Algorithmic Trading

Algorithmic Trading Strategies

Volume 32
Step-By-Step System Making

Research Paper Series by
Fudancy Research Group
& M. Schoeffel

http://Fudancy.com

Algorithmic Trading Strategies

Step by step system making

Booklet series of the

Fudancy Research Group

VOLUME 32

M. Schoeffel (Nyon, Switzerland)

Algorithmic Trading Strategies

On the web @ fudancy.com

ABOUT

Fudancy Technology is a R&D entity focusing on coaching hedge funds and financial institutions and in developing turnkey trading systems. We believe it's not about who we are. Research, performance and robustness of the systems are what matters. That said, at Fudancy Technology, we all hold PHD degrees and we love what we do.

--

TABLE OF CONTENTS

Introduction

I. Introduction to algorithmic finance

II. Beyond randomness

III.a Building equations – Principles

III.b Building equations – Inverse cubic law

III.c Building equations – General view (small time scales)

IV. Unfolding

V. The main strategies and their variations

Moving averages

Trading range breakout

Filter rule

Bootstrapping

Stress tests

REFERENCES

Introduction

This booklet is a grand summary of the paper series [1-2]. It completes a former version [1] (volume 17), where ideas were described based on the first 16 volumes of the booklet series. At the present stage, we are more advanced in the understanding of how markets evolve at large and small time scales. We have also presented more advanced concepts in statistics. Therefore, a new completed version of [1] (volume 17) makes sense. Also, we present all ideas developed in [1-2] within the perspective of practical system making. This booklet makes a complete reference with clear milestones for system makers.

I. Introduction to Algorithmic Finance

As long as financial markets and data have existed, people have tried to forecast them, in the hope that good forecasts would provide a reasonable anticipation of future events, within a given error margin. Even in early 2012, the issue of forecasting is still essential. For example, the European Central Bank (ECB) is holding a dedicated workshop on forecasting techniques in spring 2012. As we can read from the announcement of the workshop, *the goal is to provide a forum for the presentation of both theoretical and applied contributions that can help identify new directions for forecasting, particularly in the light of the issues raised by the recent financial and sovereign debt crises. The ongoing turbulence in the global economy, and especially in financial markets, requires a critical assessment of existing methods and the development of new paradigms in order to provide reliable analytical support for macroeconomic policy-makers*. Then, the ability to forecast future directions in the evolution of any financial or economic indices, at macro or micro levels, appears more than ever a key feature in modern societies.

Interestingly, in a pure financial practice, what matters most is not the question whether it is possible to forecast, but how the future path of a financial time series can be forecasted. At a theory level, however, it is merely the question whether series of speculative prices can be forecasted than the question how to forecast. Therefore practice and academics have proceeded along different paths in studying financial time series data.

In the following, we will not make the distinction between practice and theory: **we define “algorithmic finance” as the quantitative analysis of financial time series with the purpose to forecast statistically the future evolution of this time series.**

We do not even enter into the standard opposition that arises from the so-called fundamental and technical analysis. In fact, we can imagine algorithms based on fundamental variables. More generally, invariants in algorithmic finance can be defined at fundamental or technical levels. The key feature in our approach is that it must be quantitative. This means that a set of rules are proposed at the end of the analysis and these rules can be used as an algorithm to extract results, either with the help of a computer based program or analytical calculations. Practical examples are described extensively in the next sections.

Let us introduce the notations used in this booklet. The financial data series is labeled as: ***x(t),x(t+dt),…*** where dt is the time step (or time unit) we consider. As we are only interested into intra-day dynamics of market prices, we shall only consider time units of the orde of a few minutes.

Algorithmic finance is the quantitative analysis of this set of data x(t) with the goal to forecast in average at least a part of the evolution of x(t). The fact that this problem has a solution in general is obviously not granted. We discuss in following sections general search methods needed to identify significant deviations from randomness in the data.

II. Beyond randomness

Identifying the presence of non-randomness effects in the data x(t) means that we need to be able to disentangle background (randomness) with respect to signal(s) (non-randomness). In the latter case, when signals are identified, we are obviously in a situation where forecasting the evolution of the time series is possible. We can then produce anticipations and build trading algorithms. In [1], we have defined a general search mechanism that can be used quite generically. In this section, we consider an example to illustrate how this works.

The separation of signal and background postpones that we have the ability to define what is background and what is a signal in the time series x(t). This is the first moment of the analysis where imagination and ideas come in. In a next section, we define more precisely the mathematical contents of equations in algorithmic finance. Here, we assume that two basic equations are known:

(a) Background hypothesis: **dx = Sigma dW** where dW is a random walk (Wiener process of Brownian motion) and Sigma the volatility of the process. This is the null hypothesis, where x(t) behaves as a pure background or equivalently a random walk.

