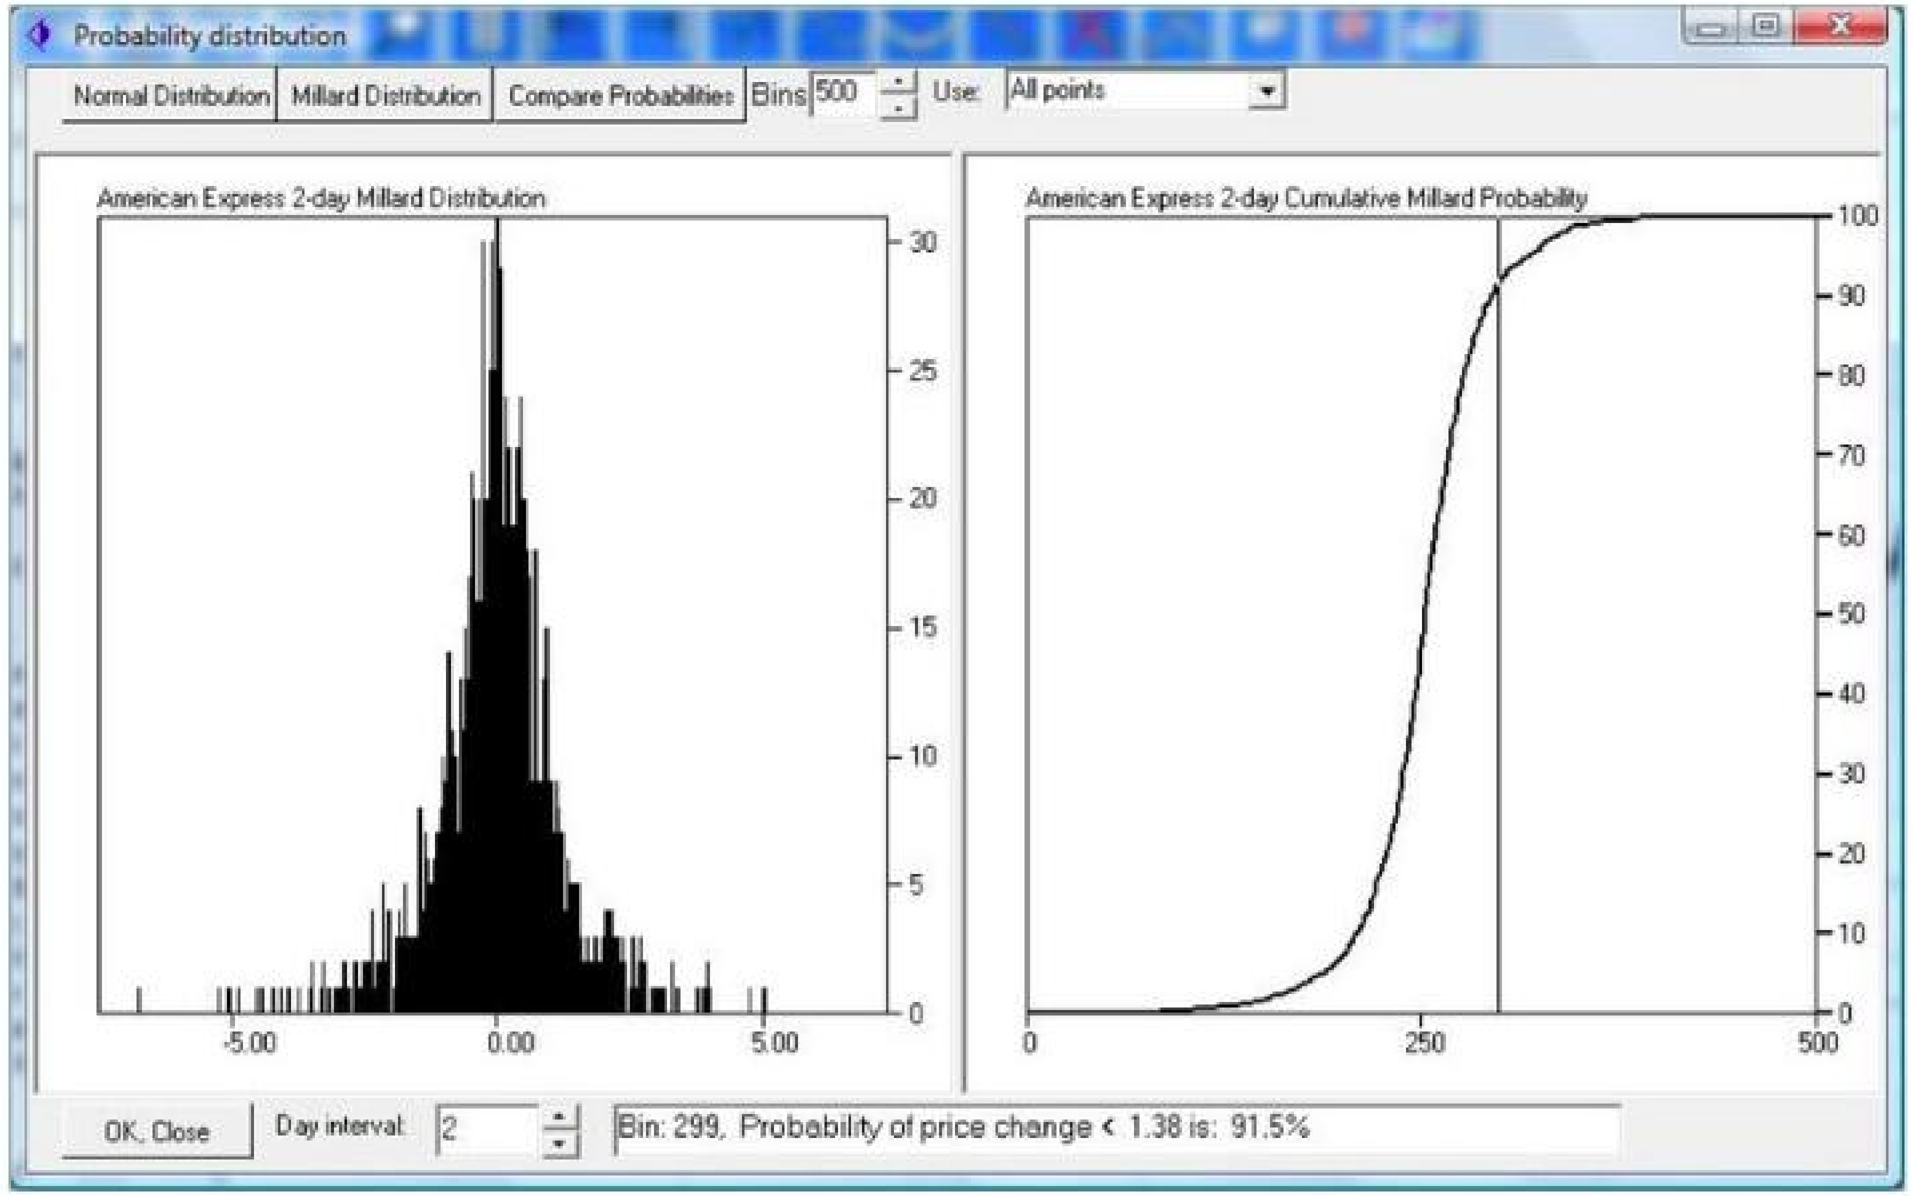

Figure 10.4 – The distribution of the American Express two-day price changes. The probability of a change of less than $1.38 can be seen to be 91.5%. Therefore the probability of a rise equal to or greater than this is 8.5%.

The next step would be to determine the distribution of two-day changes in order to estimate whether point 1355 will fall enough to cause the average also to fall. Then, a three-day distribution will give us this information for point 1356 and finally a four-day distribution will give us the information for point 1357, which will bring the average up to date.

The data obtained by this process is shown in Table 10.1.

The probabilities in Table 10.1 would lead one to believe that the average is likely to change direction by the time the estimated data for point 1357 is determined.

Table 10.1 – The probability of the price changes being sufficient to keep the nine-point

centred average rising.

The probabilities in the last column of Table 10.1 were obtained from the one-, two- and three-day distributions as discussed in Chapter 4. The two-day distribution is shown in Figure 10.4. This shows that the probability of the price change being greater than $1.38 is only 8.5%. Thus the probability of it being less than $1.38 is 100 - 8.5 = 91.5%.

From Table 10.1 it can be seen that the probability of the average still rising at point 1350 is over 97%, so that we do not expect it to change direction at that point. However, the probability of it still rising at point 1351 is now only 8.5%, so that we have a high expectation that the average will have turned down at that point. This means that the maximum point should occur at point 1350 on the previous day. In Figure 10.5 we show how the average actually moved when its values were calculated from data which terminated much later on.

Figure 10.5 – The actual path of the centred nine-day average in American Express. The average topped out at point 1350, 7 May 2008.

