PART I

Standard MIDAS Support and Resistance Curves

CHAPTER 1

MIDAS and Its Core Constituents

The Volume Weighted Average Price (VWAP) and Fractal Market Analysis

Andrew Coles

It was emphasized in the introduction that this book is not about Volume Weighted Average Price (VWAP) but a particular development of it in the MIDAS approach of Paul Levine. This point requires re-emphasis at the start of the book because at the time of writing there's a lively surge of interest in VWAP. As a result, it's becoming harder for newcomers to this area to differentiate between what lies within the ambit of Levine's contributions and what lies outside of it. A timely first aim of this chapter therefore will be to highlight a number of boundaries to the MIDAS approach in relation to its VWAP background.

A second theme will be to look at the main ideas underlying Levine's philosophy of price movement, especially his fractal conception of the markets and the application of multiple hierarchies of curves. This application adds a powerful ubiquitous forecasting capability to the curves and requires separate attention. The discussion will be partly academic in tone in its brief outline of the fractal conception of the markets that was becoming popular when Levine was working on his approach in the early 1990s.

A final theme lays the groundwork for the practical emphasis throughout this book on trading with MIDAS curves. One of the major shortcomings in Levine's lectures is his emphasis purely on the forecasting implications of the MIDAS method. Never at any time did he consider the trade-management implications of using the curves. The final theme of this chapter begins a trend in this book that focuses heavily on using the curves in practical trading contexts.

This chapter is more theoretical than other discussions in this book in outlining Levine's debt to fractal interpretations of the markets and various approaches to VWAP.

However, these deeper perspectives are helpful in understanding the MIDAS method historically as a product of two unique and very different approaches in the markets, which were just beginning to be felt in the early 1990s.

MIDAS and Its Two Key Backdrops: VWAP and Fractal Market Analysis

The MIDAS approach consists of two primary indicators, the basic MIDAS support and resistance (S/R) curves and the more complex topfinder/bottomfinder curves. Let's make a start by considering very generally the relationship these two indicators have to the broader VWAP background prior to their development and that are still very much a part of the professional market trading context today.

Before MIDAS: Initial Motivations for VWAP

There have been several motivations behind the application of VWAP to the financial markets which emerged prior to Levine's development of the MIDAS method. None of them initially involved technical market forecasting, but since they're still very much a part of today's market environment it will be worth outlining them briefly.

Distortion and Price Manipulation

One motivation has stemmed from a closing price free of distortion due to unusual transactions or even intentional price manipulation. An anomalous transaction could be caused by a large accidental buy or sell at a very high or low price level prior to market close.

As an extreme illustration, while this section is being written $1 trillion was temporarily wiped off the market value of U.S. equities on Thursday May 6, 2010, in the so-called 2010 Flash Crash. During a six-minute period the S&P 500 fell nearly 5 percent and the crash was the largest one-day point decline (998 points) in Dow Jones Industrial Average (DJIA) history. By the day's close the markets had recovered to a degree, but the S&P 500 was still 3.2 percent lower. Various reasons have been put forward for the crash, including an errant "fat fingered typo" sell order that set off a chain reaction, a sudden movement in JPY/USD, and even market manipulation.[1] Eventually, in a formal statement published in October 2010, the SEC and CFTC blamed the crash on a liquidity crisis caused by a computer trading algorithm.[2]

Circuit breakers are now being tested to halt such anomalies in the future, but one motivation for calculating the VWAP would be to remove very unusual distortions from the closing price, even if such distortions involve complex intermarket relationships in the currencies and bonds markets through sophisticated computer networks.

Alternatively, direct market manipulation may involve the intentional placing of orders during late market hours at various extreme prices. Again the reasons could be various. For example, closing prices are used for formal statements of the value of a portfolio in a company's annual report and are also occasionally used to calculate directors' remuneration as well as the settlement values of derivatives.[3] Again the VWAP is said to help prevent such skewing of market data.

Guaranteed VWAP Executions

A second motivation for VWAP calculations has emerged from the brokerage industry and bears on the ever-demanding relationship between broker and client. Many brokers will now guarantee their clients that orders are executed at the VWAP (so-called guaranteed VWAP execution) in "targeted VWAP" trading. For example, Euronext, the pan-European stock and derivatives exchange, has available what it calls a "VWAP transaction," based on an average price weighted by security volumes traded in a central order book. A large number of brokerage firms will also guarantee the VWAP for large domains of stocks, especially large caps. Due to the growing popularity of VWAP executions data, vendors such as Bloomberg will also display VWAP prices after market close.

The Minimization of Market Impact and Trader Assessment

A third and fourth motivation have arisen from the very heavy volume trading undertaken in the mutual and pensions industry. Here large investors aim to be as passive as possible in their executions and use the VWAP to ensure that they are entering the market in line with typical market volume. This minimizes market impact, which in turn reduces transaction costs. Thus, a final related motivation would be the actual assessment of trading performance: a large institutional trade entry beyond the VWAP may be criticized in light of higher transaction costs; similarly, an order filled above the daily VWAP would be regarded negatively in view of the slippage implications.

Standard VWAP Calculations

Now that the nontrading motivations for VWAP are understood, it would be helpful before turning to Levine's MIDAS approach to obtain a basic understanding of how the VWAP is calculated and how basic VWAP curves appear on a chart. In part, this discussion should also alleviate some of the confusion that has arisen around the relationship between VWAP and the MIDAS approach.

The VWAP is calculated by multiplying the volume at each price level with the respective price and then dividing by the total volume. The more volume traded at a certain price level, the more impact it has on the VWAP.[4] Here is the basic formula for VWAP calculations:

$$\Sigma(Pn * Vn)/\Sigma(Vn)$$

where

P = price of instrument traded
V = volume traded
n = number of trades (i.e., each individual trade that takes place over the selected time period)

There are variations on the basic formula. For example, George Reyna finds the following version more useful:[5]

$$(((Hc * Lc)/2) * Vc)/(Vc - V(c - s))$$

where

H = high price
L = low price
V = volume
c = current bar
s = launch point[6]

As a simple illustration of calculating the VWAP, we can go back to the original VWAP formula and calculate the VWAP over 15 minutes on a 5m chart of the DAX March 2010 futures. We'll use the closing price of three 5m bars:

Bar #1: 5,827 with a volume of 2,856 contracts
Bar #2: 5,819.5 with a volume of 1,729 contracts
Bar #3: 5,816.5 with a volume of 2,271 contracts

The average price over this 15-minute period is the total number of contracts divided by 3, or 5,821 contracts. But let's calculate the VWAP. The result obtained will depend on which method of utilizing the formula we choose. Day trading software firms will probably use one of two algorithmic procedures to derive it.

The first, usually assumed to be the more accurate method, is known as "cumulative VWAP." The first step would be to multiply the closing price with the volume for each of the three bars, arriving at the following numbers:

16,641,912
10,061,915.5
13,209,271.5

The next step would be to add them together to arrive at 39,913,099. To arrive at the denominator, the volume numbers would be summed to get 6,856 contracts. With the division, the cumulative VWAP would therefore be 5,821.630 (this method is usually calculated to three decimal places).

A second method of arriving at the VWAP is known as "iterative VWAP." It uses the last value of the VWAP as the basis for calculating the VWAP on the next trade.