(b) Signal hypothesis: **dx = A(x-Rx) + Sigma dW** where the second term is described in (a) and the first term represents the signal we intend to search for in x(t). In this equation, Rx is the price range and "**A**" is a coefficient that quantifies the deviation with respect to the pure random walk. When A is significantly non-zero and positive, this means that prices are expulsed from the range Rx. Then, we tend to have a range breakout phenomenon. On the contrary, when A is negative, this means that prices have a tendency to be contained within the range Rx. The signal hypothesis is thus what we can call a range breakout mechanism.

The imagination of the system maker comes within the expression (b). The idea is to test a range breakout mechanism in x(t) and we need to know how to build this idea into an equation. We could start from a very different idea and also obtain reasonable results at the end. One idea does not exclude others. However, the important point is that we need an input from our own views.

In [1] we have then explored this idea based on range breakout for intra-day data on commodity futures based on (b). We have proposed a simple generic algorithm that translates the idea into a well defined procedure:

i. Define a range Rx on a given time unit;
ii. Search for a range contraction or expansion;
iii. Use the item (ii) as a first trigger to search for an expulsion of the range. This means buy when the price is above the highest highs of the range and sell is the price falls below the lowest lows of the range. We can add an offset to the previous requirement depending on the tick size;
iv. In addition, define stop loses and profit limits (fixed or trailing) that protect and optimize the strategy.

We have shown explicitly that this idea and consequently the algorithm based on it provides very good results on different commodity futures, like Natural Gas, Platinum, Oil and Silver [1]. We remind some of these trading algorithms in further sections.

In [1] (volume 26), we have shown how random-walks can perfectly mimic market movements and technical analysis figures. In particular, the display below is interesting:

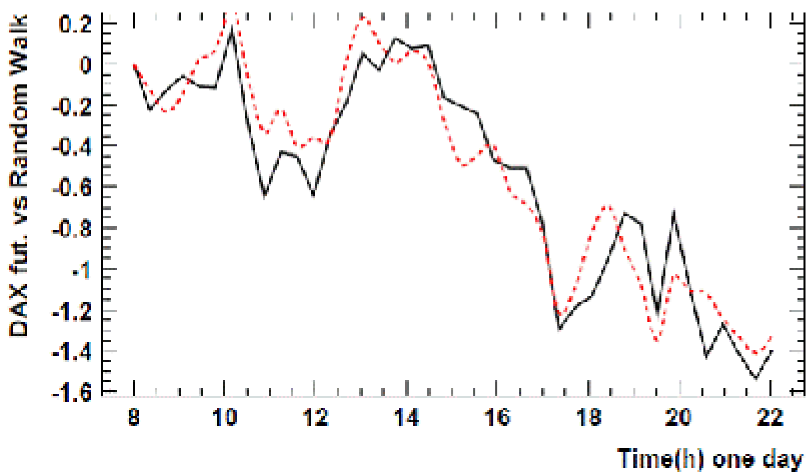

This correspond to DAX future contract on a one day basis (full line) compared to a random walk (dashed line) generated for the same sampling and period of time. **We observe that the movement of the price is essentially random** as it can be described as such.

This shows in a clear and simple way the necessary rigorous treatment of data on a statistical basis to reject with high confidence level potential fake moves.

III.a Building equations - Principles

As explained in the previous section, it is an important task to build and write equations starting from ideas on the behavior of the market x(t). In this booklet, we are not interested by the academic formalism or any *ad hoc* theoretical formulations. On the contrary, we want to show how our intuitions and ideas can be put into formal terms and embedded into market equations.

Formalism in itself is meaningless. In other words, if we take an equation from a reference paper and try to use this equation with no input from our own views, this cannot work. We need to focus on the ideas and then and only then build the formalism to fit with our views.

We start from the most basic equation, a pure random walk plus a drift. Then, the time series defined above **x(t)** follows:

x(t)-x(t-dT)=dx=mu+Sigma dW(t) where **mu** is the mean return of prices, **Sigma** the standard deviation and W(t) is Wiener (Gaussian) process [1]. We can express dW(t) quite simply as:

dW(t) = dT$^{1/2}$ G(t) where **G(t)** is a Gaussian variable of mean zero and unit variance. To a large extend, this corresponds to what we observe in most financial data [2]: a price increment of the data (dx) follows approximately a random walk (Sigma G(t)). In general, for small time units at intra-day level, the drift can be neglected. This equation is what we call in general the background hypothesis, or the default hypothesis (no signal). At intra-day time scales (mu=0), the above equation reads:

dx=Sigma dW(t)= Sigma dT$^{1/2}$ G(t)