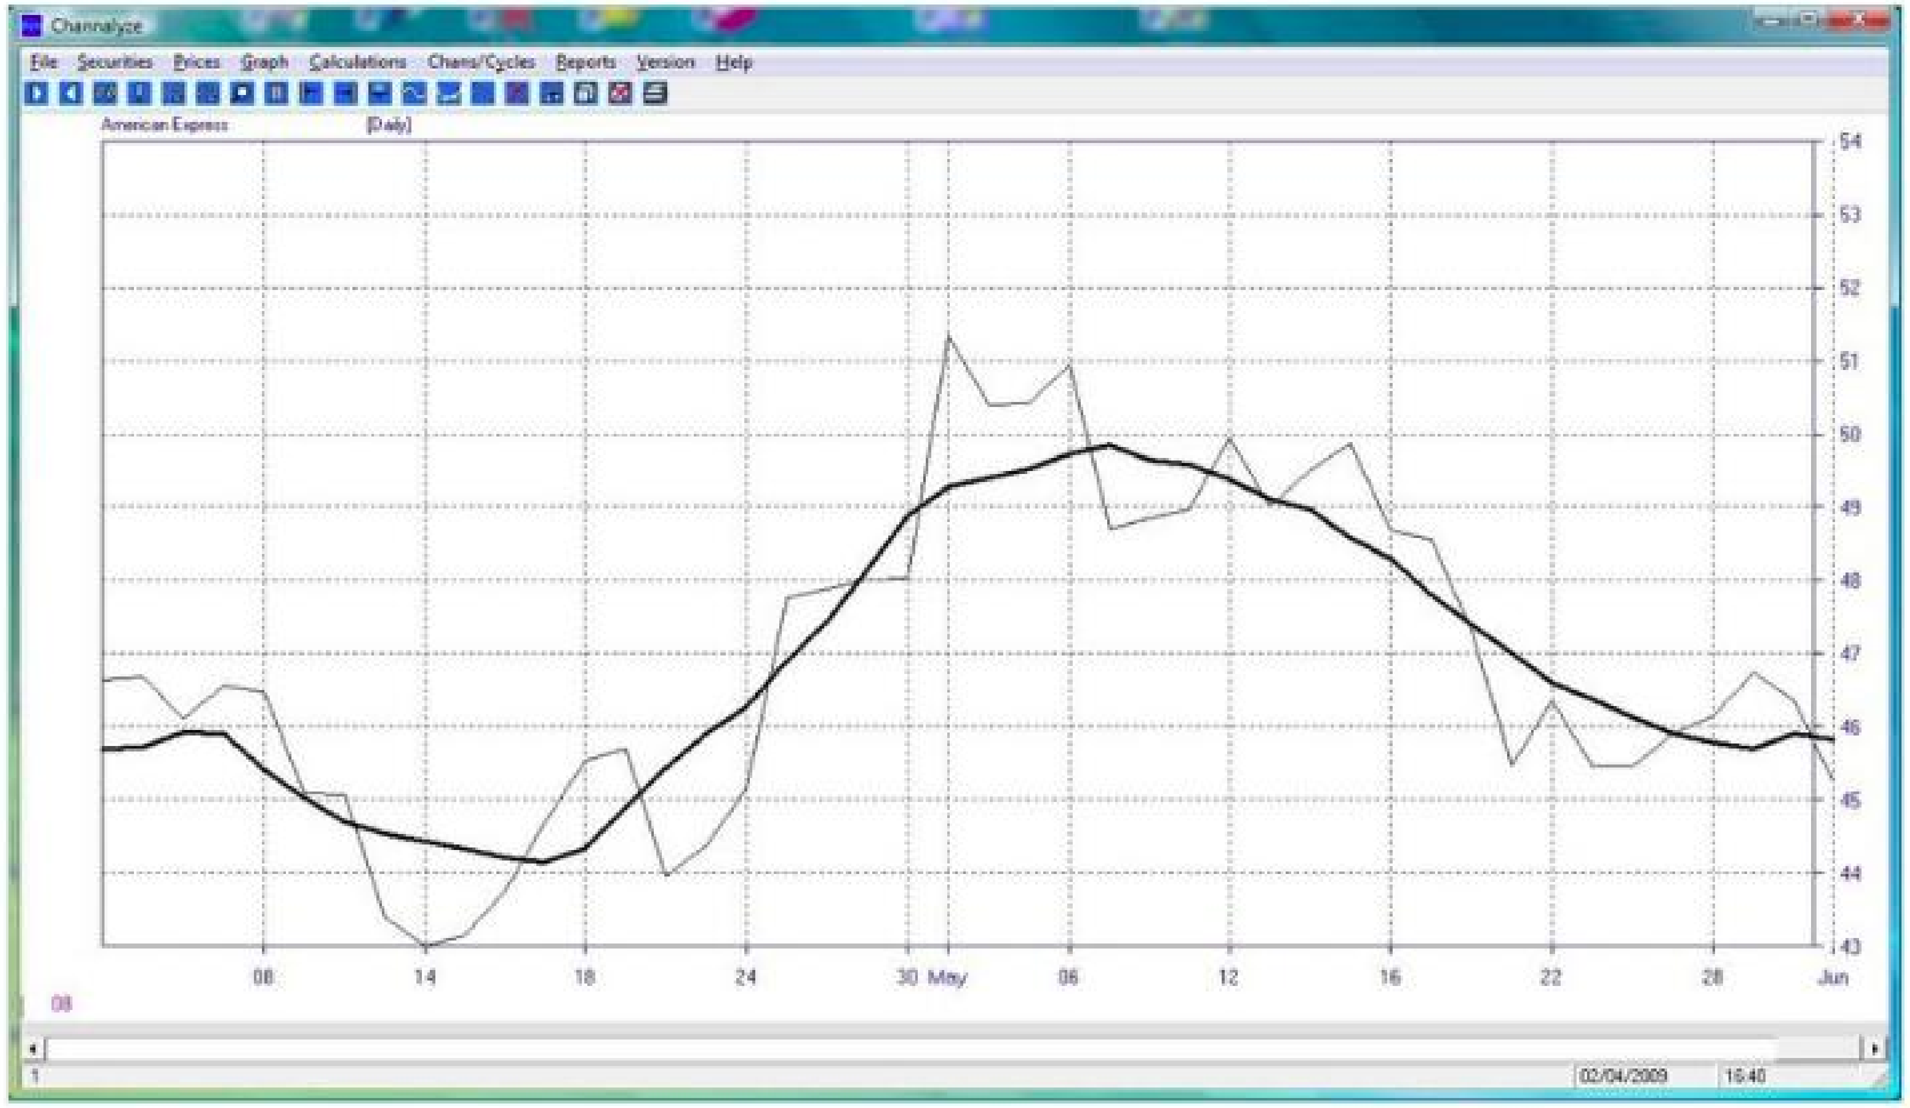

It can be seen that the average did indeed top out at point 1350 on 7 May 2008.

Thus it appears that we have a useful method of determining the probability that an average has changed direction. However, as discussed in Chapter 4, there will be a limit for the spans (periods) of the averages to which this method might be applied. For averages with longer periods, the Millard version of the Monte Carlo simulation should be employed (see Chapter 11).

TURNING POINT IN FALLING AVERAGES

We can now test this method on a longer period of average – 15 days – and also on a situation where we are trying to predict whether the average will change direction and start to rise. In Figure 10.6 we show a plot of the AT&T stock price with a 15-day centred average superimposed.

Figure 10.6 – The centred 15-point average calculated with the last data point 1307, 7 March 2008.

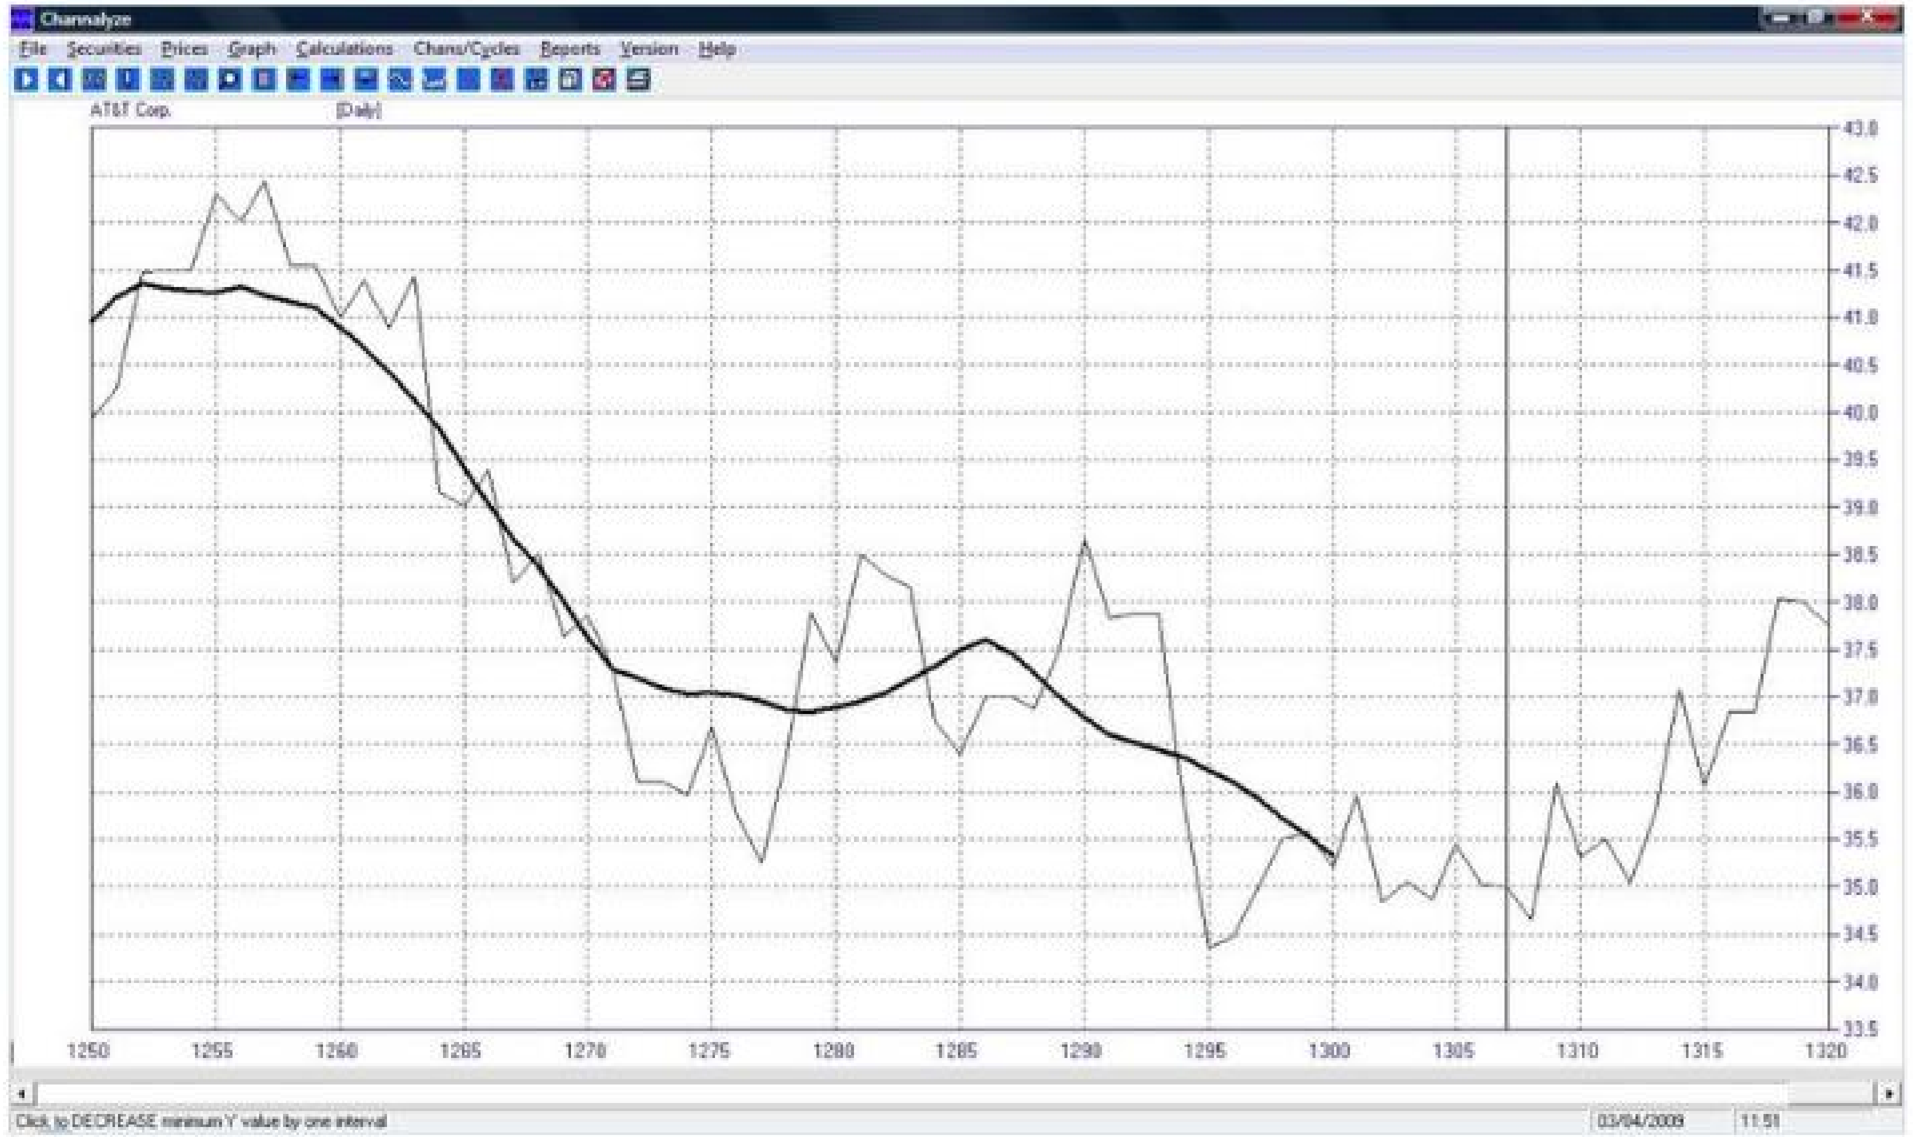

Since in Figure 10.6 the average is falling, our interest is in seeing when or if it has changed direction between point 1292 and 1307. When we compute a new value for the average, we will bring in the estimated value at point 1307, but will drop the point 15 days back, i.e. point 1292. Thus, in order for the average to start rising, the new value must be greater than that at point 1292, i.e. more than $37.88. Since the value at point 1307 is $35.01, then this can only be achieved if the one-day change from point 1307 to 1308 is a rise of more than $2.87.

We can follow the same procedure for this example as we did for the previous one. The various data points and probabilities are listed in Table 10.2.

From Table 10.2 it can be seen that the probability of the average still falling at point 1301 is over 99%, so that we do not expect it to change direction at that point. The same applies to point 1302, in which the probability of it still falling is over 94%. However, when we come to point 1303, the probability of it still falling is now only 13.25%, so that we have a high expectation that the average will have turned up at that point. This means that the minimum point should occur at point 1302.

Table 10.2 – The probability of the price changes being sufficient to keep the 15-point centred average falling.

The distribution of three-day changes for AT&T is shown in Figure 10.7. It can be seen that the probability of a change of less than -0.7 is 13.25%. This value is used in Table 10.2 because it is the closest to the required value of -0.65.