This is an example of the procedure:

First iteration: (5,827 ∗ 2,856) / 2,856 = 5,827
Second iteration: 5,827 + [(5,819.5 − 5,827) ∗ 1,729] / (2,856 + 1,729) = 5,824.172
Third iteration: 5,824.172 + [(5,816.5 − 5,824.172) ∗ 2,271] / (2,271 + 2,856 + 1,729) = 5,821.830

Thus, the iterated VWAP for this same time period is 5,821.830, as opposed to 5,821.630 in the cumulative VWAP approach. As more trades (iterations) are made, the closer the two VWAP calculations will become.[7]

Aside from there being variations of the VWAP formula and calculation differences, another potential source of confusion is that the basic VWAP formula is identical to the one for the volume weighted moving average (VWMA).[8] The two differ only indirectly in terms of the calculation procedure in trading platforms, with the VWMA relying on the "sum" (summation) function and the VWAP utilizing the "cum" (cumulative) function. The difference this makes will be illustrated in a moment in Figure 1.1. It's also worth pointing out that some platforms additionally calculate the Volume Adjusted Moving Average (VAMA), a slightly different curve that's based on the "mov" (moving average) function and that results in a variation of the VWMA

FIGURE 1.1 5m chart of Eurex DAX September 2010 futures showing the DAX as a basic line plot (heavy black).

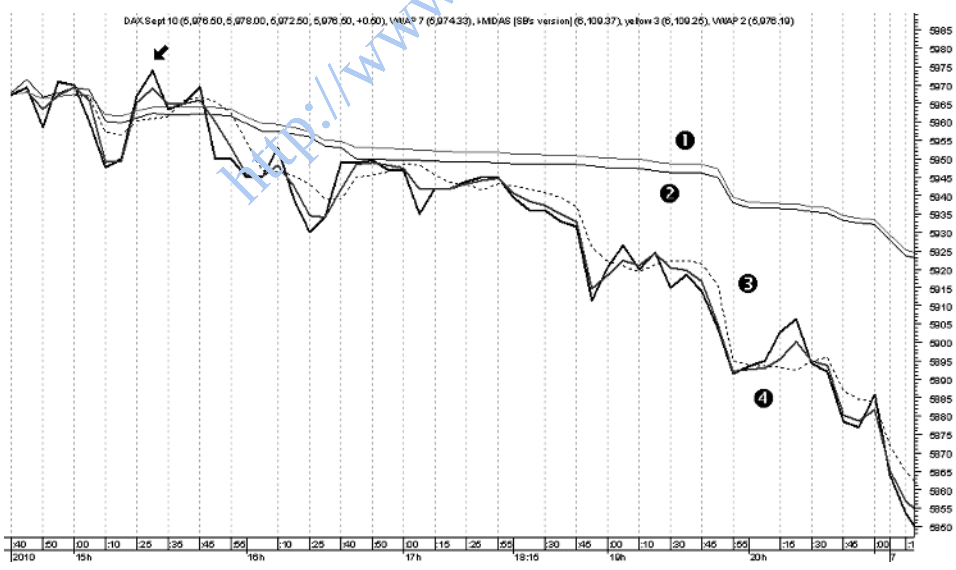

Plot (1) (gray) = standard VWAP; plot (2) (black) = MIDAS; plot (3) (dotted) = VWMA; and plot (4) (heavy gray) = VAMA.
Source: eSignal and Metastock. www.esignal.com and www.equis.com.

plot. Figure 1.1 is a 5m chart of the German DAX September 2010 futures illustrating four curves alongside the dark black line plot of the DAX. Plot (1) (gray) is a standard VWAP curve anchored to the start of the chart. Plot (2) (black) is a basic MIDAS curve. Plot (3) (dotted) is the VWMA, and plot (4) (heavy gray) is the VAMA. We'll come to the discussion of MIDAS curves shortly, but the purpose of this chart is to illustrate how different these curves appear on a chart even though there is so much conflation over the use of the terms used to describe them.

The conflation is at its worse with regard to the terms "VWAP" and "MIDAS." Indeed, many traders who use MIDAS analysis techniques are actually using VWAP curves without realizing it. Yet there are four reasons why traders who deploy MIDAS techniques should ensure that they're using the MIDAS formula (see below) and not the standard VWAP formula:

1. As illustrated in Figure 1.1, the first plot (standard VWAP) is quite different from the second (basic MIDAS).
2. There are variations of the basic VWAP formula (Reyna's version is a good illustration). There's the potential therefore for an even greater difference between VWAP and MIDAS curves.
3. There are even alternatives to the way the standard VWAP formula is calculated, as illustrated in the difference between the cumulative and iterative methods. These methods can give rise to further variations between a standard VWAP and MIDAS plot.
4. VWAP utilizes the average price whereas many who use MIDAS curves use the low price in uptrends and the high in downtrends (Hawkins is an example). This again will create significant differences between a standard VWAP plot and a MIDAS curve.

Trading Applications of VWAP

As already noted, the earliest motivations for establishing the VWAP were not related to technical market forecasting. The first published use as a market entry criterion appears to be trader Kevin Haggerty's in a 1999 interview. Haggerty stated that he favored a simple methodology of choosing long positions when price is above its daily VWAP and short positions when it's below.[9] However, in the past few years there has been a blossoming of interest in VWAP and now there are seemingly as many ways of utilizing it for trading purposes as there are traders taking an interest. As noted, the problem is that many traders use the term "VWAP" erroneously to refer also to MIDAS curves, so when trading ideas are being discussed it's often hard to know which particular curve a trader has in mind.

Bob English, of The Precision Report, has argued that the previous day's closing VWAP is a powerful support and resistance pivot for the current day, often determining the absolute high and low.[10] The trader Brett Steenbarger, PhD, plots the VWAP from the start of the new day's futures session and views its direction as giving a sense to the intraday trend. In trending market conditions, he'll stay to one side of the VWAP, whereas if the market is in a trading range he'll consider trading both sides of it.[11] Participants in the trader forums are also busy with new ideas. For example, one long and influential thread on the Traders Laboratory web site outlines a trading system based on combining the daily VWAP with a volume distribution histogram similar to market profile.[12]

VWAP and Paul Levine's MIDAS System

In relation to the VWAP backdrop there are two main aspects to MIDAS support/resistance curves that differentiate them from it.

The Formula Difference

First, there's Levine's variation of the basic VWAP formula. Second, which we'll come to below, there's his view of how to launch MIDAS curves. As for the variation, in his twelfth lecture he gave the formula for his MIDAS S/R curves as follows:

$$\Sigma(Pn * Vn) - \Sigma(Ps * Vs)/\Sigma(Vn) - \Sigma(Vs)$$

where

Pn and *Vn* are the current cumulative price and volume
Ps and *Vs* are the cumulative price and volume at the MIDAS curve launch
Vn is the current cumulative volume
Vs is the cumulative volume at the MIDAS curve launch

In plain English the formula reads: (cumulative average price)(volume at a given instant) – (cumulative average price)(volume at a period d units of cumulative volume earlier), all divided by d, where d is the cumulative volume displacement measured from the launch point to the given instant.

We've already seen from Figure 1.1 that Levine's variation of the VWAP formula results in a curve that differs from a standard VWAP curve. The question is why Levine felt it necessary to introduce this minor modification to the original VWAP formula. He never tells us in his lectures, but it's possible to speculate accurately as to his reason. To do so, we need to look at an important theoretical idea that distinguishes the MIDAS method from more basic approaches involving VWAP.

Paul Levine's Philosophy of How Market Prices Evolve

This theoretical idea lies in two factors that were of fundamental importance to Levine:

(i) The critical choice of where to launch MIDAS curves, and
(ii) The multiple applications of MIDAS S/R curves based on a fractal conception of price movement

P1: OTA/XYZ P2: ABC
JWBT458-c01 JWBT458-Coles March 9, 2011 23:10 Printer Name: Yet to Come

It's the combination of (i) and (ii) that turns the MIDAS approach into a genuine trading system as opposed to a set of indicators on a chart.

We can better understand these two features by reducing Levine's philosophy of price movement implicit in his lectures to five key tenets:

1. The underlying order of price behavior is a fractal hierarchy of support and resistance levels.
2. This interplay between support and resistance is a coaction between accumulation and distribution.
3. This coaction, when considered quantitatively from raw price and volume data, reveals a mathematical symmetry between support and resistance.
4. This mathematical symmetry can be used to predict market tops and bottoms in advance.
5. Price and volume data—the volume weighted average price (VWAP)—subsequent to a reversal in trend, and thus to a major change in market (trader) sentiment, is key to this process of chart prediction.