The probability distribution function (PDF) associated to the price increment dx(t)=x(t)-x(t-dT) is then a Gaussian distribution:

P[dx]=1/(2 pi Sigma^{2}dT)$^{1/2}$ exp[-1/[2Sigma^{2}dT] dx^2]

Then, modifications to this basic equation can be done.

i. If we want to study how a moving average MA(x) reacts with respect to the price (expulsion or attraction), we can write: **dx=k[x-MA(x)]+Sigma dW(t)**. For positive k, this means that the MA(x) tends to expulse the price and inversely for negative values of k. Of course, k needs to be significantly non-zero for the effect to be visible and not hidden by the volatility-stochastic term 'Sigma dW'.

ii. If we are interested on the crossing of 2 moving averages MA1(x) and MA2(x) of different lengths and how this crossing affects price increments, we can write: **dx=k'[MA1(x)-MA2(x)]+Sigma dW(t)**. The reasoning is exactly the same as above. Again, the respective orders of magnitude of terms ||k'[MA1(x)-MA2(x)]|| and ||Sigma dW(t)|| determine the visibility of the effect of the MA crossing.

iii. As shown above, the trading range breakout idea can be encoded very easily into an equation.

Till now, we have considered the stochastic term as a pure random walk in writing, **dx=x(t)-x(t-dT)=Sigma dW(t)**, with W(t) a Wiener Gaussian process: $<dW^2>=dT$. In this case, increments are completely independent. It is possible to relax this hypothesis by considering more sophisticated stochastic term which verifies: $<dW_H^2> = dT^{2H}$ with H=0.5 in case of pure random walk and H non equal to 1/2 in the generalized case. The interest is that it is then possible to encode novel effects like persistence or mean-reverting phenomena into the equation of motion for x(t) even if only the stochastic term is present into the equation.

We can show that:

$<dW_H(t+h)dW_H(t)>=2H(2H-1)h^{2H-2}$ for large h.

Using standard definitions, H is called the Hurst exponent and the case W[H.eq.1/2](t) is called a Brownian Motion (BM) whereas the case W[H.ne.1/2](t) is called a fractional Brownian Motion (FBM). From the formula above, we conclude immediately that **H>1/2 corresponds to persistence:** a positive (negative) increment is most likely to be followed by a positive (negative) one, which gives a trend following feature to the time series $W_H(t)$. On the contrary, **H<1/2 corresponds to anti-persistence:** a positive (negative) move is most likely to be followed by a negative (positive) one, which gives a mean reverting feature to the time series $W_H(t)$.

The equation of motion: **$dx=x(t)-x(t-dT)=\text{Sigma}\ dW_H(t)$** can therefore be used to test novel effects in the data. For example, we can test the hypothesis of persistence effects by testing the parameter space H>1/2 on the financial data series. The search can be done following statistical methods described in references [1]. A similar search can be followed for mean-reverting effects with H<1/2 parameter space.

We can generalize even further the above equation. From the previous formalism we can write:

$||x(t)-x(t-dT)|| = O(dT^H)$ where $||u||$ represents a norm like the average of the absolute value: $<|u|>$. Then, we define the moment of order q of the time series x(t): **$K_q(dT) = <|x(t)-x(t-dT)|^q> / <|x(t)|^q>$.**

As a first approximation, using the previous formalism, we can write: $K_q(dT) = O(dT^{qH})$. In fact, a more general form of this result can be considered with the formula: $\mathbf{K_q(dT) = O(dT^{qH(q)})}$ where **H:=H(q) is a priori a function of q**. It is called the generalized Hurst exponent.

The interest is that using moments of any orders q, **we can filter out the type of fluctuations of the time series x(t)**: large values of q filter out to large size fluctuations and low values of q filter out small size fluctuations.

The idea is very simple: from the data series, we extract $K_q(dT)$ for different values of q. If H(q) deviates from 0.5 for small values of q (for example), we can conclude that small fluctuations in the series x(t) are not compatible with a pure BM. In such case, small size fluctuations are correlated. Also, if we find that H(q)=0.5 for large values of q (for example again), we can conclude that large fluctuations are compatible with a pure BM and thus uncorrelated.

In addition, we have shown in [1] (volume 31) that the exponent H is not only a measure of persistence but also can be interpreted as a measure of the mood on the market. Indeed, as H<1/2 characterizes the anti-persistent behavior, a decreasing tend on H can be seen as an increasing nervousness on the market. Similarly, an increasing H can be seen as a support of the trend that has just started. This is a very interesting property, derived from our previous knowledge. It can bring clearly new ideas in trading strategy making. For example, we can use an increasing H as a trend seeker or inversely a decreasing H as a turning point in a market environment. More precisely, we can write typical conditions of turning points:

(1) Decreasing trend on H

(2) $H<0.5$ and typical moving averages on $H<0.5$

Do not forget also the general 'theorem' we have established in [1] (volume 21) concerning the values of H that we expect to be close to 0.5 in general. Regarding this proposition, any significant deviation would be a clear trigger that something new happens.