Figure 10.7 – The distribution of the AT&T three-day price changes. The probability of a change of less than $0.7 can be seen to be 13.25%. Therefore the probability of a rise in the average is 86.75%.

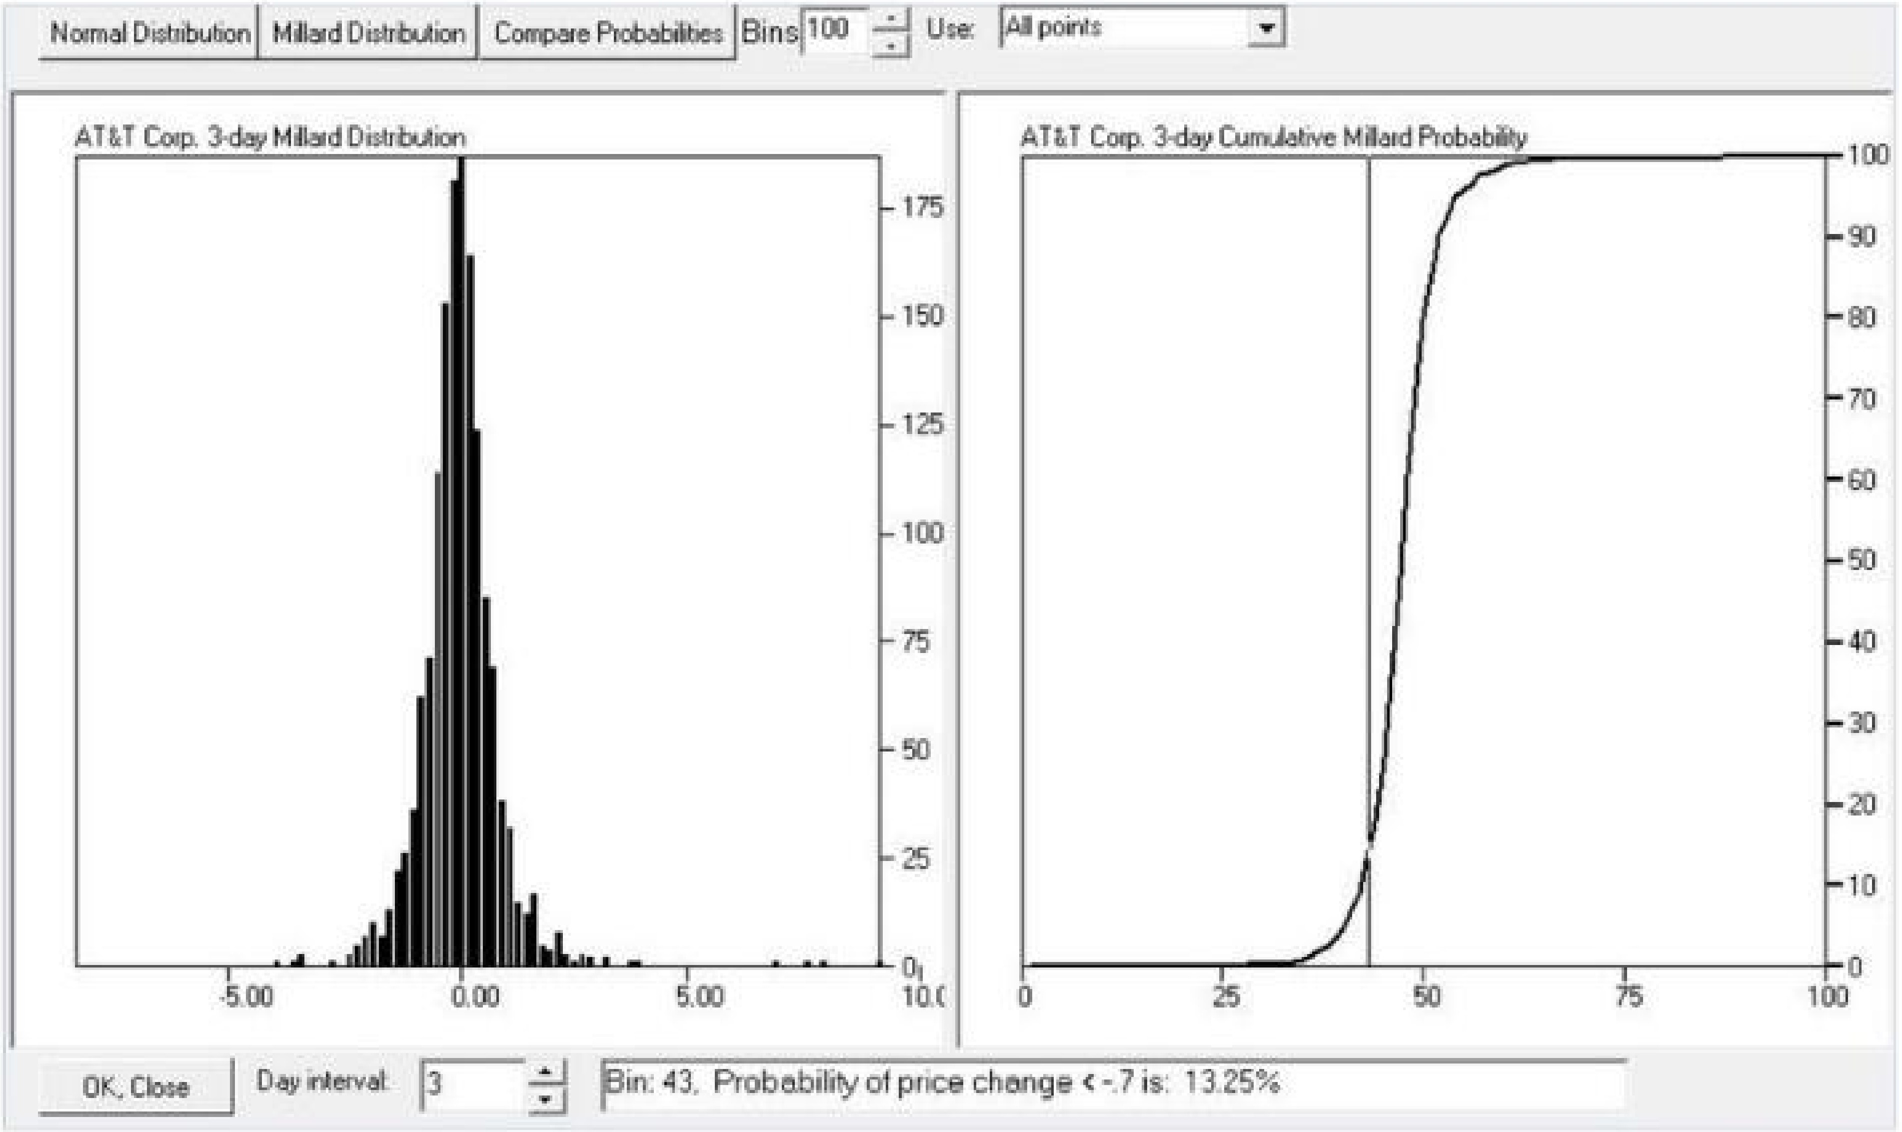

In Figure 10.8 we show how the average actually moved when its values were calculated from data which terminated much later on. The minimum occurred at point 1301 on 28 February 2008. We were only one day out in predicting the turning point of this average!

Figure 10.8 – The actual movement of the 15-point centred average over the period for which the prediction was being made.