The Critical Choice of Where to Launch VWAP Support/Resistance Curves

According to factor (i), Levine believed that when charted all price behavior can be reduced to multiple hierarchies of support and resistance. What this means is that as price moves forward at all degrees of trend, it is either testing existing support or resistance or breaking out from them to create new hierarchies. Accumulation therefore amounts to price respecting existing support, breaking out of overhead resistance, and moving up the chart to create new levels of resistance and support. Distribution amounts to its opposite. According to tenet (4), this repetitive price behavior can be captured using the MIDAS support and resistance curves with the same formula. In other words, it makes no difference to the algorithm whether price is rising (accumulation) or falling (distribution).

With tenet (4) in mind, the question is how MIDAS can be used maximally to highlight these hierarchies of support and resistance. This is where tenet (5) assumes importance. It's this tenet that marks the main distinction between standard applications of VWAP and Levine's specialized use. It's also why these MIDAS support and resistance curves have come to be known as "anchored VWAP" curves. Levine focuses on this topic in lecture eight. He ends lecture seven with the following remark:

> We have not yet specified the interval over which the averages are to be taken. In fact, it is this *choice of averaging interval which uniquely distinguishes the MIDAS method*.[13]

In lecture eight he first identifies and then justifies this averaging interval. He argues that where price finds subsequent support or resistance is directly associated with where there was a change in the underlying psychology, otherwise there'd be no change in trend. This is where the averaging must start and hence where a MIDAS curve should be launched, or "anchored."

With this information, we can now answer a question left unanswered earlier, namely why Levine felt it necessary to introduce a minor modification to the original VWAP formula. As we've just seen, Levine believed that the launch bar of a MIDAS curve was the last bar—and hence the bottom—of the previous trend. Since for him the VWAP subsequent to a reversal in trend is the critical data, he subtracted the VWAP of the launch bar from subsequent data because he believed that the launch bar VWAP was a part of the previous market psychology before it changed direction and thus marked a new change in sentiment. He might have omitted the VWAP of the launch bar from the equation entirely instead of subtracting it from the subsequent VWAP. Or he might instead have launched MIDAS curves from the price bar subsequent to the last bar of the previous trend and simply used the original VWAP formula. For reasons he doesn't specify, he does neither, and opts for the approach that underlies the MIDAS formula provided earlier. Possibly Levine had done research on these alternatives and found them wanting. He never tells us one way or the other.

When it comes to the actual plotting of the curves, subsequent reversals in trend, which the MIDAS S/R curves are intended to capture, are connected mathematically to this change in sentiment, since subsequent trader mood is intimately linked to it. Here is Levine again:

> Our "message" is that instead of "moving" averages, one should take fixed or "anchored" averages, where the anchoring point is the point of trend reversal.[14]

The implication for trading is this. If I know that certain points on a chart are trend reversals and that the corresponding changes in psychology are associated with subsequent levels of support and resistance, I can use this information to trade these subsequent levels, provided I have the right tool—in this case, a MIDAS curve—to identify these subsequent levels. By contrast, nothing this precise is implied by the VWAP itself.

Compare, for example, Figure 1.2 with Figure 1.3. Figure 1.2 is a 5m chart of the March 2010 Xetra DAX futures and has a standard anchored VWAP curve plotted throughout the day from the market opening. As noted earlier, some traders will start what is actually an anchored VWAP curve from the market open and stay to one side of it in trending days or trade both sides of it in rangebound conditions. Now there's nothing wrong with these suggestions, but they're not MIDAS strategies. For one thing, the curves are standard VWAP curves not MIDAS curves. For another, today's open (or yesterday's close) would figure in MIDAS thinking only if it represented a change in market psychology. Where it doesn't, I showed in a previous article that plotting a MIDAS curve from the previous day's close or today's open is ineffectual in relation to the MIDAS method.[15] Figure 1.2 is a case in point. Here there's no significant swing high or low involving the open; as a result, the MIDAS curve drifts through the opening hours of trading and then displaces as prices make a sharp upside move. The two pullbacks circled represent good opportunities to join the ongoing trend. However, it's clear that the MIDAS curve has displaced far too much to be of any help and we get little aid from indicators, such as the stochastic, which is already

FIGURE 1.2 5m chart of Xetra DAX March 2010 futures with a standard VWAP curve plotting from the open.

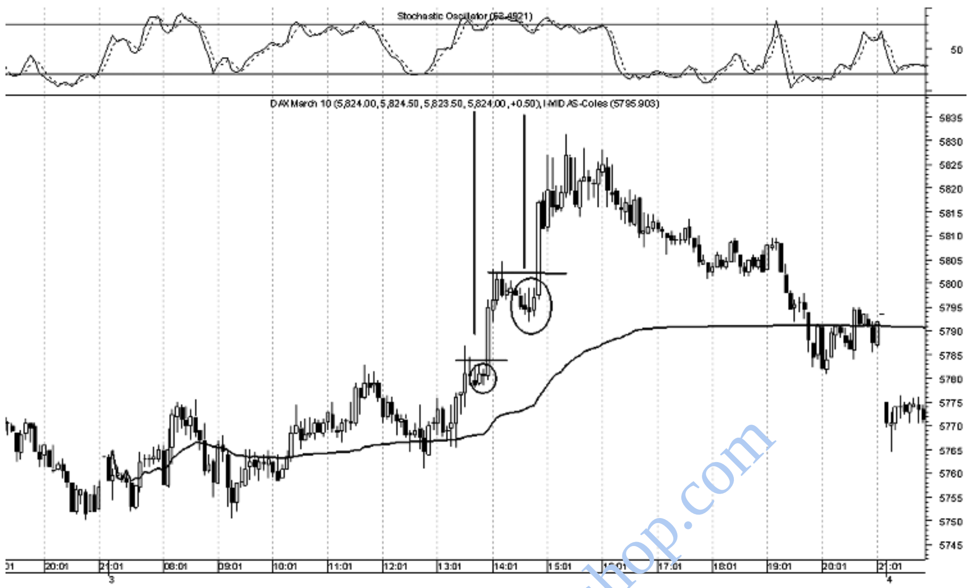

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

FIGURE 1.3 The same 5m chart with an anchored MIDAS support curve accurately capturing the two pullbacks.

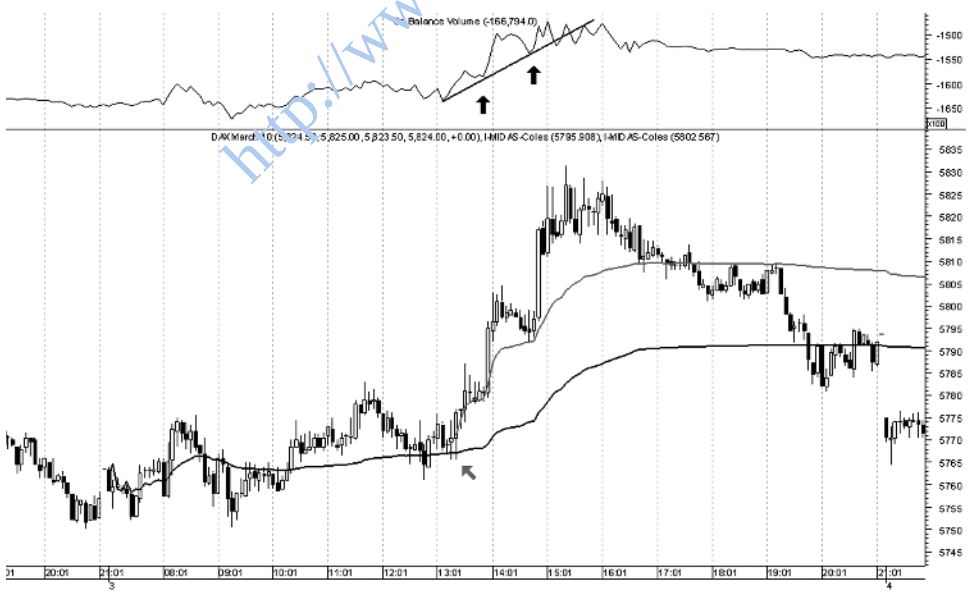

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

overbought. The best we could do is trade basic breakouts while the MIDAS curve itself is irrelevant.