III.b Building equations – Inverse cubic law

Returns in financial time series are the most fundamental inputs to quantitative finance. To a certain extent, they provide some insights in the dynamical content of the market. In this booklet, we consider several financial series on futures, using always a five minutes sampling.

Each future contract is characterized by a price series S(t), from which we extract the log-returns $x(t) = \log[S(t + 1)/S(t)]$. The analysis is then driven on these log-returns x(t).

From standard equations [1] (volumes 4-9) we know that the distribution of log-returns x(t), namely P(x), can be written quite generally as

$$P(x) = \frac{1}{Z}\exp(-\frac{2}{D}w(x)/2)$$

where w(x) is an objective function and Z a normalization factor. In particular, it can be shown easily that w(x) can be derived by minimizing a generating functional F[w(x)], subject to some constraints on the mean value of the objective function. It reads

$$F = \int dx P(x) \left[\log P(x) + w(x)/D - \lambda\right]$$

where lambda is an arbitrary constant.

In addition, the expression given in Eq. (1) for the probability distribution can also be seen as the outcome of an equation of motion for x(t). From Eq. (1) and (2), we can express the stochastic process x(t) as in references [1] (volumes 4-9)

$$\frac{dx}{dt} = f(x) + g(x)\epsilon(t)$$

where epsilon(t) (e(t)) is a Gaussian process satisfying < e(t)e(t′) >= D.delta(t− t′) and < e(t) >= 0. Functions f and g depends only on x(t). The distribution function P(x, t), associated with this equation of motion is given by the following Fokker-Planck equation (see [1])

$$\frac{\partial P(x,t)}{\partial t} = \frac{\partial^2}{\partial x^2}\left[\frac{D}{2}g^2(x)P(x,t)\right] - \frac{\partial}{\partial x}\left[f(x)P(x,t)\right]$$

We can finally extract the stationary solution for P(x)

$$P(x) = \frac{1}{Z}\exp\left[-\frac{2}{D}\int dx \frac{Dg\frac{dg}{dx} - f}{g^2}\right]$$

Whether the above equations can be related to real data on financial markets is not granted.

Therefore, we need to compare predictions derived from these equations to real data. As mentioned above, we use financial time series on different futures, using a five minutes sampling. **See also [1] volumes 7 and 9.**

In the above equations, functions f and g are not specified and any choice can be considered. Obviously, only specific choices will have a chance to get a reasonable agreement with real data. For example, let us consider the different cases below:

(i) If $f(x) = -x$ and $g(x) = 1$, we obtain $P(x)Z = \exp(-x^2/D)$, and thus we predict a Gaussian shape for the log-returns distribution.

(ii) In the more general case where $f(x) = K\, g\, dg/dx$ and g is not constant, we obtain

$$P(x)Z = \frac{1}{g^{2(1-\lambda/D)}}$$

and thus we predict non-Gaussian shape for the log-returns distribution.

Using a specific functional form for g(x), we can get

$$P(x)Z = \frac{1}{(1+\frac{x^2}{\nu})^{(\nu+1)/2}}.$$

In this case (ii) P(x) follows the so-called t-distribution. It **depends on one parameter nu to be fitted on real data**, for normalized log-returns.

In Fig. below, we present the log-returns x (**five minutes sampling**) for a large set of futures. To make the comparison, we have scaled x for all futures to the volatility of the DAX future (FDAX).

On the top of Fig. below, we observe that futures on DAX, Bund, Yen, Euro, Gold present the same probability distribution for x, hence P(x) is universal for all these data series once the volatility is normalized to the same value. On the bottom of Fig. below, we provide comparisons with futures on commodities. We observe that data series on CL (Crude-Oil) follows the same P(x) as FDAX, but other data series on Wheat and NG (Natural Gas) exhibit some larger tails.

We can use results developed in the previous section in order to compare with experimental (real) distributions P(x).