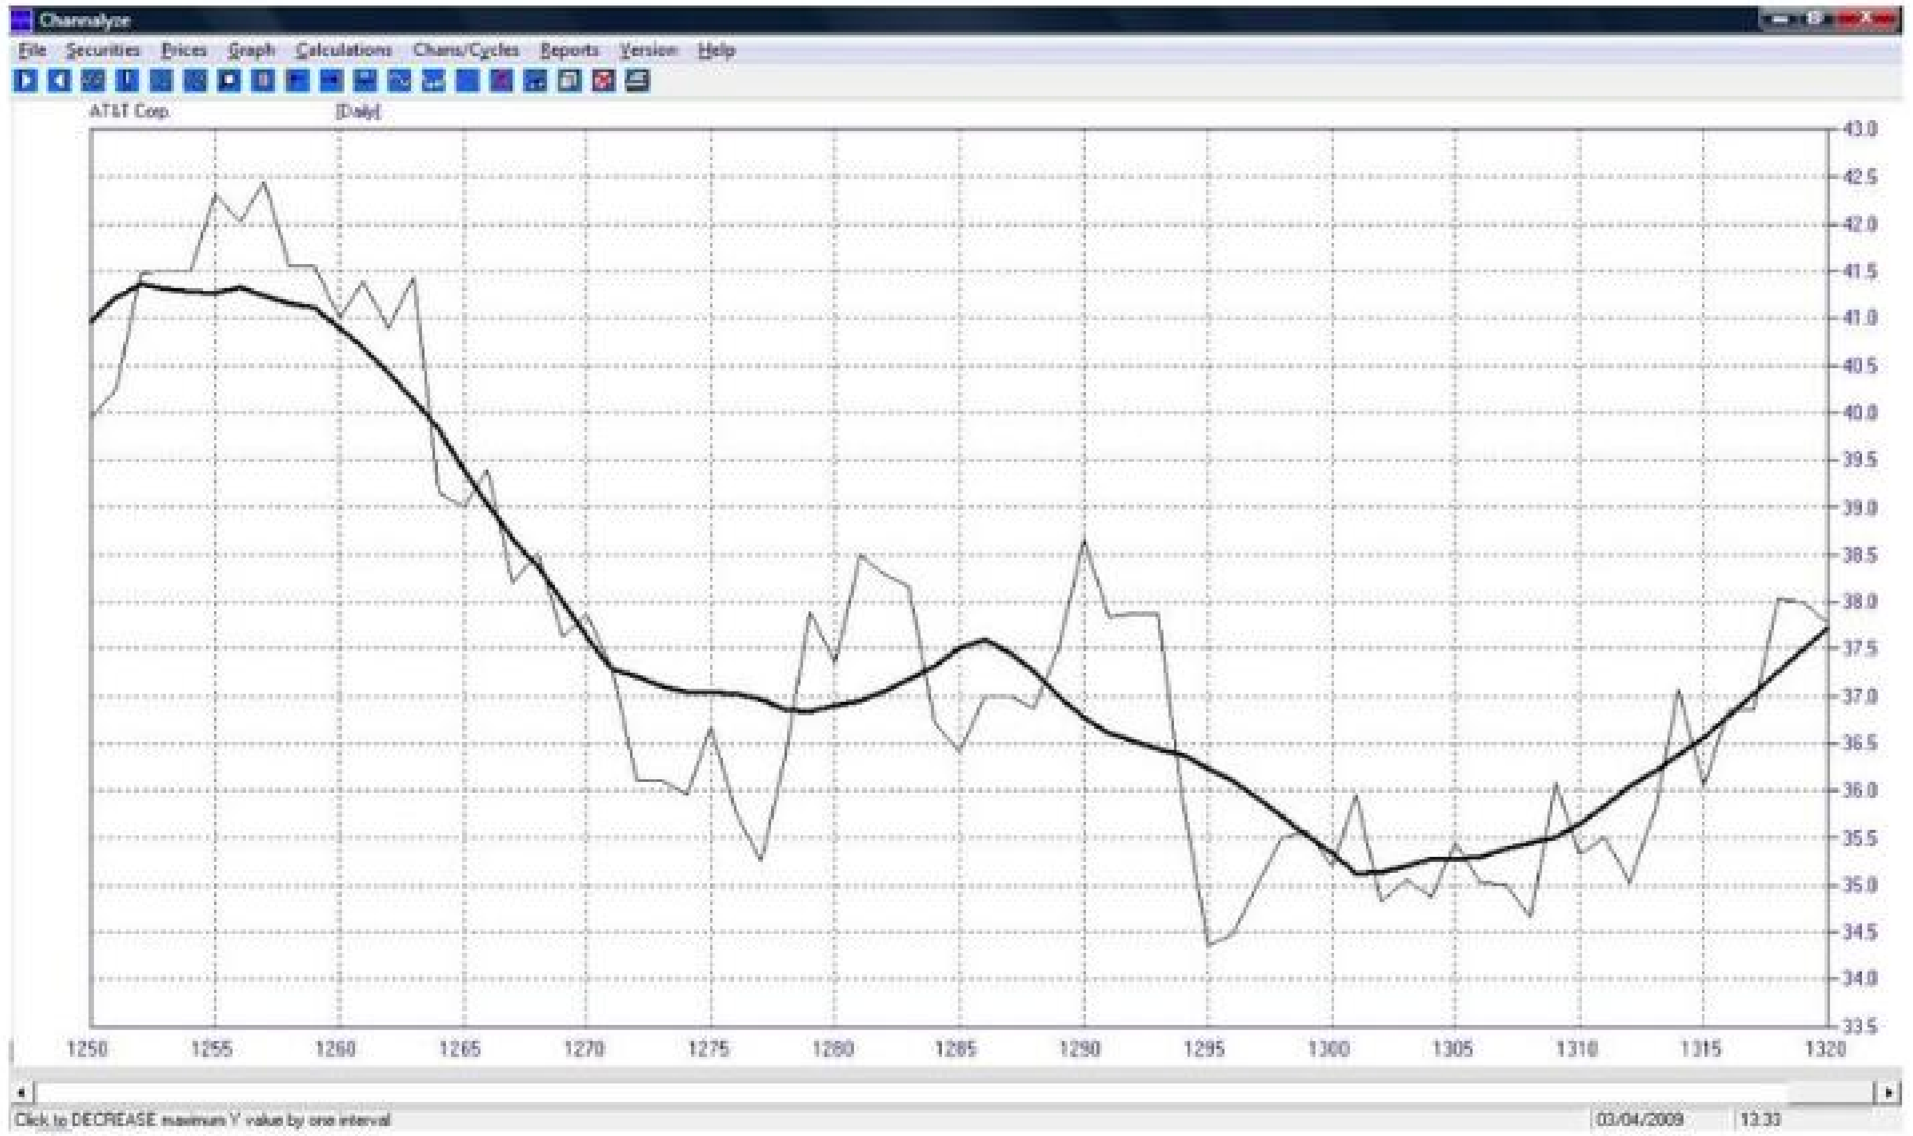

RISING AVERAGE CONTINUING TO RISE

As well as being able to determine that an average may have changed direction in the gap between the last calculated point and the present time, it is also useful to be able to determine that an average will not change direction.

Shown in Figure 10.9 is a section of the American Express stock price with the nine-point centred average superimposed. The average has been rising for some time, so the issue is whether it is still rising at the point corresponding to the last data point.

Figure 10.9 – A rising nine-point centred average in the American Express stock price. The issue is whether it is still rising at the current time. This point is at the position of

the vertical line.

Using the same procedure as was shown in Table 10.1 for a rising nine-point centred average in American Express, the data shown in Table 10.3 were obtained.

Table 10.3 – The probability of the price changes being sufficient to keep the nine-point centred average rising.

The last column shows that the probability is extremely high that the average will keep on rising. That it did indeed rise is shown in Figure 10.10.

Figure 10.10 – The nine-point centred average in the American Express stock price continues to rise as predicted by the probability calculations shown in Table 10.3.

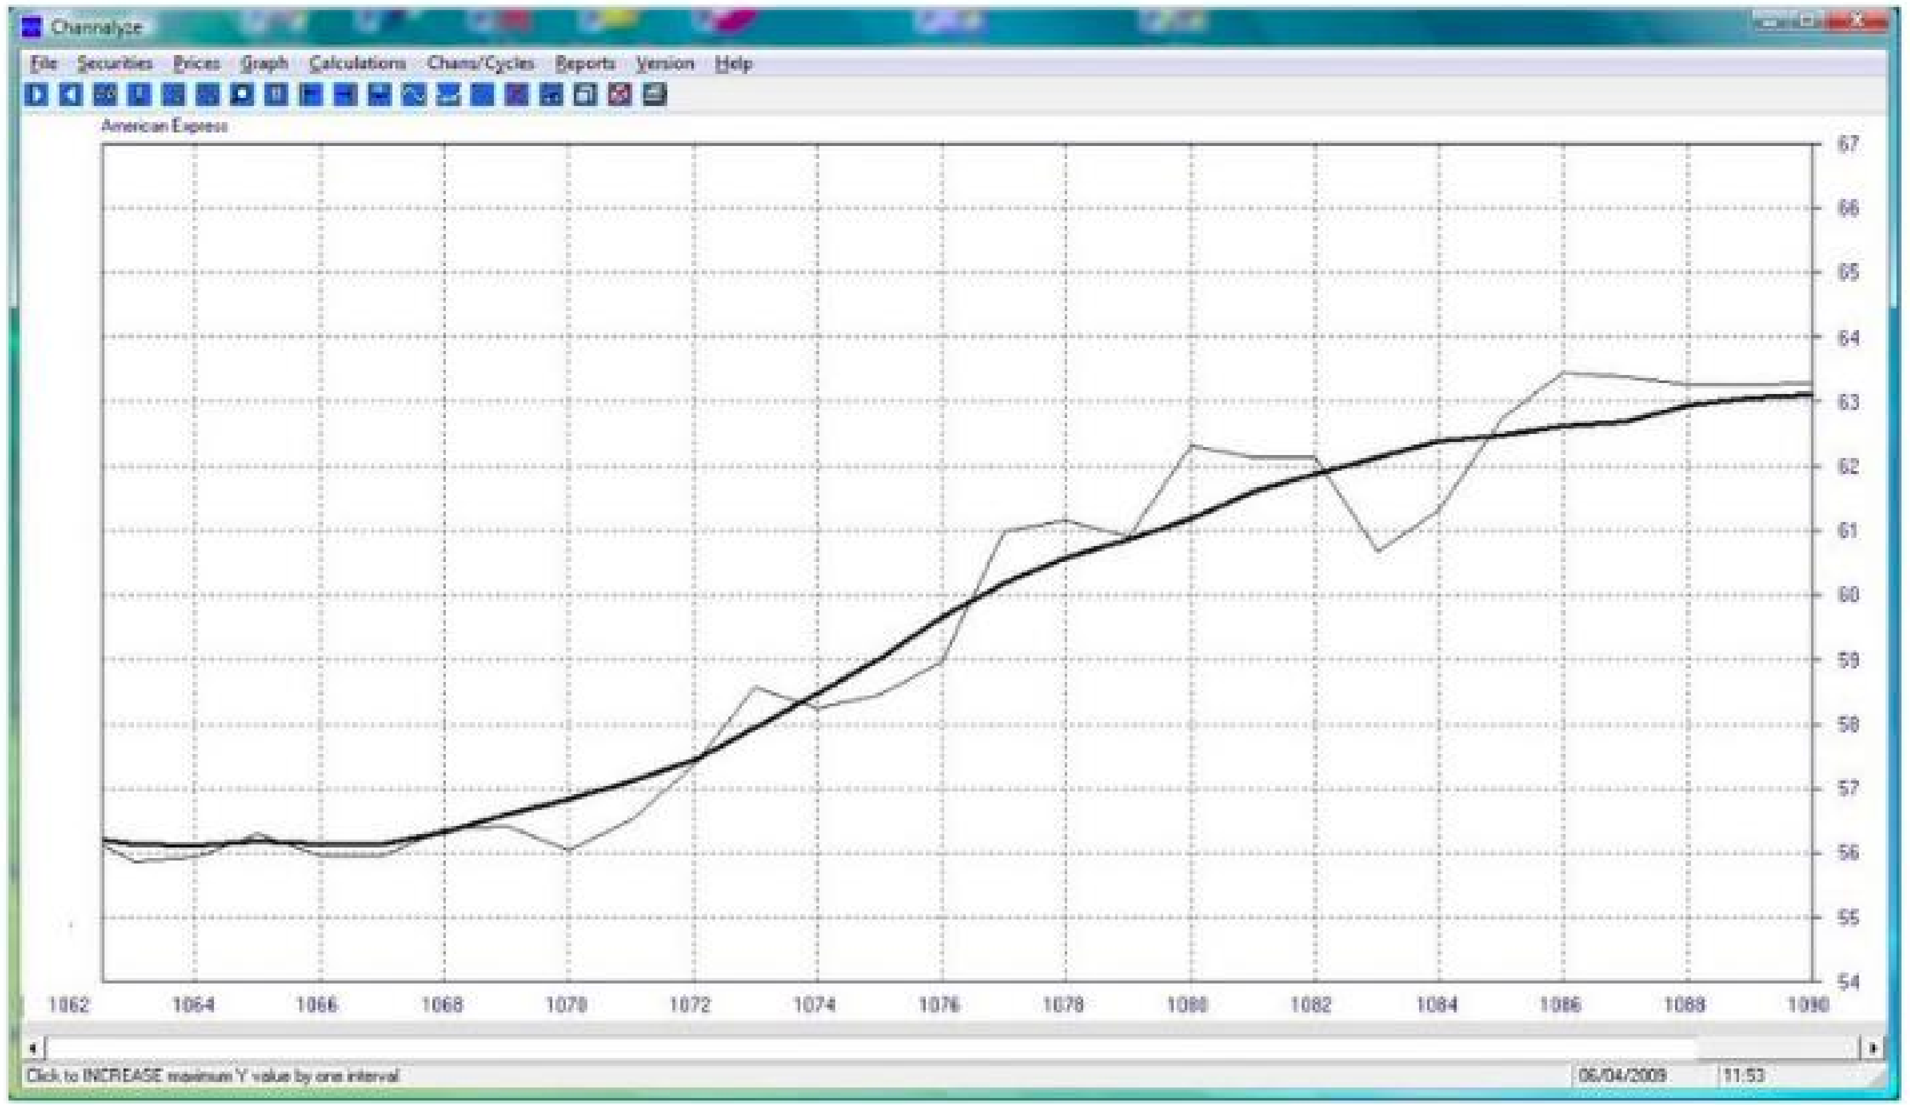

FALLING AVERAGE CONTINUING TO FALL

Just as it was possible to determine that the average was still rising, then so we should be able to determine the probability that a falling average is still falling. This time we can use the example of the Eastman Kodak stock price with a 15-point average superimposed. This is calculated at 2 June 2008 and is shown in Figure 10.11.

Figure 10.11 – A falling 15-point centred average in the Eastman Kodak stock price as calculated on 2 June 2008. The issue is whether it is still falling at the current time. This point is at the position of the vertical line.

The calculation in Table 10.4 brings us right up to the last data point at 1376 and it can be seen that the high percentage probabilities would suggest that the centred average is still falling, although the probabilities are decreasing slowly. This is seen to be the case from the plot of this average calculated further on in time in Figure 10.12. The average did not stop falling for another eight days, finally bottoming out on 12 June 2008.

Table 10.4 – The probability of the price changes being sufficient to keep the 15-point centred average falling. Point 1368 on 2 June 2008 is the last real data point.

Figure 10.12 – The 15-point centred average in the Eastman Kodak stock price continues to fall as predicted by Table 10.4.