By contrast, Figure 1.3 is the same chart with a MIDAS support curve meaningfully anchored to the start of the new phase of the uptrend highlighted by the gray arrow and interacting directly with its pullbacks. By a judicious use of Japanese candlesticks, both to gauge reversals and to set stops, a properly anchored MIDAS curve checks every box a trader requires, including trend direction, trade timing and entry, plus trade-management in clear risk levels.[16] In Figure 1.3 the On Balance Volume indicator also significantly enhances the MIDAS signals in virtue of its trend line properties, as can be seen at the arrow highlights (see also Chapter 3).

Multiple Applications of MIDAS S/R Curves Based on a Fractal Conception of Price Movement

Moving on to factor (ii), anchoring MIDAS curves to clear points on a chart where there's a change in psychology isn't the only theoretical element that distinguishes the MIDAS system from basic VWAP. The other major determinant is Levine's insistence on the application of multiple curves to the same chart. In his lectures, Levine maintained that support and resistance levels connected with earlier points of trend reversal should be associated with a hierarchy of theoretical curves. I summarized this idea in terms of the first of the five tenets earlier. This is one of the factors that truly establish the MIDAS approach as a genuine standalone trading system, since the concept of hierarchy presupposes multiple levels of price action, none of which are beyond the analytical reach of the anchored MIDAS curves. The concept of the market as a hierarchy of support and resistance levels presupposes in turn that price formations are fractal. Levine uses the term "fractal" four times in his lecture series, with the main passage being this:

> The foregoing properties [namely, similar zigzags in price behavior at all degrees of trend] of *self-similarity* and *scale-independence* are characteristics of *fractal* behavior. *The fractal nature of stock price fluctuations has been recognized for some time on purely empirical grounds*. What has been missing is an understanding of why markets should behave fractally (i.e., beyond the obvious fact that they are *complex non-linear dynamic systems*). In the Midas method, we have seen that the complex zigzags in price behavior can be (to quote article #8) "understood with respect to a single algorithmic prescription: support (or resistance) will be found at the VWAP taken over an interval subsequent to a reversal in trend." The psychological elements of greed and fear, whose quantification led to this algorithm, *apply to investors/traders across all time scales* (my italics throughout).[17]

What is meant by "fractal" in this context, and how precisely is it linked to the notion of a hierarchy of support and resistance levels? This is an important question because without its fractal capabilities MIDAS would be a shadow of its true forecasting potential. Consequently, we'll complete the first half of this chapter by focusing on the crucial role that fractal market analysis plays in the MIDAS method.

Levine refers to the fractal nature of markets as a self-similar, scale-independent, nonlinear dynamic system, and of this fractal nature as being proven empirically. As a research physicist publishing his lectures online in 1995, Levine would not have been deferring to Elliott Wave theory in claiming that the fractal market hypothesis had been proven empirically. He would have been referring to a particular statistical method affirming this hypothesis. It is worth spending a section or two on this topic, not only to enlighten the role played by the fractal market hypothesis in Levine's thinking but also to allow other relevant discussions of it in later chapters.

MIDAS and Fractal Market Analysis

The empirical grounds Levine refers to have their origin in the pioneering work of the British hydrologist H. E. Hurst (1880–1978) and subsequently in the applications of Hurst's ideas to the financial markets by Benoit Mandelbrot. From 1913 Hurst had spent his early career as head of the Meteorological Service working on the Nile River Dam Project with its focus on the control and conservation of Nile waters. Working with vast records of contemporary and historical rainfall and river flow patterns in the Nile and its network of tributaries, Hurst came to believe that the Nile's overflows weren't random and that there was evidence of nonperiodic cycles (one of several hallmarks of a fractal process (see below)). As a result, Hurst developed his own statistical methodology to test this assumption known as Rescaled/Range (R/S) analysis. His work was formally published in 1951[18] and was subsequently refined by Mandelbrot and others when it began to be applied extensively to financial market time series.[19] As a practicing physicist with an abiding interest in the financial markets, it's possible that by the 1990s Levine was familiar with some of this work. However, it's more likely that he was drawing on the recently published books of Edgar Peters in 1991 and 1994,[20] although there was also other material on fractals discussing the financial markets in more or less detail of which Levine might have been aware.[21] Much of this work describes R/S analysis as proving empirically that the financial markets are fractal time series. For reasons that will emerge later in the book, it will be worth explaining the nature of this empirical evidence in a little more detail as well as linking it to several core ideas in Levine's market philosophy.

R/S analysis claims to show that the financial markets are fractal because it is a statistical methodology for distinguishing between random and nonrandom (fractal) time series. When Einstein looked at the random path followed by a particle in a fluid (Brownian motion), he discovered that the distance covered increases with the square root of time used to measure it ($R = T^{0.50}$, or the "T to one-half rule," where R = distance covered and T = a time index).[22] This equation is now commonly used in finance to annualize volatility by standard deviation. For example, the standard deviation of monthly returns is multiplied by the square root of 12 on the assumption that the returns increase by the square root of time. Here markets are assumed to follow a random walk (i.e., exhibit Brownian motion). By adapting the T to one-half rule and embedding it within a larger statistical procedure,[23] Hurst arrived at the R/S methodology that produces an exponent he called the K exponent and which has since been labeled the Hurst exponent by Mandelbrot in honor of Hurst. It's the Hurst exponent, then, that estimates the degree of nonrandomness in time series to which it is applied.[24] A vast amount of recent work has focused on the international financial markets using this technique,[25] albeit with varying results in regard to the actual Hurst exponent for each market.[26]

If the R/S analysis applied to a given time series results in a Hurst exponent of 0.5, it means that the time series is a pure random walk; in other words, it increases with the square root of time as Brownian motion. However, if $0.50 < H < 1.00$, it implies a "persistent" time series covering a greater distance in the same timespan than a random walk—hence the term "fractional Brownian motion"—and it is characterized by a long-term memory effect. In other words, what happens today affects what happens tomorrow, and the changes are correlated. This means that there is sensitivity to initial conditions (another hallmark of a chaotic system) and that this long-term memory effect affects changes at all degrees of trend (daily changes are correlated to later daily changes, weekly changes to weekly ones, and so on). There is thus no characteristic timescale, yet another hallmark of a fractal time series.[27] If $H < 0.5$, it implies that the time series is antipersistent, meaning that it covers less distance than a random walk because it is reversing itself far more frequently. In the financial markets antipersistent price activity would be typically found in tight congested (rangebound) markets.