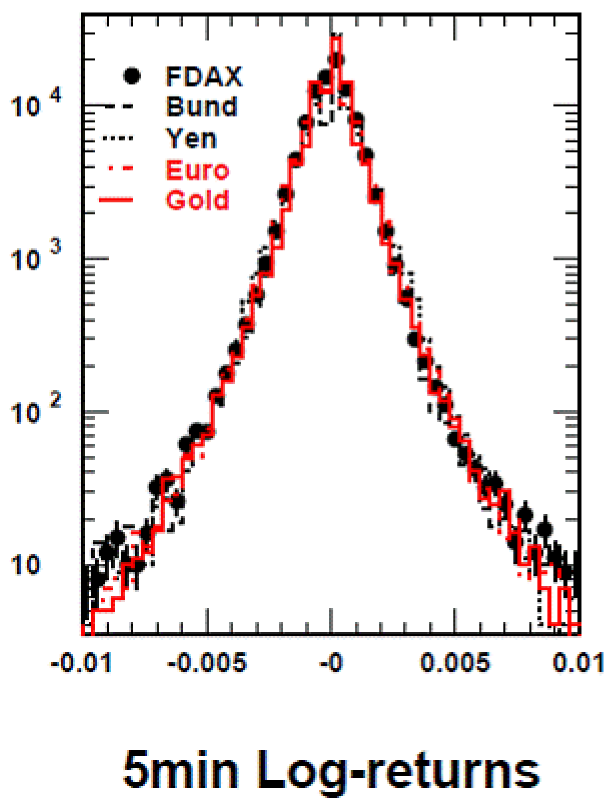

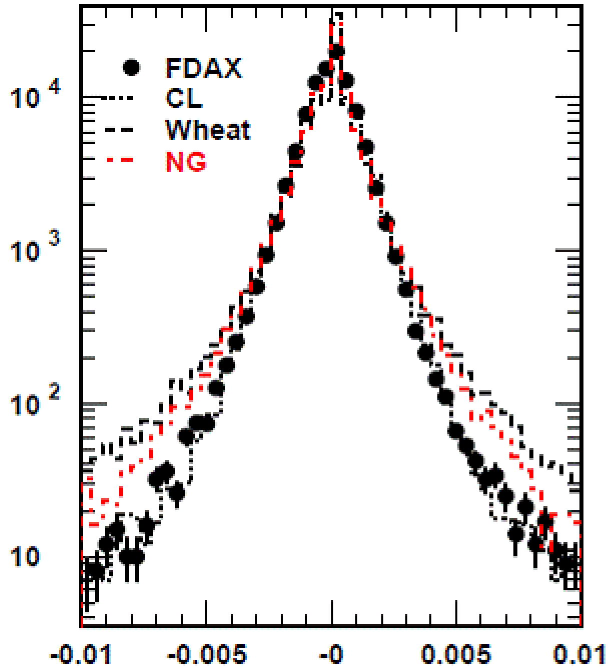

5min Log-returns

Above we display the distributions of Log-returns x (five minutes sampling) for a large set of futures. To make the comparison, we have scaled x for all futures to the volatility of the DAX future (FDAX). Top: we observe that futures on DAX, Bund, Yen, Euro, Gold present the same probability distribution for x, hence P(x) is universal for all these data series once the volatility is normalized to the same value. Bottom: we observe that data series on CL (Crude Oil) follows the same P(x) as FDAX, but other data series on Wheat and NG (Natural Gas) exhibit some larger tails.

Distributions above can be fitted we obtain nu=3 and P() as:

$$P(x) \propto \frac{1}{(1+\frac{x^2}{\nu})^{(\nu+1)/2}} \quad , \quad \nu \simeq 3.$$

This formula holds for 5 minutes sampled data for all futures contracts studied! This result is thus quite universal for small intra-day time scales

III.c Building equations – General view (small time scales)

In reference [1] (volume 23), we have extended this argument to general functional forms. Let us write e(t) as a Gaussian process satisfying < e(t)e(t′) >= D.delta(t− t′) and < e(t) >= 0. Then a general expression for a stochastic variable can be:

dx/dt = f(x) + g(x) e(t)

Here, functions f and g depends only on x(). The distribution function P(x,t) associated with this equation of motion is given by the following Fokker-Planck equation (as discussed above):

$$\frac{\partial P(x,t)}{\partial t} = \frac{\partial^2}{\partial x^2}[\frac{D}{2}g^2(x)P(x,t)] - \frac{\partial}{\partial x}[f(x)P(x,t)]$$

With specific choices of f() and g() simpler than D1() and D2() expressions above, we have been able to describe the data using a very simple function PDF (with nu=3) for a large variety of future contracts (see previous section).