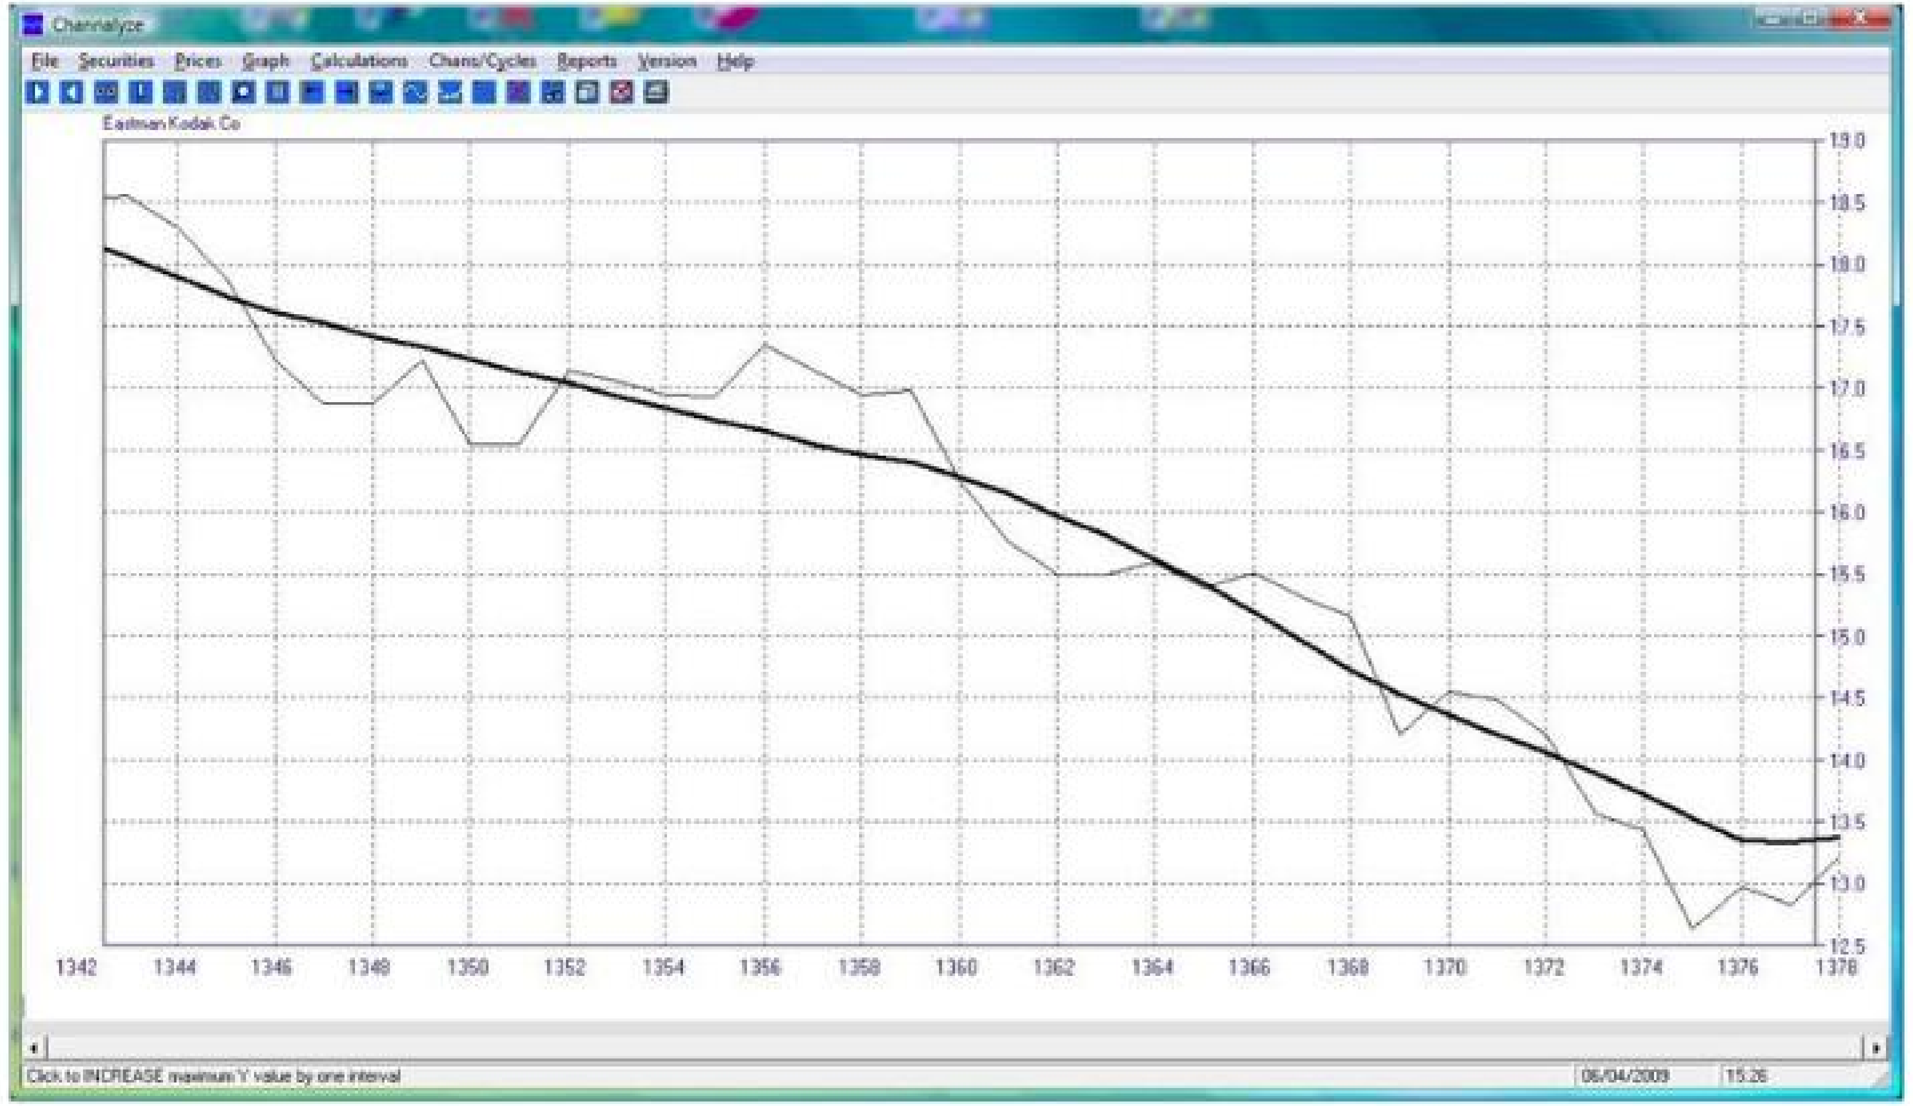

Quite clearly, we have a very powerful method for determining how a centred average, and thus the trend which it represents, has behaved from the last calculated point up to the present time. Without this knowledge many traders will take a position assuming a trend is rising or falling when in fact it has already changed direction. This change of direction will not become obvious until a few days into the future, when more data has arrived to update the position of the centred average.

PREDICTING TURNING POINTS IN THE NEAR FUTURE

The examples discussed so far have been those in which we investigated whether the average had changed direction in the gap between the last true value of the average and the latest data point. It was seen quite clearly that this method gave an excellent indication of changes of direction which occurred and also indicated where a change of direction did not occur.

It will have been noted that as far as the known drop points are concerned, we only needed to use enough to cover half a span of the average, which would bring us up to the present time. Since there are n drop points, where n is the span of the average, then we have enough extra drop points to be able to predict the path of the average up to half a span into the future. The question is, therefore, are we going to be as successful in these predictions?

In Table 10.5 we see the calculation of the probabilities of the average continuing to fall. There is no suggestion from these figures that the average will change direction. However, it was determined later from new data arriving that the average bottomed out at point 1376.

Table 10.5 – The probabilities of the 15-point centred average of Eastman Kodak stock continuing to fall out over the period of half a span into the future from the last data point at point 1372 on 6 June.

In order to check that this was not just a problem with particular values of the last data point and the drop points, the exercise was repeated from the next data point at 1373 (see Table 10.6), with a value of $13.57. Even without calculating the probabilities, it can be seen that the values labelled 'Must be<' are such that the probabilities will be high that the values will be less than these. As can be seen from the final column, there is no indication that the average changed direction on or about point 1377.

Table 10.6 – The probabilities of the 15-point centred average of Eastman Kodak stock continuing to fall out over the period of half a span into the future from the last data point at point 1373 on 9 June.

Finally, in Table 10.7 we see the data estimated where we have made the last data point to be that at 1384, 13 June. It was not until we reached point 1388 that the probability fell below 50% that the average was still continuing downwards, i.e. a reasonable probability that the average had now bottomed out.

Table 10.7 – The probabilities of the 15-point centred average of Eastman Kodak stock continuing to fall out over the period of half a span into the future from the last data point at point 1377 on 13 June.

The exact path of the average over this period of time is shown in Figure 10.13. It can now be seen that the reason for the change in direction of the average was the large increase in price between point 1383, value $12.34, and point 1385, value $15.12 – a change of $2.78 over two days. The probability of this happening is less than 1% and hence it would not be predicted on the grounds of probability.

Figure 10.13 – The turning point in the 15-point centred average in the Eastman Kodak stock price was driven by the large rise in the stock price ($2.78) between points 1383 (23 June) and 1385 (25 June). The estimated probability of such a strong rise is less than 1%.