If $0.50 < H < 1.00$ (that is, a persistent time series with long-term memory), it also means that H is scaling according to a power law as there is a shift from smaller to larger increments of time in the time series. Power laws are common to all fractal time series as well as to fractals in the natural world as diverse as city population size, earthquake magnitudes, clouds, coastlines, word frequency in languages, in addition to thousands of other natural phenomena. In virtue of these power laws, all fractals have in common the fact that they don't scale up or down according to the same ratio, hence the term "scale invariant" which Levine refers to in the passage quoted. For example, trees and coastlines are well-known fractal systems because although they scale up and down, each scaling level is similar to but not identical with the others. Trees have branches that resemble one another (global determinism), but none are identical close up (local randomness).[28] Applied to examples such as these, and also to time series such as the financial markets, the term "qualitative self-similarity" is used. As we have seen, the power law that explains this is related to the Hurst exponent. This power law scaling feature is also sometimes called the fractal dimension. The fractal dimension is related to the Hurst exponent by the equation $D = 2 – H$. Thus, a Hurst exponent of 0.7 is equivalent to a fractal dimension of 1.3. The fractal dimension is often used as a means of describing how fractal objects, such as coastlines, fill the space around them and how they scale in relation to it. Fractal time series, on the other hand, scale statistically in time,[29] and so the fractal dimension of a time series measures how jagged or rough it is ("statistical self-similarity"). A straight line would have a fractal dimension of 1, while a random time series would have a fractal dimension of 1.50. A fractal

FIGURE 1.4 Dietmar Saupe's illustration of time series ranging from antipersistence to a time series exhibiting clear long-term memory processes (the lines added to the final time series are my own).

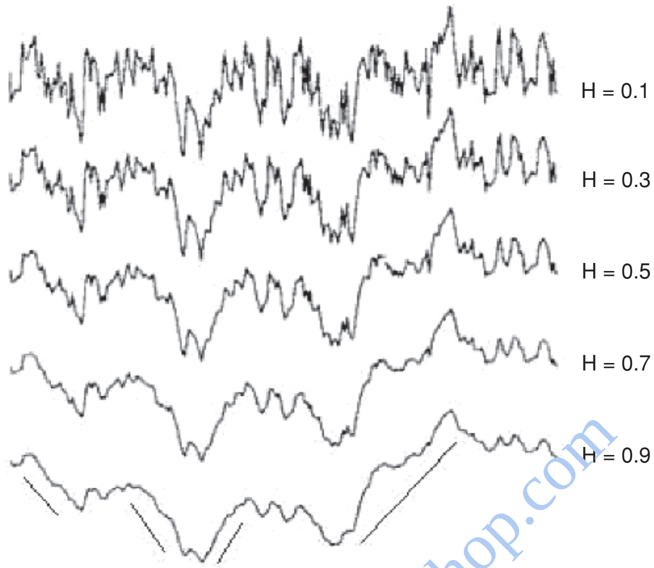

Source: Adapted from D. Saupe, "Random Fractal Algorithms," Chapter 2, in *The Science of Fractal Images*, ed. D. Saupe and Heinz-Otto Peitgen (New York: Springer Inc, 1988).

time series would therefore always have a fractal dimension between 1 (indicating a pure deterministic process) and 1.50 (indicating a random walk or Brownian motion). Thus, a fractal time series increases at a faster rate than the square root of time (= H = 0.5). Anything between 1.50 and 2 would imply the antipersistence mentioned earlier.

In the remainder of this section, let's highlight more clearly the relationship between the fractal interpretation of time series and the notion of "anchoring" MIDAS curves and how the latter depend critically on the former to work at all. A very helpful visual appreciation of statistical self-similarity can be seen in Figure 1.4, which is derived from an illustration by Dietmar Saupe in his chapter "Random Fractal Algorithms" in Saupe and Peitgen.[30]

From a MIDAS viewpoint, what is interesting about the first of these two time series is that they're both antipersistent (H < 0.5) and not particularly amenable to MIDAS curves. The middle time series, with a Hurst exponent of 0.5 (D = 1.50), is a pure random walk. Here we begin to see opportunities to launch MIDAS curves from certain highs and lows. However, the last two time series are fractal (D = 1.3 and 1.1 respectively). It can be seen straightaway how inviting they are to MIDAS analysis. The second of the two, with a fractal dimension of 1.1, gets close at certain points to being a deterministic straight line (hence the high Hurst exponent), especially in the four areas highlighted. Here, because of the high Hurst exponent, the trends are

P1: OTA/XYZ P2: ABC
JWBT458-c01 JWBT458-Coles March 9, 2011 23:10 Printer Name: Yet to Come

showing distinct signs of acceleration and as such are suitable for the launch of the topfinder/bottomfinder indicator. This is an important point (see Chapter 4 where this indicator is examined in relation to the fractal characteristics of time series components to which it should be applied).

The Real-Time Fractal Dimension and MIDAS Curves

In the meantime, we round off this discussion by looking at an actual time series in relation to Figure 1.4 that illustrates a full application of standard MIDAS S/R curves. Figure 1.5 is a 15m chart of the September 2010 Eurex DAX futures spanning nearly six trading days from July 5 to July 12. This entire period has a Hurst exponent of 0.526138 and thus a fractal dimension of 1.473862. The Hurst exponent is graphically illustrated in Figure 1.6, a common way of presenting the Hurst exponent in the financial markets as discussed in Peters (1991 and 1994).

With a fractal dimension of 1.473862, the DAX futures are barely more than a random walk over this timeframe and should be compared with the third time series in Saupe's illustration in Figure 1.4. Yet as Figure 1.5 reveals, it's still easy to apply standard MIDAS support and resistance curves to this chart as well as three topfinder/bottomfinder curves (points (1), (2) and (3), even though the latter function correctly only when the market is exhibiting a very high degree of persistence. In fact, the overall Hurst exponent in Figure 1.5 is misleading, since the price series

FIGURE 1.5 15m chart of DAX September 2010 futures showing six trading days with a Hurst exponent of 0.526138 and hence a fractal dimension of 1.473862.

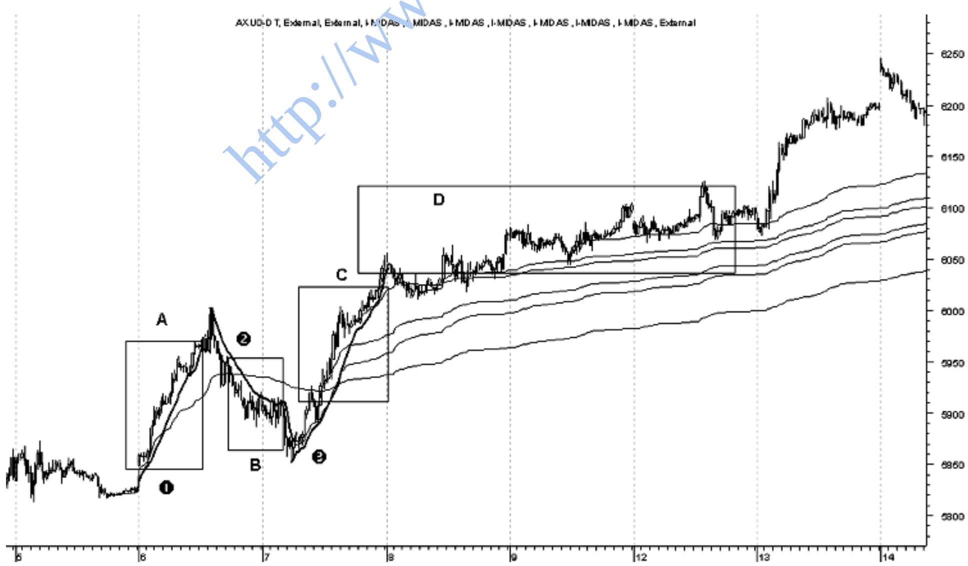

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

FIGURE 1.6 R/S chart of the DAX 15m September 2010 futures over six trading days.