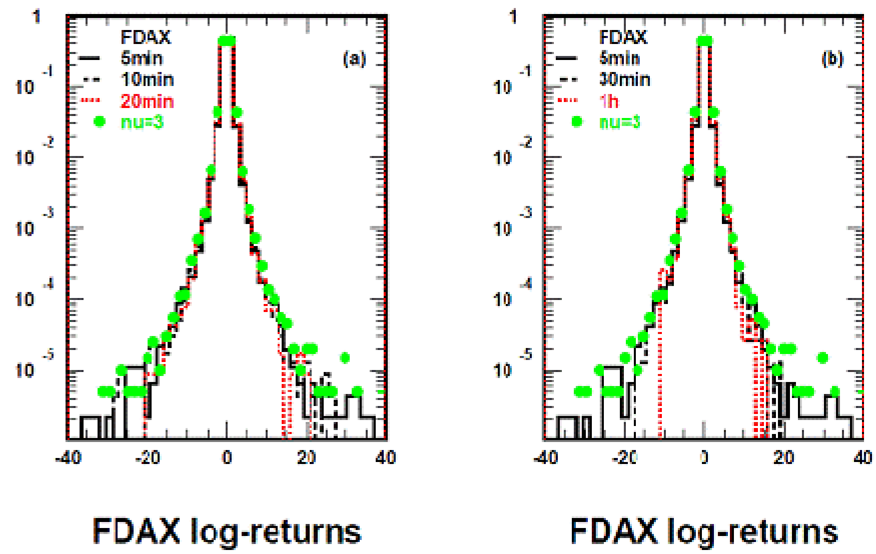

We also observe of the figure above that small time scales mean below 1 hour (typically), which is in good agreement with the values extracted from the factorial moment analysis [1] (volume 20).

Let us consider the inverse formulation of the problem. We start from an equation of motion for the stochastic variable x() as:

$$dx/dt = -\text{gamma}\; x + x\; k(t) + e(t)$$

where is gamma a positive constant, k(t) and e(t) white noise (BM), of zero average and standard deviations: $<k^2>=M^2$ and $<e^2>=A^2$. This equation is more general that in the previous sections as we consider 2 uncorrelated BM here.

In previous references [1] (volumes 2-6), we have shown how to start from market equations to extract practical trading strategies! Thus, the inverse problem of this section is interesting and efficient to solve.

Generalizing the expressions given above, we can write the associated Fokker-Planck equation as:

$$dP(x,t)/\; dt = -d/dx[\; \text{gamma}\; x\; P\;] + 0.5\; d^2/dx^2[\; (A^2+M^2x^2)/2\; P\;]$$

If $M^2=0$, we do not have k() term in the equation, we find a Gaussian process and if gamma is also null, then we find back the standard Wiener process.

In practice, if we write an index q as: **q:=(1+2M²/gamma)/(1+M²/gamma)**, the data at small times scales correspond to **q around 1.7**. This is a conclusion we can make from the figure above, when fitting data with the equations above. **This gives typical values for effects beyond the Wiener process at small time scales (from few minutes resolution to 1 hour)**, which generalizes also the expressions for D1 and D2 given above.

IV. Unfolding

In principle, unfolding is a statistical tool which allows extracting a true distribution from a reconstructed (visible) one. In our approach, we have generalized this view to algorithmic trading with the simple equivalence: reconstructed == sequence of prices, true == equity (sequences of trades of the system).

This is an interesting way to present a trading system (formally). In this booklet, we show very simply how to push this idea into practice and optimization of trading strategies.

Let us define the trading system as: P(t) -> E(t), where P(t) are the prices and E(t) the cumulative equity based on the trade sequence. The arrow represents the triggers of the system: time unit, trigger decisions, money management etc.

In the transformation (unfolding) P(t) -> E(t), we expect to transform a Hurst exponent on prices of about ½ to a Hurst exponent on the equity above ½. Then, let us rewrite: P(t) [H=1/2] -> E(t) [H'>1/2]. This is the trivial aim of the trading system: obtaining a better persistence once the strategy is built than before any trigger decisions.

First, this can be checked directly once a strategy is established and this can be used as a rejection item.

Second, nothing prevents us to follow the procedure iteratively, namely:

P(t) [H=1/2] -> E(t) [H'>1/2] -> E'(t) [H">H'>1/2] -> ...

V. The main trading strategies and their variations

The main trading strategies are described as well as the necessary conditions needed to ensure the qualification of these strategies on any financial data series. As explained in sections II and III, this is the second step of the system making process. First, the idea and the search for any possible signals inside x(t) based on this idea. This first part can rely on model building from equations discussed in section III. Once the idea has been identified, for example repulsion of the prices by a moving average or trading range breakout, the system making process follows with the strategies as encoded in this section (second step). The third and fourth steps are then the validation of the strategy and stress tests. These steps are also discussed below.