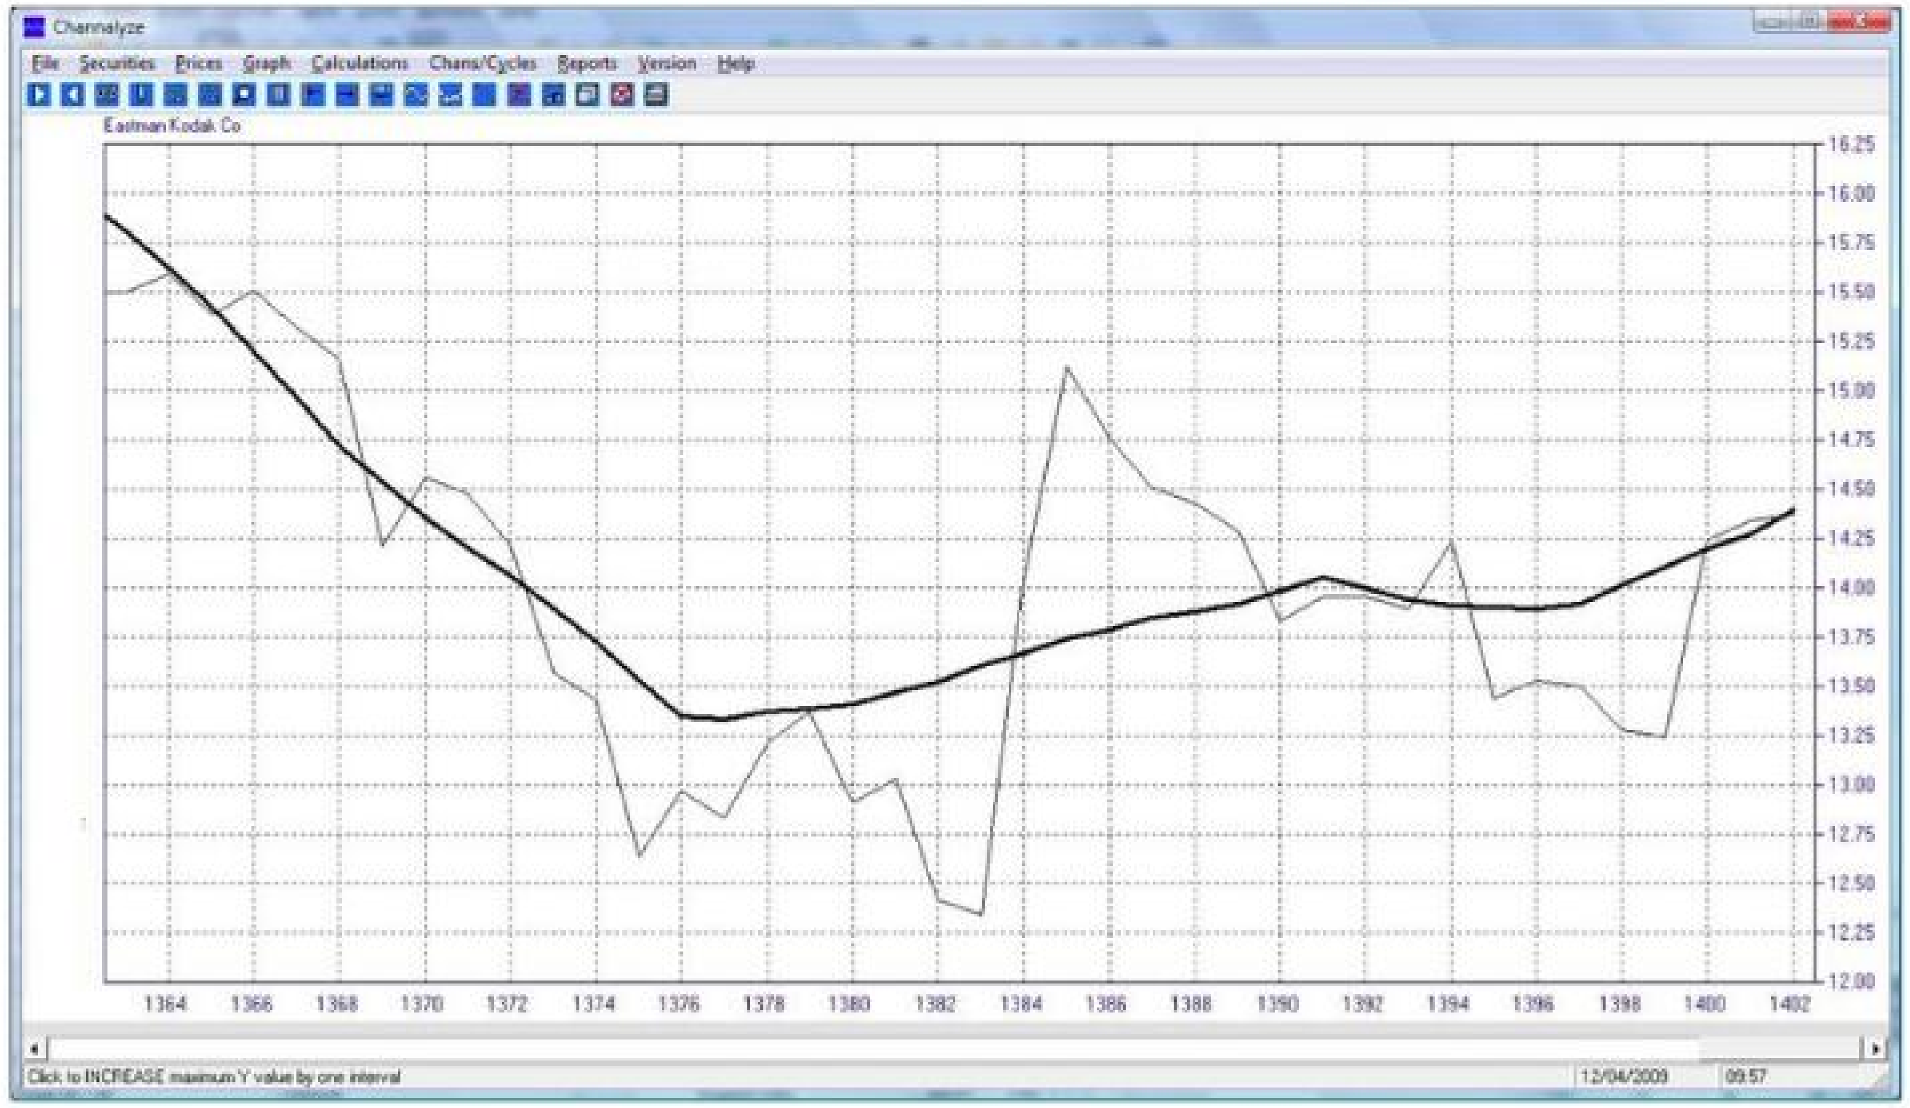

In conclusion, we have shown that the use of probability estimation is a very valuable tool in deciding whether an average has changed direction somewhere between the last calculated average point half a span in the past and the latest data point. It is not very useful in investigating if the average will change direction in the very near future, i.e. up to half a span into the future from the last data point.

11. Trend Turning Points (II)

In this chapter we will be examining more closely the properties of channel boundaries. The channel boundaries are of course constructed by drawing replicates of the centred average above and below the position of the average, and equidistant from it. Thus the channel is of constant depth throughout. This can be done by a computer program or by eye with a pencil and paper on the chart. The depth of the channel depends solely upon the level required for the probability that a new data point will lie within the channel. A good general value for this is 97.5%.

It is interesting that while the centred average is used to produce the boundaries, we can then use the boundaries which have been extrapolated to estimate the path the average might have taken across the gap. It can be seen that a few iterations of this process might be needed before a good estimate of the average can be obtained. This estimate will be particularly valuable when it comes to deciding whether the average, and hence the channel, has changed direction in this gap, or even if it might change direction in the very near future.

Estimation of Channel Direction

Shown in Figure 11.1 is the 101-point channel of the Brunswick Corporation stock price. A value of 96.5% has been chosen for the number of points allowed to be outside the channel. The vertical lines show the gap between the last true calculated point and the latest data point. The subsequent movement of the stock price is shown so that we can discuss how the analysis can proceed.

As the channel approaches the position of the last true point, it can be seen that in this case it is almost a straight line. This makes it easy to extrapolate, since the general rule in the first pass of extrapolation of channels is to maintain the previous rate of curvature. However, it is obvious that the extrapolation is not correct, since the stock price violates the upper boundary almost immediately at the position of the latest price. Since only 3.5% of points lie outside of the boundaries, then the upper boundary (and hence the lower boundary to preserve constant depth) will have to be raised at this point in order to preserve this allowable penetration.

Figure 11.1 – The 101-point channel of the Brunswick Corporation. The left-hand vertical line is the position of the last true calculated point of the average from which the channel is derived. The right-hand vertical line is the latest data point (4 September 2006). The subsequent movement of the stock price (not used in the calculation) is shown in order to clarify the discussion.

Of course, the boundary cannot simply be raised and its appearance as a straight line maintained at the same time. If this was attempted, then the trough which is shown within the gap would then violate the allowable level of penetration of the lower boundary. The only solution is to begin to bend the channel so that its rate of fall will begin to lessen.

The general approach therefore is to check boundaries for penetration by the stock price and make the appropriate adjustment while checking that this does not cause a penetration of the opposite boundary. This will inevitably mean that the rate of curvature will be changed. The position of the extrapolated boundary when new data are available for 12 September for the Brunswick Corporation is shown in Figure 11.2. This simply takes the curve of the actual channel and extrapolates it to the current last data point. Note that the additional data points have now introduced a slight bend into the channel even before any further adjustment is carried out.

Figure 11.2 – By 12 September the new data has caused the channel to curve slightly. The channel has not yet been adjusted to account for the data point at the latest position being above the upper boundary.

Shown in Figure 11.3 is the effect of a minor adjustment in the curvature of the extrapolated portion of the channel. The curvature has now increased and the channel appears to be about to bottom out.

Figure 11.3 – The channel has now been adjusted to account for the data point ($30.54) at the latest position (12 September) being above the upper boundary. Note the curvature has now increased and there is now the possibility that the channel is about to bottom out.

A few days later (20 September) the bend in the channel had to be increased by an adjustment in order to accommodate the fact that the new data point ($31.4) was again above the upper boundary. The result of adjusting the boundaries is shown in Figure 11.4. Quite clearly, therefore, the adjustment which was needed to preserve the low percentage of points allowed outside of the boundaries has shown that the channel has bottomed out and is now on its way upwards. The estimated low point of the channel from this analysis is at point 912 on 1 September.

This estimated low point is about halfway across the gap between the last calculated point and the last data point. It is interesting to check with later data how accurate this estimation was.

Figure 11.4 – The channel has now been adjusted to account for the data point ($31.4) at the latest position (20 September) being above the upper boundary. Note the increasing curvature now shows that the channel has bottomed out and is expected to rise.

It can be seen from Figure 11.5 that the actual turning point was at point 903, on 21 August. When it is considered that the gap between the last calculated point and the last data point is 50 days, then to estimate the turning point of the channel (and hence the average) to within nine days of the correct value is a vindication of this method of analysing the relationship between the data and the channel.

Figure 11.5 – The channel drawn once more data was available. The actual bottom was at point 903 (21 August).

In this type of visual analysis, what we are doing is unconsciously using the properties and relationship of the cycles of wavelength less than the span of the average to modify the position of the cycles of longer wavelength.

To be successful, it is necessary that these cycles of lesser wavelengths should be prominent across the gap that is to be extrapolated. Without this, there will be no features to help in establishing the upper and lower boundaries. The centre section of Figure 11.5 is a case in point. This section is expanded in Figure 11.6. There are no obvious features to help to determine the positions of the upper or lower boundaries. In such a case the only conclusion is that the channel is not about to change direction.

Figure 11.6 – The extrapolated section of the channel from Figure 11.5 shows no features that will help to determine whether the channel has changed direction. It must be assumed that it has not.

The relationship between the channel we have been discussing and shorter wavelength cycles is shown in Figure 11.7.

Here we have plotted the nominal 15-day cycle. As expected, there is a correlation between those places in the chart where the cycle is prominent, i.e. its amplitude is above the typical value, and those points in the channel where the price touches the boundaries. In order to highlight these touching points, the channel depth has been decreased to contain 92% of the data points instead of 97.5%.

Figure 11.7 – The lower panel shows the nominal cycle of wavelength 15 days in a section of the Brunswick Corporation stock price. It can be seen that those points where the cycle is of higher amplitude are the points in the channel where the stock price approaches the appropriate boundary. The channel depth has been decreased in order to accentuate the channel touching points.

REDUCING THE CHANNEL DEPTH

If there are some fairly prominent features that do not reach a channel boundary, then there is the option of reducing the constant depth of the channel. This is more likely to be required when a large span is used for the underlying average. Thus shown in Figure 11.8 is the channel derived from the 301-point average in the British American Tobacco stock price.

Figure 11.8 – The 301-day channel in British American Tobacco. There are very few touching points in the long upward run until the end of 2007.

The way in which the channel topped out in May 2008 is of interest since, as shown in Figure 11.9, the initial extrapolation of the channel by continuing its rate of curvature seems to fit in with the oscillations of the price data.

Figure 11.9 – The 301-day channel in British American Tobacco. The channel depth has been reduced to allow 10% of the points to lie outside of the boundaries. As a result there are now many more touching points.

In order to see this more clearly, this section is enlarged in Figure 11.10. Here the channel depth has been set back to its initial value where 3.5% of the data points are allowed to be outside the boundaries.

Figure 11.10 – The topping-out section of the 301-day channel in British American Tobacco. The channel depth has been reset to allow just 3.5% of the points to lie outside of the boundaries.

The interesting point about this chart is that there are now five areas (not necessarily a sharp, single peak or trough) where the price either touches or just penetrates the boundary. Comparison with Figure 11.9 shows that the reason is the increased amplitude of shorter-term cycles in this section. Two questions arise from this. Firstly, what is the particular short-term cycle or band of cycles which are now so prominent? Secondly, how far back in time can we move the latest data point and arrive at the fact that the channel is now topping out?

THE SHORT-TERM CYCLES IN BRITISH AMERICAN TOBACCO

A rough approximation of the wavelength of the more recent short-term cycles can be obtained by inspection of the peaks and troughs in the extrapolated section of the channel in Figure 11.10. The peaks are about 100 days apart, as are the troughs.

In order to check this, the nominal 101-day cycle is plotted in the lower panel in Figure 11.11. Now the relationship can be seen quite clearly. The amplitude increases from 38p to 90p over this period. As might be expected, the peaks and troughs in the cycle are in line with the peaks and troughs in the data which approach the channel boundaries.

Figure 11.11 – The topping-out section of the 301-day channel in British American Tobacco. The nominal 101-day cycle is shown in the lower panel. The time axis is labelled with numbers so that the wavelength of the cycles can be checked. See text for explanation of vertical lines.

The vertical line in the lower panel of Figure 11.11 shows the position of the last true calculated point of the cycle. The second vertical line in the lower panel represents the latest data point. The vertical line in the upper panel is the start of the extrapolation of the channel. The cycle has been extrapolated into the future.

Predicting whether the turning point has occurred

Figure 11.12 shows the position at 28 May 2008 at point 1365. The initial extrapolation of the upper boundary seems to be perfectly correct, since the peaks at point 1341 (24 April) and point 1352 (9 May) with values of 2023p and 2027p are just touching the upper boundary. There is no reason to expect the channel to do anything other than continue to rise.

Figure 11.12 – The channel extrapolation at 28 May 2008 at point 1365. There is no reason to believe that the channel is topping out.

The latest data point at 1365 (28 May) has a value of 1889p and this is just below the lower boundary. It is only slightly below the boundary and so this level of penetration is acceptable.

The position at 18 June is shown in Figure 11.13. The latest price is now well below the lower boundary and there is also a sign that the rate of climb of the channel is decreasing. Since the price is so far below the lower boundary, the channel will have to be adjusted to reduce the amount of penetration.

Figure 11.13 – The channel extrapolation at 18 June 2008 at point 1380. There is now just a suggestion that the rate of climb of the channel is decreasing. The value of the last data point on 18 June is 1831p, which is well below the lower boundary.

The shape of the channel after an adjustment is shown in Figure 11.14. The adjustment which has been made brings the lower boundary down to the level of the latest point at 1831p. This means that the two peaks at 1341 and 1352 mentioned earlier now penetrate the upper boundary by a considerable amount. However, experience of channel turning points shows that there is almost always an overshoot of boundaries in such cases. Thus in the current circumstances the adjusted position of the channel is acceptable. Therefore we can now see quite clearly that the channel is now topping out.