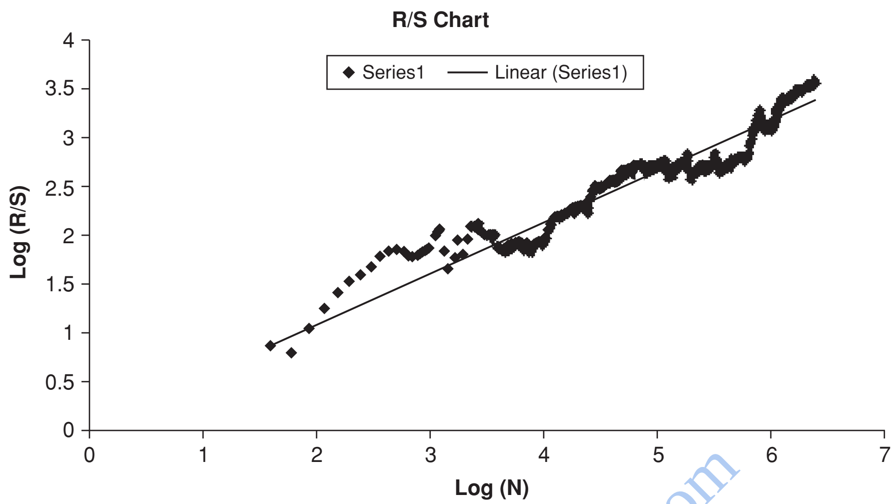

Source: Xlpert. www.elpert.com.

clearly shows signs of acceleration in the trend portions highlighted by rectangles A, B, and C where the three topfinder/bottomfinder curves have been launched successfully.

To get a more accurate real-time perspective on the fractal dimension of the market, and hence an accurate mathematical context in which to use MIDAS, it's now possible to obtain real-time Hurst estimates[31] and the corresponding fractal dimension in virtue of indicators such as the Fractal Dimension Index (FDI). Figure 1.7 is the same 15m chart with the FDI programmed into eSignal. In relation to the 1.5 random walk level, the real-time fluctuation of the indicator clearly shows the fractal dimension decreasing during trending periods and increasing into antipersistent levels during periods of deceleration and rangebound conditions. In her 2007 study of the FDI, Radha Panini argued that the indicator is a much better filter than other trend-measuring indicators such as Wilder's Average Directional Index (ADX) and the Vertical Horizontal Filter (VHF) when used alongside moving average systems, breakout systems, and oscillator trading systems such as the RSI.[32] If this is true, there is undoubtedly an even better theoretical synergy between MIDAS and an indicator that actually measures the fractal dimension of markets in relation to which MIDAS was primarily developed.

Since, as Saupe's diagram in Figure 1.4 illustrates, there's a critical relationship between the fractal dimension of a market and the successful application of MIDAS curves, a real-time fractal measuring device such as the FDI would prove to be a much better fit with MIDAS than would any other technical tool. I suspect that Paul Levine would have approved strongly of it, given his tendency to choose other indicators selectively to work alongside MIDAS.

FIGURE 1.7 Same 15m chart of the DAX futures covering the same trading period with the Fractal Dimension Index plotting in real-time in the lower pane.

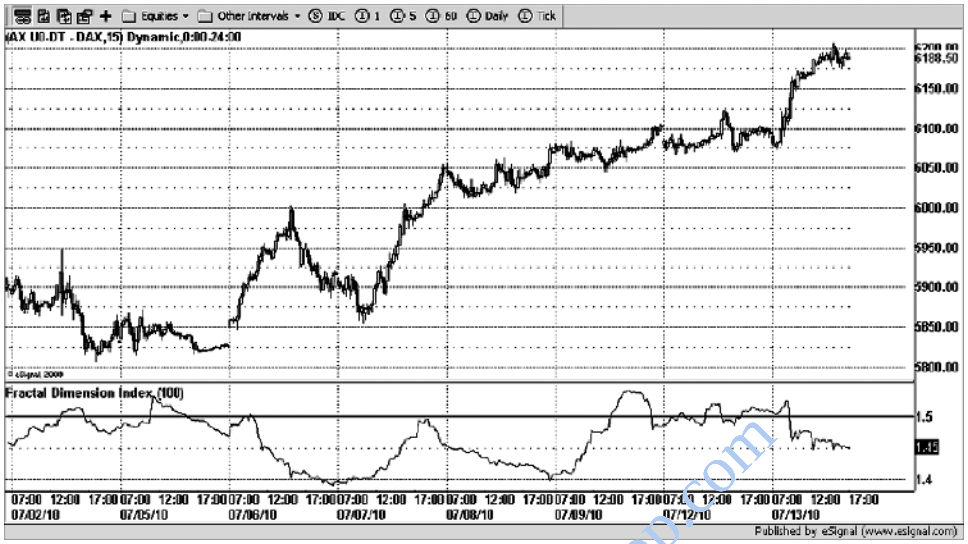

Source: eSignal. www.eSignal.com.

The Background Influence on Levine

Finally, on a point of pure academic interest I've suggested that Levine's view that trader emotions are fractal probably has its origins in Peters's fractal market hypothesis. First, let's remind ourselves of what Levine said:

> The psychological elements of greed and fear, whose quantification led to this algorithm, apply to investors/traders *across all time scales* (my italics).[33]

In so far as fractal market activity is assumed to be linked in the MIDAS approach by means of predictable human emotion, it's the opposite of what has come to be known as the Efficient Market Hypothesis, and it's very tempting to see it as a radically foreshortened statement of the Fractal Market Hypothesis put forward by Peters in *Fractal Market Analysis*.

In the Efficient Market Hypothesis (EMH) price changes are noncorrelated (serially independent) from period to period and timeframe to timeframe, with the semistrong version stating that the market's random walk is due to a rational dissemination of all known news and fundamental information uniformly across timeframes. This is radically unlike the fractally dispersed emotional psychology Levine believed in.[34] Moreover, it's certain that he would have seen the mathematics of MIDAS as being inconsistent with standard deviation and the normal distribution curves of the EMH.[35]

In putting forward his Fractal Market Hypothesis in his book *Fractal Market Analysis* Peters[36] argued that prices aren't interpreted univocally across timeframes because only information relevant to a particular time horizon will be judged relevant. In general, technical information will be weighted much more highly in the shorter term. When markets occasionally do weight information equally across timeframes the consequence is a loss of market liquidity, resulting in a market crisis, since longer-term investors either stop participating in the market or else lose faith in fundamental data and trade short-term.[37] Thus, liquidity and market stability cannot be accounted for by the EMH.

Levine was relating the mathematics of VWAP to the fractal ideas inherent in distinctive multi-timeframe market psychology in virtue of the "anchoring" methodology implicit in his hierarchies of support and resistance. Ultimately the formulation of these ideas goes back to Peters's book *Fractal Market Analysis*, though they have their origin too in several earlier studies.

The MIDAS Approach as a Genuine Standalone Trading System

In the introduction to this chapter, I observed that a major weakness in Levine's presentation of the MIDAS approach is his lack of attention to the practical implications of trading with MIDAS curves, whereas the emphasis in this book is very much on practical trading implications. With this in mind, the second half of this chapter lays a little groundwork for what is to come by discussing how the MIDAS approach can be converted into a practical trading system.

In the following discussion, trading system criteria will be outlined alongside a brief discussion of how the fractal nature of MIDAS meets each one. This discussion can also be read alongside the discussion in Chapter 3 of how standard MIDAS S/R curves could be used with relative ease by an advanced beginner or an intermediate-level trader as a standalone day trading system.

Van K. Tharp, PhD,[38] defines a trading system in terms of the following eight components:

1. A market filter
2. Setup conditions
3. An entry signal
4. A worst-case stop loss
5. Re-entry when it is appropriate
6. Profit-taking exits
7. A position-sizing algorithm
8. The possibility of multiple systems for different market conditions

We can omit criterion (7) because it's not so relevant to the present discussion and replace it with the requirement that a system (especially in a context such as day trading) generate sufficient signals throughout the trading day (or timeframe of interest) to ensure that the system is self-reliant; that is, that it doesn't require the input of outside elements to generate an appropriate number of signals. Criterion (7) can be reframed as follows:

7. That the system be capable of generating sufficient signals over the timeframe of interest

A Market Filter

This criterion has to do with how a market is moving (i.e., trending up, down, or sideways) and whether the system works adequately in relation to it. Volatility will also play a role here insofar as markets can be more or less volatile, regardless of their direction.