Moving-average

Moving-average (MA) trading rules are the most commonly used and most commonly tested technical trading strategies. Moving averages are recursively updated averages of past prices. They yield insight in the underlying trend of a price series and also smooth out an otherwise volatile series. In this booklet we use equally weighted moving averages

$$MA_t^n = \frac{1}{n}\sum_{j=0}^{n-1} P_{t-j},$$

where MA(n)(t) is the moving average at time t of the last n observed prices. Short (long) term trends can be detected by choosing n small (large). The larger n, the slower the MA adapts and the more the volatility is smoothed out. Technical analysts therefore refer to a MA with a large n as a slow MA and to a MA with a small n as a fast MA.

MA trading rules make use of one or two moving averages. A special case is the single crossover MA trading rule using the price series itself and a MA of the price series. If the price crosses the MA upward (downward) this is considered as a buy (sell) signal. The double crossover MA trading rule on the other hand uses two moving averages, a slow one and a fast one. The slow MA represents the long run trend and the fast MA represents the short run trend. If the fast MA crosses the slow MA upward (downward) a buy (sell) signal is given. The signal generating model is given by

$$\begin{aligned} Pos_{t+1} &= 1, && \text{if } MA_t^k > MA_t^n \\ Pos_{t+1} &= Pos_t, && \text{if } MA_t^k = MA_t^n \\ Pos_{t+1} &= -1, && \text{if } MA_t^k < MA_t^n, \end{aligned}$$

where k < n and Post+1 = −1, 0, 1 means holding a short, neutral respectively long position in the market in period t + 1.

Thus, moving average cross-over (MAC-O) rule compares a short moving average to a long moving average. The MAC-O rule tries to identify a change in a trend. There are two categories of the MAC-O rule: variable length moving average (VMA) and fixed length moving average (FMA). The FMA stresses that the returns for a few days following the crossing of the moving averages should be abnormal. The VMA generates a buy (sell) signal whenever the short average is above (below) the long average.

In this booklet, we call the single and double crossover MA rules described above, the basic MA trading rules. These basic MA rules can be extended with a per cent-band filter, a time delay filter, a fixed holding period and a stop-loss. The per cent-band filter and time delay filter are developed to reduce the number of false signals. In the case of the per cent-band filter, a band is introduced around the slow MA. If the price or fast MA crosses the slow MA with an amount greater than the band, a signal is generated; otherwise the position in the market is maintained. This strategy will not generate trading signals as long as the fast MA is within the band around the slow MA. The extended MA model with a b per cent filter is given by

$$
\begin{aligned}
Pos_{t+1} &= 1, && \text{if } MA_t^k > (1+b)MA_t^n \\
Pos_{t+1} &= Pos_t, && \text{if } (1-b)MA_t^n \leq MA_t^k \leq (1+b)MA_t^n \\
Pos_{t+1} &= -1, && \text{if } MA_t^k < (1-b)MA_t^n.
\end{aligned}
$$

According to the time delay filter a signal must hold for d consecutive days before a trade is implemented. If within these d days different signals are given, the position in the market will not be changed. A MA rule with a fixed holding period holds a long (short) position in the market for a fixed number of f days after a buy (sell) signal is generated. After f days the market position is liquidated and a neutral market position is held up to the next buy or sell signal. This strategy tests whether the market behaves different in a time period after the first crossing. All signals that are generated during the fixed holding period are ignored. The last extension is the stop-loss. The stop-loss is based on the popular phrase: Let your profits run and cut your losses short. If a short (long) position is held in the market, the stop-loss will liquidate the position if the price rises (declines) from the most recent low (high) with x per cent. A neutral market position is held up to the next buy or sell signal.

Trading range break-out

Our second group of trading rules consists of trading range break-out (TRB) strategies, also called support-and-resistance strategies. The TRB strategy uses support and resistance levels. If during a certain period of time the price does not fall below (rise beyond) a certain price level, this price level is called a support (resistance) level. According to technical analysts, there is a battle between buyers and sellers at these price levels.

The market buys at the support level after a price decline and sells at the resistance level after a price rise. If the price breaks through the support (resistance) level, an important technical trading signal is generated. The sellers (buyers) have won the battle. At the support (resistance) level the market has become a net seller (buyer). This indicates that the market will move to a subsequent lower (higher) level. The support (resistance) level will change into a resistance (support) level. To implement the TRB strategy, support- and-resistance levels are defined as local minima and maxima of the closing prices. If the price falls (rises) through the local minimum (maximum) a sell (buy) signal is generated and a short (long) position is taken in the market. If the price moves between the local minimum and maximum the position in the market is maintained until there is a new breakthrough. The TRB strategy will also be extended with a per cent-band filter, a time delay filter, a fixed holding period and a stop-loss. The basic TRB strategy, extended with a per cent-band filter, is described by

$$
\begin{aligned}
Pos_{t+1} &= 1, && \text{if } P_t > (1+b)\max\{P_{t-1}, P_{t-2}, \ldots, P_{t-n}\} \\
Pos_{t+1} &= Pos_t, && \text{if } (1-b)\min\{P_{t-1}, \ldots, P_{t-n}\} \leq P_t \leq (1+b)\max\{P_{t-1}, \ldots, P_{t-n}\} \\
Pos_{t+1} &= -1, && \text{if } P_t < (1-b)\min\{P_{t-1}, P_{t-2}, \ldots, P_{t-n}\}
\end{aligned}
$$