Figure 11.14 – The channel extrapolation at 18 June 2008 at point 1380 has now been adjusted by increasing the curvature to accommodate the low value (1831p) of the latest data point.

If we move on to point 1399 (15 July), where the value has fallen to 1693p, then we get the position shown in Figure 11.15. The penetration of the upper boundary is now quite large, so an adjustment is needed to attempt to reduce this. The adjusted channel is shown in Figure 11.16.

Figure 11.15 – The channel extrapolation at 15 July 2008 at point 1399. The channel has not been adjusted, but the fact that it has topped out is obvious.

Figure 11.16 – The channel extrapolation at 15 July 2008 at point 1399. The channel has now been adjusted in order to try to reduce the penetration of the upper boundaries.

If we now look at the channel from the point of view of exactly where we have found the topping point to be, then this point is about halfway across the gap, with its maximum at around point 1308 on 10 March 2008. Thus we have been able to establish that the channel, and hence the average upon which it is based, topped out some four months in the past.

That this estimate of the position of the top of the channel is correct can be seen from a plot of the channel once sufficient data had arrived for its calculation. This is shown in Figure 11.17. The actual channel top occurred at point 1326 on 3 April 2008. Although this is some 18 days later than predicted by the channel analysis, it should be pointed out that since we are using a 301-point average, the gap is half a span, i.e. 150 days wide. An error of 18 days in 150 days is more than acceptable in warning that the current direction of the long-term trend is now downwards.

Figure 11.17 – A later calculation enables the exact top to be determined at point 1326 (3 April 2008) as shown by the arrow.

Predicting if the turning point will occur in the near future

It would be even more useful if we could use this method of analysis to determine whether the channel was about to change direction in the near future. However, this is quite difficult for channels based on a high value for the period of the average, such as the 301 days that we have used in the previous examples. It will not work if we have a gradual increase in the amplitude of the shorter-term cycles contained within the channel. This is because in such a case there is probably a prominent cycle about to develop over the near future that is not yet part of the calculation and so will have a large effect once it is included in the near future.

Thus the most likely scenario to give an indication of a change of direction in the near future is where successive peaks are aligned with the upper boundary. An example can be seen in Figure 11.18 for the H.J. Heinz Company. The last data point is at 20 February 2007 with a value of $47.69.

Figure 11.18 – At this time in the H.J. Heinz Company stock price (point 1034 (20 February 2007, $47.69)) it is predicted that the channel will top out in the near future.

The indication is that the channel is showing a decreasing rate of climb and should top out somewhere in the future. Only a slight adjustment is needed to align the upper boundary with the recent peaks, as shown in Figure 11.19.

Figure 11.19 – A small adjustment has been made to bring the upper boundary to the level of the recent peaks.

The indication is still that the channel will top out in the near future and the estimation by extrapolating the current amount of curvature is that this will happen by the time the right hand end of the chart is reached, around point 1100.

As shown in Figure 11.20, the actual top of the channel occurred at the point indicated by the arrow, point 1106 on 31 May 2008. Our estimation of topping out, although rough and ready since it is based on extrapolating the curve of the channel into the future, was therefore very helpful. Since the price was at the upper boundary, the decision would be taken to sell, since the probability is high of a rebound back towards the centre of the channel.

Figure 11.20 – A later calculation enables the exact top to be determined, point 1106 (31 May 2008) as shown by the arrow.

Rising or falling?

An interesting chart is that of the British pound vs US dollar, as shown in Figure 11.21, with the latest data point being at point 770 on 12 February 2006. It is interesting because, unlike the previous examples in which the short-term cycles contained within the channel were either of fairly constant amplitude or increasing in amplitude, it can be seen in Figure 11.21 that the short-term cycles have diminished greatly. Even so, the indication is quite obvious that the channel has already bottomed out, with a turning point at about point 718 on 1 December. Notice the long gap from May 2005 until the latest point where there is no close approach to the upper boundary. It is the sequence of troughs from July 2005 that have defined the lower boundary and, hence, by virtue of the constant channel depth, also the upper boundary.

Figure 11.21 – The 301-day channel in a chart of the British pound/US dollar with the last data point at point 770 on 12 February 2006. The indication is that the channel bottomed out around point 718 on 1 December 2005.

A few weeks later, the newly calculated channel appeared as shown in Figure 11.22. Now the bend has disappeared and the indication is for a falling channel. A close inspection shows that the lower boundary has now moved away from the succession of troughs that had defined it on the earlier occasion. It can be seen that the channel can be adjusted so that these troughs are at the lower boundary without causing a difficulty with the peaks in this region.

Figure 11.22 – The 301-day channel in the chart of the British pound/US dollar with the last data point at point 794 on 17 March 2006. The extrapolated channel has now straightened out with no indication of a turn having occurred at around point 718 (1 December 2005).

The adjustment which causes the channel to bend upwards is shown in Figure 11.23.

Figure 11.23 – The 301-day channel in the chart of the British pound/US dollar has now been adjusted to bring the recent troughs to the lower boundary.

It could also be argued, if not quite as convincingly, that the channel could be adjusted downwards, so that the bunch of small peaks around point 750 could sit on the upper boundary. This is shown in Figure 11.24. However, the fact that the data itself is trending sideways would perhaps suggest that the fall in the channel is becoming less steep.

Figure 11.24 – The 301-day channel in the chart of the British pound/US dollar has now been adjusted to bring the recent peaks to the upper boundary.

This ambiguous situation is one in which no decision can yet be taken about the future path of the 301-day channel. If raising the boundary is the correct view to take, then we would expect the price to rebound upwards from this adjusted lower boundary.

If lowering the boundary is the correct view, then we would expect the price to rebound downwards back from this upper boundary. To solve this dilemma it is simply necessary to wait for a few more days until the situation becomes resolved. Never try to force a conclusion out of inconclusive data.

Figure 11.25 – The 301-day channel in the chart of the British pound/US dollar as at point 823 on 27 May 2006. The direction of the channel has now been verified as changing from falling to rising. The low point in the channel is at point 764, on 3 March 2006.

As can be seen in Figure 11.25, the position on 27 May 2006 shows that the channel has bottomed out and is now rising. Since such a long-term channel is involved, it can be expected to continue to rise for some time into the future, giving excellent trading opportunities.

The actual position of the channel as determined much later shows that the estimation of the channel direction shown in Figure 11.25 was correct. This is shown in Figure 11.26. The channel actually bottomed out on 6 December 2005 at point 721, compared with the estimated bottom being on 3 March at point 764. While this is a relatively large discrepancy, it should be pointed out once again that the gap between the last calculated point and the last data point is one of 150 days, so that this error of 43 days represents only about a quarter of the gap.

What is extremely interesting is that the first estimation of the turning point as shown in Figure 11.21 for 12 February gave 1 December as the turn. This is only three days away from the actual value!

Figure 11.26 – Enough new data has arrived to enable the change in direction of the 301-day channel in the chart of the British pound/US dollar to be confirmed. The channel bottomed out at point 721 on 6 December 2005.

SHORTER-TERM CHANNELS

The advantage of using averages with long periods is that they change direction relatively slowly and therefore it is fairly easy to extrapolate the curve across the gap and up to the present time. The disadvantage is, as we have seen, that a long time elapses between the actual turning point and having enough data to be sure that it has turned. Using 301-day channels, this delay amounted to 73 days in the case of British American Tobacco and 59 days in the case of the British pound versus the US dollar.

In the case of the H.J. Heinz Company, there was an indication of the turning point 72 days in advance. However, as was pointed out earlier, it is only on very rare occasions that we are able to do this. Quite clearly, if we reduce the period of the underlying average, we will reduce this time lag, but on the other hand we should expect the extrapolation to the present time to be more difficult.

MULTIPLE CHANNELS

It might be thought that channels based on different periods will all top out or bottom out at the same time. However, since channels are based on the sums of cycles this is not the case, since the positions of the peaks and troughs in cycles which are reasonably well separated in wavelengths are independent of each other. For channels, this is illustrated in Figure 11.27.