Detailed studies of how the standard MIDAS S/R curves and the topfinder/bottomfinder algorithm work in various market environments is discussed in forthcoming chapters, while Chapter 14 specifically explores conditions in congested markets and increased volatility. For now, however, we have seen in Figures 1.5 and 1.6 that any market with a Hurst exponent above 0.5 will provide maximum opportunities for MIDAS curves. Figure 1.8 is a 5m chart of the Xetra DAX March 2010 futures illustrating this point. The majority of the movement is captured by the standard S/R curves, while accelerated portions of the trend where the fractal dimension is reduced

FIGURE 1.8 5m chart of Xetra DAX March 2010 futures showing one trading day (March 9) and extensive applications of MIDAS with Granville's OBV in the top pane.

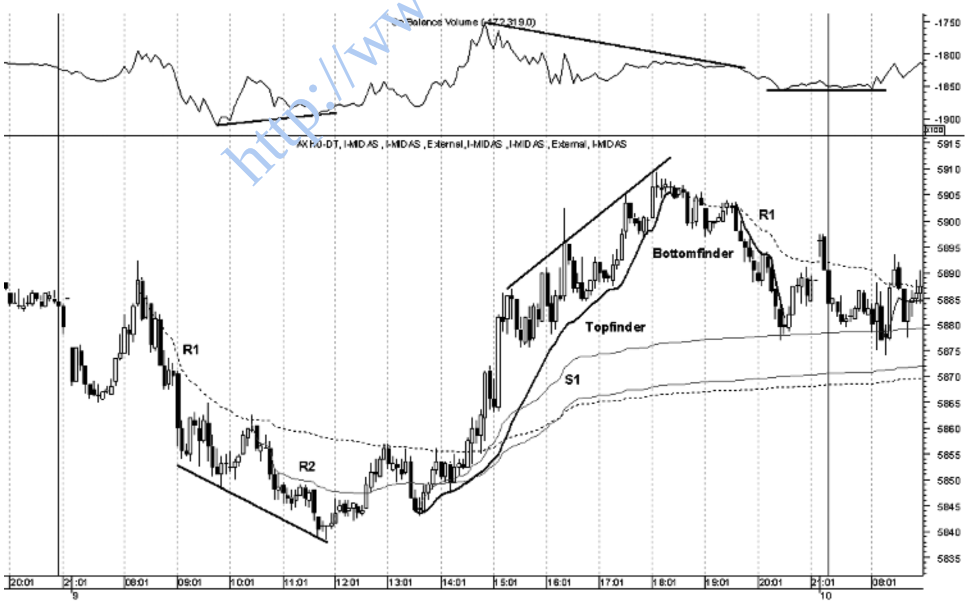

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

P1: OTA/XYZ P2: ABC
JWBT458-c01 JWBT458-Coles March 9, 2011 23:10 Printer Name: Yet to Come

to a minimum are captured by the two topfinder/bottomfinder indicators. The top pane contains Granville's On Balance Volume indicator, which was favored by Levine as providing a background indication (in the form of divergences) of whether standard S/R curves are likely to continue holding price.

The trading opportunities illustrated in Figure 1.8 should again be compared with the single VWAP curve in Figure 1.2 and how it must be bolstered by additional analysis in order to generate a timely market filter. No such additions are required with the MIDAS method.

Setup Conditions

As noted, it's a very significant weakness in the original MIDAS method that Levine never discussed trade-management issues. Typically, a setup would occur with respect to a standard MIDAS curve when price pulls back to it and we establish a pure support/resistance-based contrarian play. For example, if we look again at Figure 1.8 we can see a variety of instances where this occurs. In the case of the topfinder/bottomfinder curves, a setup would be when the curve is launched and it provides a certain cumulative volume prediction while price is trending above or below the curve respectively. The trade-management implications relating to the topfinder/bottomfinder curves are discussed in Chapter 4. As discussed earlier, no comparable setup conditions are available in the case of a standard VWAP curve unless additional indicators and/or chart analysis are brought in.

The Entry Signal

This is another important topic that is discussed in more detail in Chapter 3, but for now it's necessary to concede that Levine fails to meet this criterion in his discussion of the MIDAS method. My own view is that Japanese candlesticks are ideally suited to MIDAS setups insofar as one can use the well-known candlestick reversal signals as filters for price reversals off standard MIDAS S/R curves. Readers should consult Steve Nison's book *Japanese Candlestick Charting Techniques*,[39] especially Chapter 11 ("Candlesticks with Trend lines") and Chapter 12 ("Candlesticks with Retracement Levels"), as a primer for what is being proposed here.[40]

Figure 1.9 is a 3m chart of the same DAX March 2010 contract. A standard MIDAS support curve is already running on the chart from an earlier time, and our interest is in boxes (1) and (2).

In box (1) price reverses on the support curve in a hammer candlestick. By the time of this reversal, it has also been possible to draw a trend line from the low at the arrow, thus strengthening the MIDAS support. An entry signal is subsequently produced when price breaks above the hammer's high. This also coincides with the breaking of the small downtrend line. Where the downtrend line is above the break of the high of the reversal bar, a trader could wait until the trend line has been broken at the cost of a later entry.

FIGURE 1.9 The two boxes highlight how combining MIDAS with Japanese candlesticks produces robust criteria for entry signals.

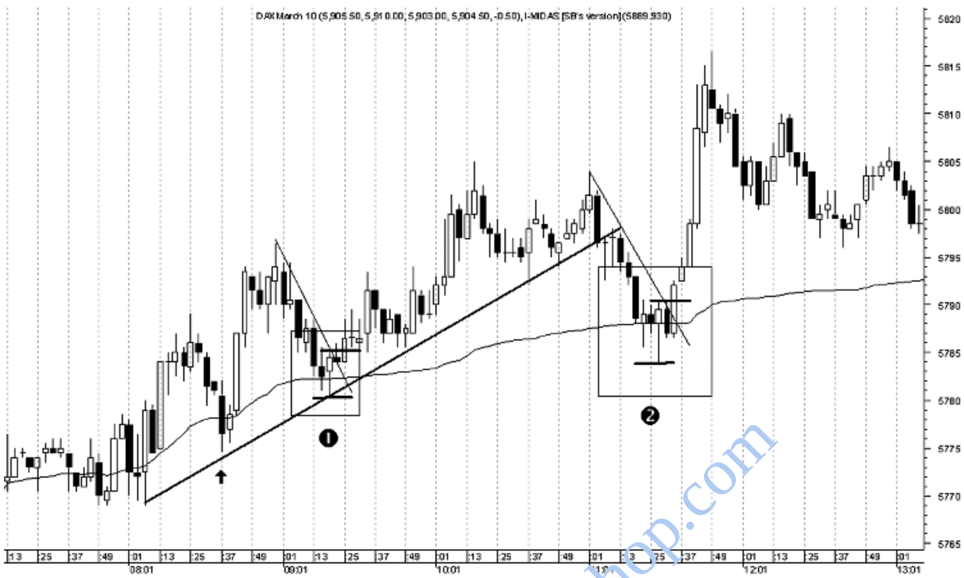

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

In the box marked (2), we have another hammer candlestick with a longish lower shadow. Here the entry signal is the same: price breaks above the hammer's high at the same time as it breaks the downtrend line, thus triggering the entry. For uptrend reversals, of course, the conditions would be reversed.

The topfinder/bottomfinder curves require a more subtle approach as regards an entry signal (see Chapter 4).