Filter rule

The final group of trading strategies we test is the group of filter rules, introduced by Alexander (1961). These strategies generate buy (sell) signals if the price rises (falls) by x per cent from a previous low (high). We implement the filter rule by using a so called moving stop-loss. In an upward trend the stop-loss is placed below the price series. If the price goes up, the stop-loss will go up. If the price declines, the stop-loss will not be changed.

If the price falls through the stop-loss, a sell signal is generated and the stop-loss will be placed above the price series. If the price declines, the stop-loss will decline. If the price rises, the stop-loss is not changed. If the price rises through the stop-loss a buy signal is generated and the stop-loss is placed below the price series. The stop-loss will follow the price series at x per cent distance. On a buy (sell) signal a long (short) position is maintained. This strategy will be extended with a time delay filter and a fixed holding period.

Bootstrapping

Bootstrapping is a technique wherein generally a new series of asset prices are created by the random reordering of the original series. In general, we consider 300 random arrangements of the original series as it is sufficient to prove statistical significance. Moving average rules, range breakout triggers and filter rules are applied to each of the new series. The profits of the original series are compared to the distribution of results from the numerous random series.

In this method no assumptions are made about the volatility of the assets over time, or about the distribution of the asset prices. Only the price changes from data are randomly reshuffled keeping the exact same starting and ending values of the time series so that each of the new series will have identical distributional properties to those of the original series. In particular, a simple buy and hold strategy over the full period provides the same profit or loss for all data and pseudo-data series.

Once each new random series is generated, the technical trading rules are applied and profits calculated. This now creates an empirical distribution of profits on which the original profits can be compared and measured. If profits from the original series are above a percent cut off level then they can be stated as significant. The null hypothesis in this situation is that there is no information for making excess returns in the original time series. This is rejected of course if the profits on the original series are significant compared to the distribution of profits.

More precisely (see also [1]): the process described above generates an empirical distribution of profits. The profits calculated on the original data sets are then compared to the profits from the randomly generated data sets. A simulated p-value is produced by computing the proportion of returns generated from the simulated series that is greater than the return computed with the actual series.

The null and alternative hypotheses are given by: (H0) the trading rules provide no useful information; (H1) the trading rules provide useful information.

The resulting p-value from the bootstrapping simulation reads as follow: if the original return has a rank of 100, then the return is the highest of any of the randomly generated returns, and has a corresponding p-value of 0.00. A rank of fifty reveals that half of the randomly generated returns were greater than the original return, resulting in a p-value of 0.50... In general, we consider as reasonable strategies, those where p-values are below 0.05 (5 per cent). Then, we consider that (H0) can be rejected and (H1) validated. In strict statistical terms, the last statement is not correct as detailed in [1]. However, we are much interested by the practical use of concepts than the pure academic (sterile) formulations.

Stress tests

As discussed in [2], robustness testing needs to be performed to mitigate the effects of data mining and to further analyze the significance of the trading rule profits. To test the returns for robustness, returns will be calculated on three sub-periods of the original data. The sub-periods are determined by arbitrarily dividing the data sets into thirds and then testing for structural breaks between the subsets. The Chow Test can be used to test for structural breaks [2]. The subsets are used to test for robustness if the parameters of each subset are determined to be non-stationary. Three new subsets are selected if the parameters of the subsets are constant. The returns from each trading rule and the buy-and-hold strategy, along with the Sharpe ratio, are computed for each sub-period. Consistent excess returns and stable Sharpe ratios across the sub-periods are associated with robust returns. In addition, a large number of stress tests on parameters of the strategy can be thought of in order to ensure its internal robustness. See [2] for details.

REFERENCES

[1] Algorithmic Trading - Algorithmic Trading Strategies – **volume 1-31** [Kindle Edition], **Fudancy Research**.

http://www.amazon.com/s/ref=ntt_athr_dp_sr-_1?_encoding=UTF8&search-alias=digital-text&field-author=Fudancy%20Research

[2] Trading Systems: How to Outperform Markets Using Algorithmic Systems [Kindle Edition], **Fudancy Research**.

http://www.amazon.com/Trading-Systems-Outperform-Algorithmic-ebook/dp/B006YKF7SA/ref=sr_1_2?s=digital-text&ie=UTF8&qid=1329813631&sr=1-2