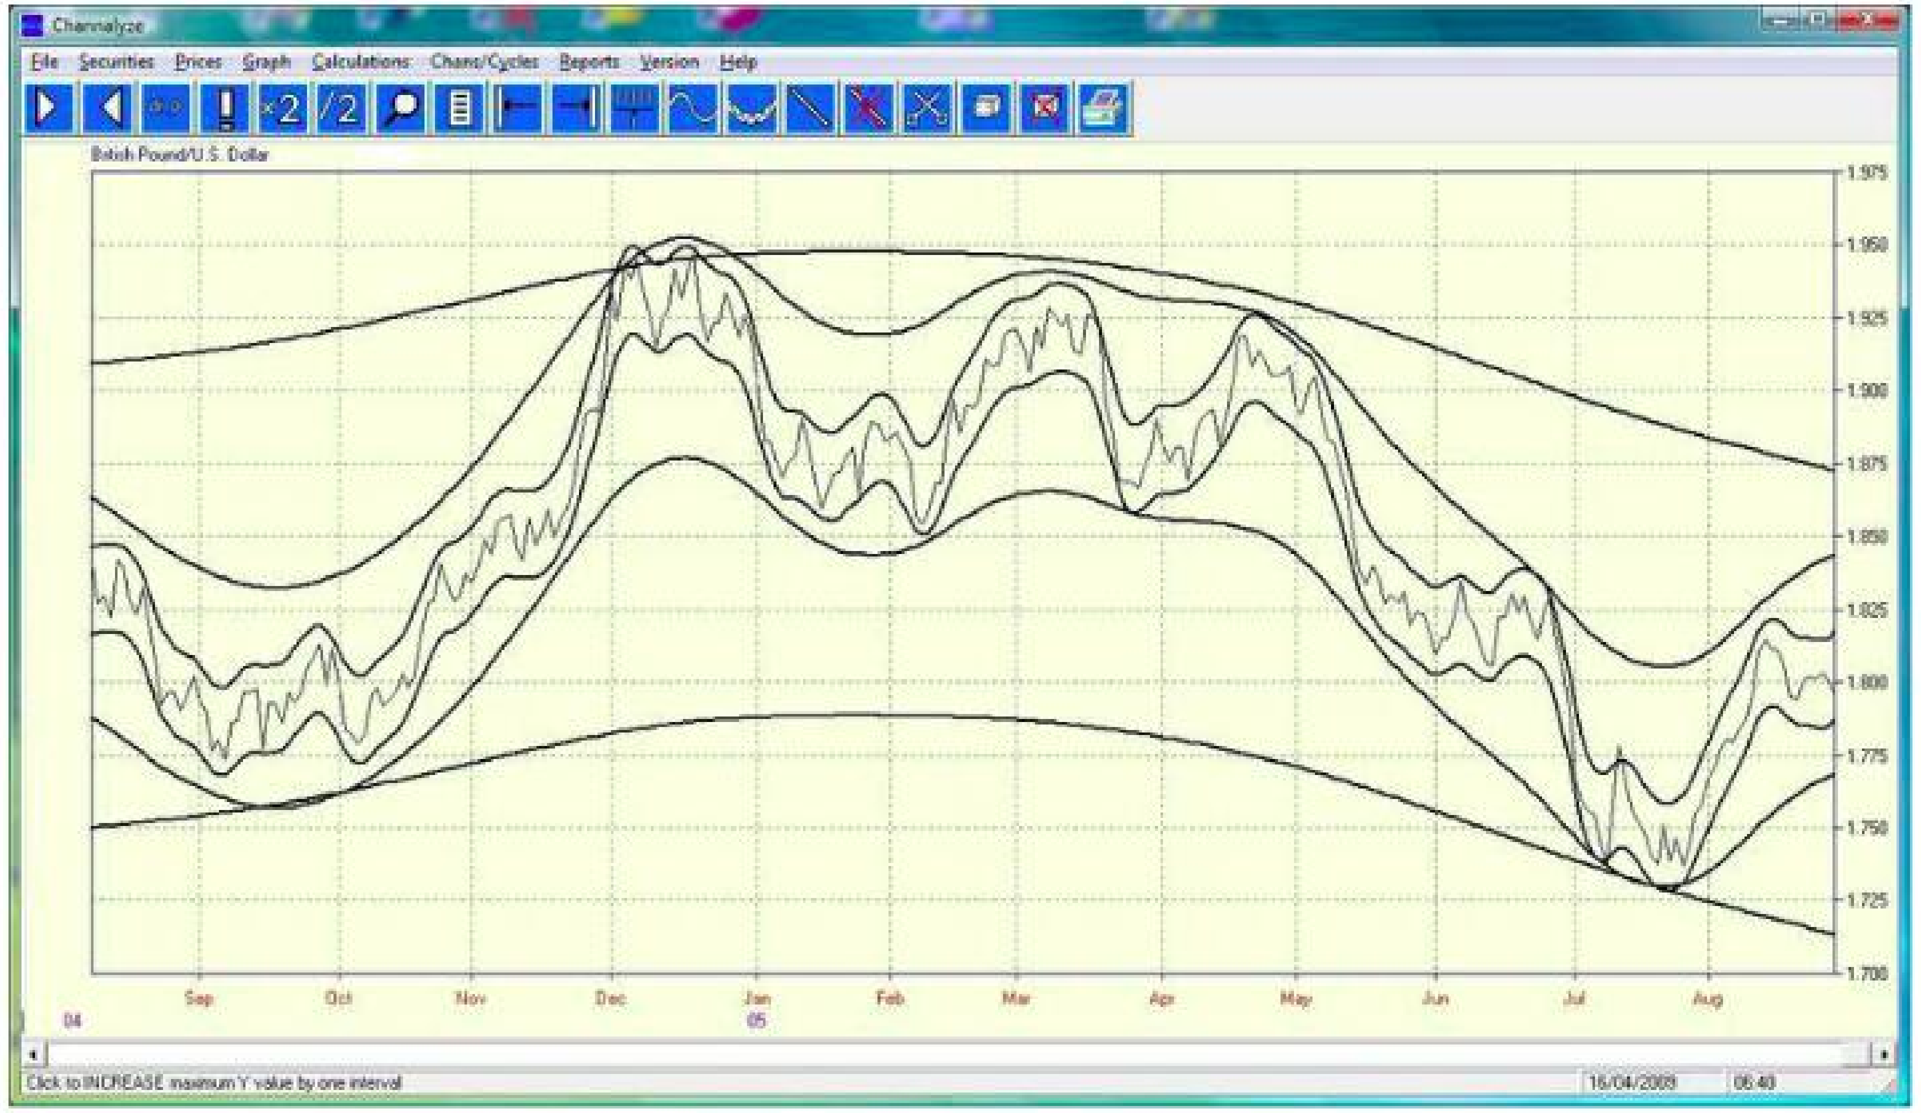

Figure 11.27 – The topping region of the 301-day channel in the chart of the British Pound versus the US Dollar. Also shown are the 51-day and five-day channels. The channels do not top at the same time. Also note the decreasing smoothness as we move to short periods in the underlying average.

Several important points can be made from a closer look at Figure 11.27:

- the three channels do not top out at the same time
- the smoothness of the channels drops off as we move to smaller values for the underlying period. In the case of the five-day channel, this will make it almost impossible to extrapolate
- what was a single top in the case of the 301-day channel breaks up into multiple tops for the shorter-term channels
- each inner channel bounces off the next outer channel.

VERY SHORT-TERM CHANNELS

In Figure 11.28 we show the nine-day channel of the British pound versus the US dollar. This was plotted with the last data point being at point 557 on 20 April 2005. There is no indication in this chart that the rising channel will not continue to rise in the near future.

Figure 11.28 – One of the topping regions of the nine-day channel in the chart of the British pound versus the US dollar. The last data point is at point 557 on 20 April 2005. There is no indication that the channel will be topping out in the near future.

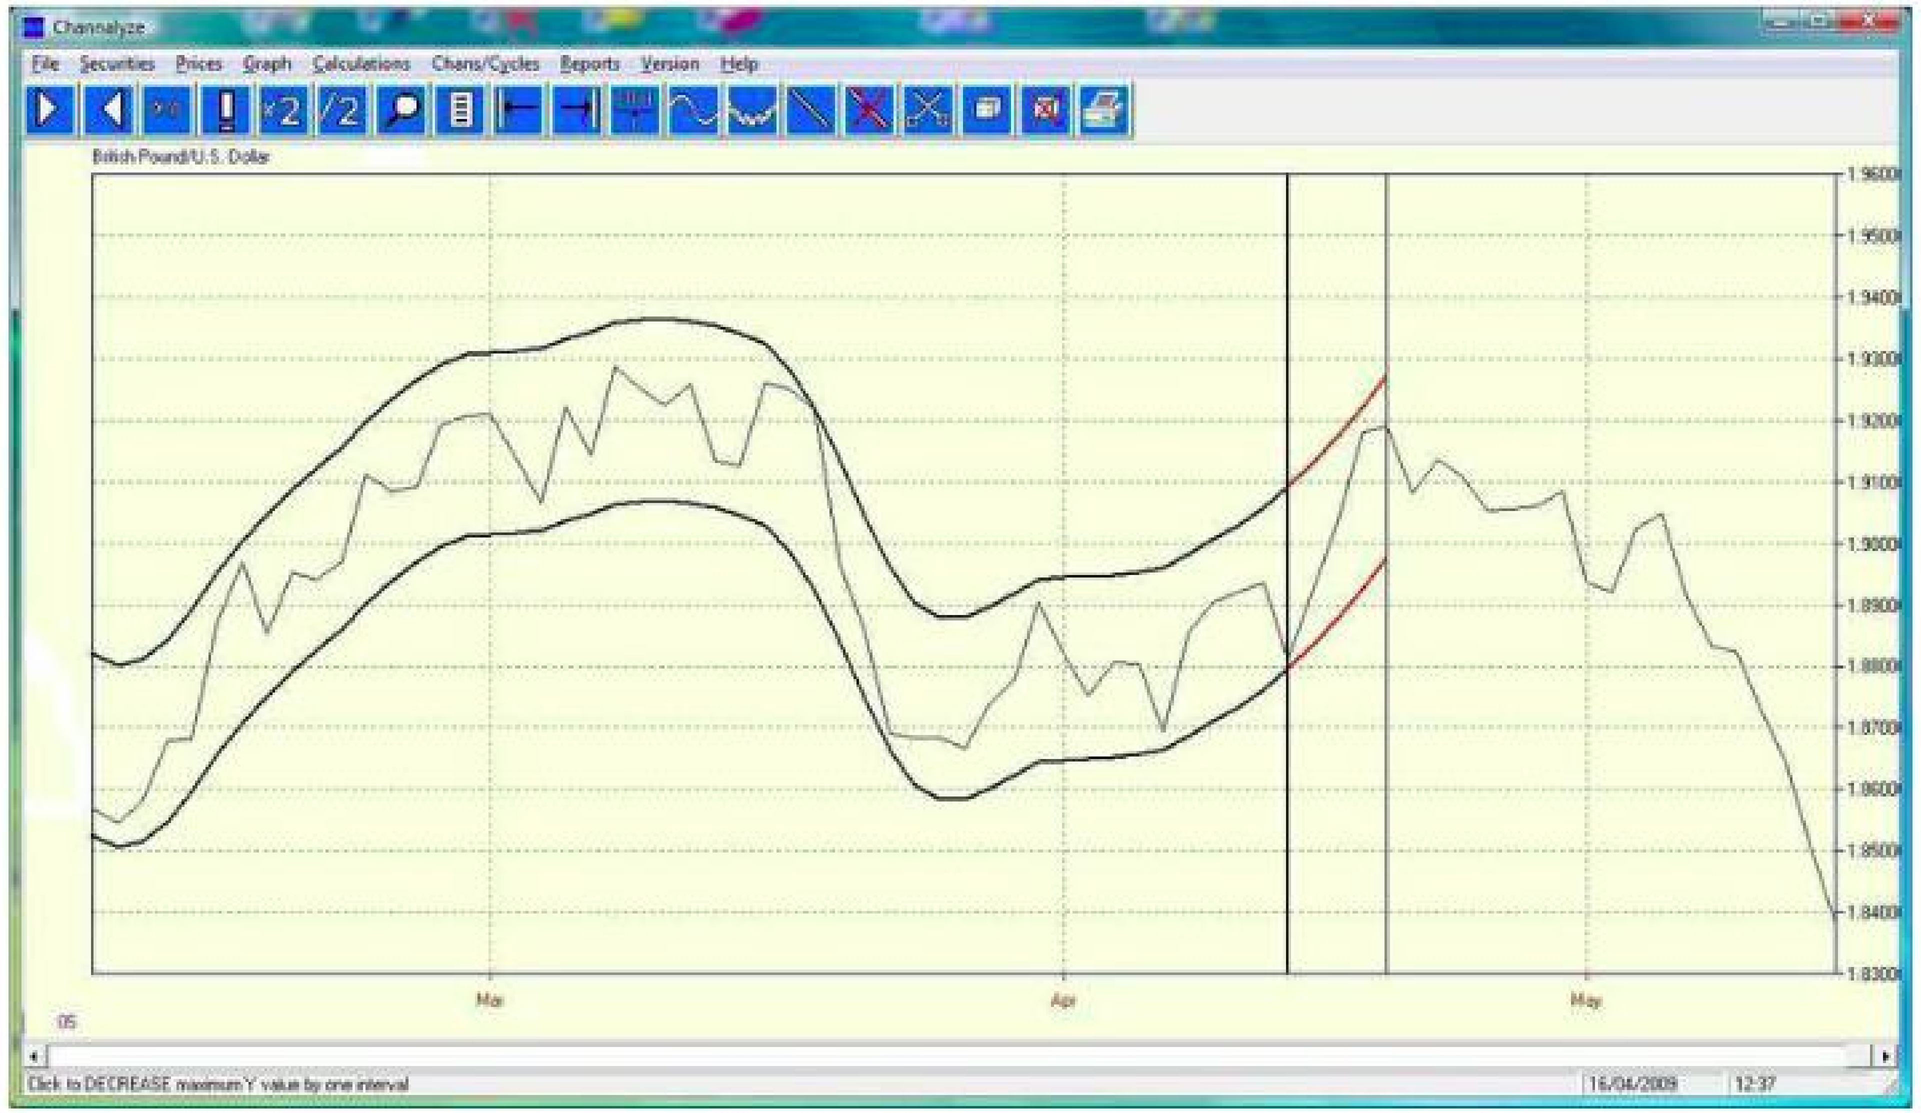

However, a few days later at point 561 on 26 April, as shown in Figure 11.29, we can see that the data has now violated the lower boundary so that the channel has to be adjusted to take this into account.

Figure 11.29 – Here we show the position a few days later at point 561 on 26 April.

The adjusted channel is shown in Figure 11.30.

Figure 11.30 – Here the channel has been adjusted by raising the lower boundary to accommodate the fact that the data has violated the latter.

The message from Figures 11.28 to 11.30 is that when short-term channels are used to aid trading decisions, the position can change very quickly. Usually it is not possible to decide that a channel has changed direction until we reach a situation where the latest data point violates the boundary, making it necessary to carry out an adjustment. More often than not, this adjustment will generate a bend in the channel so that it can be seen to have topped out or bottomed out.

However, if the data is inspected closely, a warning sign that the price itself may reverse direction is given by the relationship between the price and a boundary. In the case of a rising channel such as that shown in Figure 11.28, it can be seen that the peak value at point 556 on 19 April is very close to the upper boundary. Thus the trader would expect the price to fall at this point since it is in a low probability area for maintaining its position.

The main difficulty with short-term channels is that the price movements contained within them, since they are of such a short-term nature, do not give the range of movement that can be seen with longer-term channels. There is therefore a happy medium in which the delay in determining that a channel has changed direction can be balanced against the potential for a large rise or fall.

This chapter has covered the analysis of the channel in order to establish whether it has changed direction or not, and hence establish whether the trend represented by the channel has changed direction.

The issue of the behaviour of the data within the channel, which means an analysis of the short-term trends represented by the intra-channel movement, is covered in the next chapter. It will be seen that the best way of determining whether the short-term trends have changed direction is to establish their nearness to a channel boundary, rather than try to analyse channels based on the extrapolation of short-term centred averages, which is not recommended.