The Protective Stop

Again Levine never discussed stop-losses in relation to the MIDAS approach, but I believe that Japanese candlesticks provide the solution. If we go back to the two boxes in Figure 1.9, we see that stop-losses are placed at the low of each hammer reversal candlestick. We can't tighten them further because, as we can see, candlesticks often penetrate MIDAS curves marginally as they respond to them. This also happens when combining Japanese candlesticks with standard trend lines and support and resistance as well as with Fibonacci retracement levels, so it's a common phenomenon a trader must get used to. Again the topfinder/bottomfinder curves require a more subtle approach as regards an entry signal (see Chapter 4).

The Re-Entry Strategy

One of the topics discussed further in Chapter 13 is a phenomenon Levine labeled "porosity." This occurs when price marginally penetrates a MIDAS curve before

P1: OTA/XYZ P2: ABC
JWBT458-c01 JWBT458-Coles March 9, 2011 23:10 Printer Name: Yet to Come

responding to it. Porosity can be very misleading because once price breaks a MIDAS S/R curve the expectation is that it will move to the next most proximate curve. Consequently it can be very easy for a trader to be caught on the wrong side of the market when price does belatedly respond. Figure 1.10 illustrates this phenomenon. At the area circled, price has penetrated the curve while in a steep downtrend. Below the curve, it pulls back very marginally and then forms a doji candlestick before price very slightly turns down. Here a short position might have been taken before price reacts to the curve. If so, the short would have been stopped out on price crossing back through the curve. The first re-entry point is either at the first trend line break above the curve or in the first box when price breaks away after finding support on the curve for several bars. Another re-entry point comes at the second box when, after a long black candlestick, price creates a doji candlestick on the curve and then moves away. In both cases, the entry point and stop-loss follow the same strategy outlined earlier.

As we see later, whether porosity is actually taking place should always be considered alongside a confirming indicator such as Granville's On Balance Volume. In Figure 1.10, for example, porosity, rather than a breakdown, is virtually guaranteed by the fact that the OBV line started diverging positively from price at 9.44am (vertical line), some two hours before the porosity occurred. It's highly likely in such circumstances that the MIDAS support curve will hold price and that the downtrend line will be broken.

FIGURE 1.10 3m chart of DAX March 2010 futures with the trading implications of price porosity.

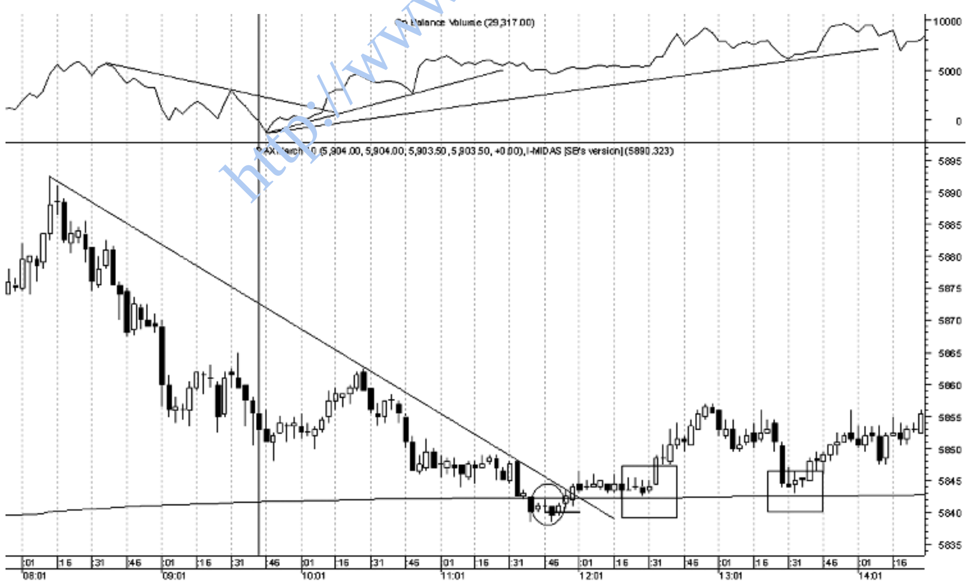

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

Profit-Taking Exits

There are two primary ways that the MIDAS system provides price targets, though it's also perfectly compatible with techniques such as trailing stops or other indicator-based approaches. The first is that, by its very definition, a MIDAS support/resistance curve will confirm the price direction until it is broken. Price also has a tendency to move between MIDAS support and resistance curves once they've been penetrated. These phenomena will be illustrated thoroughly in forthcoming chapters. The topfinder/bottomfinder curves also support and resist price in addition, uniquely, to establishing cumulative volume targets, which are easily converted into price targets. These again are features of the topfinder/bottomfinder curves to be discussed in Chapters 4 and 5.

Sufficient Signals over the Timeframe of Interest

This is an important, if underappreciated, criterion of any system. If we go back to Figure 1.2, we can quickly appreciate that a single VWAP curve running indiscriminately from the start of the day fails to provide anything beyond a superficial idea of market direction and provides little by way of trading opportunities satisfying the criteria we've been considering here. By contrast, as we've seen in Figure 1.8, the fundamental relationship between the fractal nature of the financial markets and the notion of "anchored VWAP" means that MIDAS curves can be launched at all degrees of trend and need to be if MIDAS is to provide a clear perspective on current market direction. The higher resolution analysis of Figure 1.8 can be contrasted with the broader sweep afforded by the curves in Figure 1.11. This is another 5m chart, with MIDAS curves spanning days of price activity as they pick up most of the broader price movement.

In comparing these two charts, the true scale invariance Levine referred to becomes evident; yet because the fractal dimension is less than 1.50 (the Hurst exponent on these several days of data is 0.62 (= 1.38)) MIDAS curves are highly effective. Both Figures 1.11 and 1.8 can be usefully compared with Figure 1.4, especially the last two time series with Hurst exponents above 0.5.

The Occasional Need for More than One Trading System

Obviously the more robust the system, the less this need will arise. Provided markets are trending, MIDAS curves will be applicable according to the criteria already discussed. When trends are accelerating, the topfinder/bottomfinder curves take precedence, as discussed in Chapters 4 and 5. One of the main drawbacks with the MIDAS system is that it was never designed for sideways markets and it would also be useful if a curve could be created to capture the VWAP highs in uptrends and the lows in downtrends. Chapter 14 discusses an indicator to meet these requirements.

FIGURE 1.11 5m chart of the DAX March 2010 futures with lower resolution multiday price activity captured by the broad sweep of the curves.

Source: eSignal and Metastock. www.esignal.com and www.equis.com.

Summary

- With a growing interest in VWAP among the trading community, the boundary between the MIDAS approach and the broader VWAP backdrop is becoming very unclear, especially in the labeling and use of curves. Clarity is required both on the mathematical differences between basic VWAP formulae and the MIDAS formula as well as on the theoretical differences.
- The five main tenets that define Paul Levine's market philosophy help accentuate two theoretical factors that separate MIDAS from mainstream VWAP analysis: (i) the critical choice of where to launch MIDAS curves, and (ii) the use of multiple MIDAS curves.
- The latter is based on the notion of a hierarchy of support and resistance that assumes a fractal interpretation of the financial markets that has in turn been supported by a growing body of empirical statistical studies.
- These studies do indeed show that persistent time series with a Hurst exponent higher than that signified by a pure random walk ($H > 0.5$) are best suited for MIDAS curve analysis. Indeed, the topfinder/bottomfinder curves only work if the fractal dimension of the markets is significantly less than 1.5. See also Chapter 4.

- It's likely that Levine's fractal market hypothesis owes much to the work of Edgar Peters, which in turn is derived from other academic interpretations of price movement in the early 1990s that were incompatible with the Efficient Market Hypothesis.
- Despite the absence of trade-management considerations in practical trading contexts by Levine, the fractal nature of MIDAS justifies the claim that MIDAS (alongside a robust charting method such as Japanese candlesticks) can be elevated to the status of a genuine trading system.