The EMA In Practice

© Malcolm Moore
14-Dec-2003

Comparing an EMA with a SMA

The graph below shows a 100-unit step input being applied to both an SMA20 and an EMA20 filter, and the two outputs are clearly seen. From the step input, the SMA20 output rises as a ramp till it hits the maximum value – just like a slew rate limited amplifier!

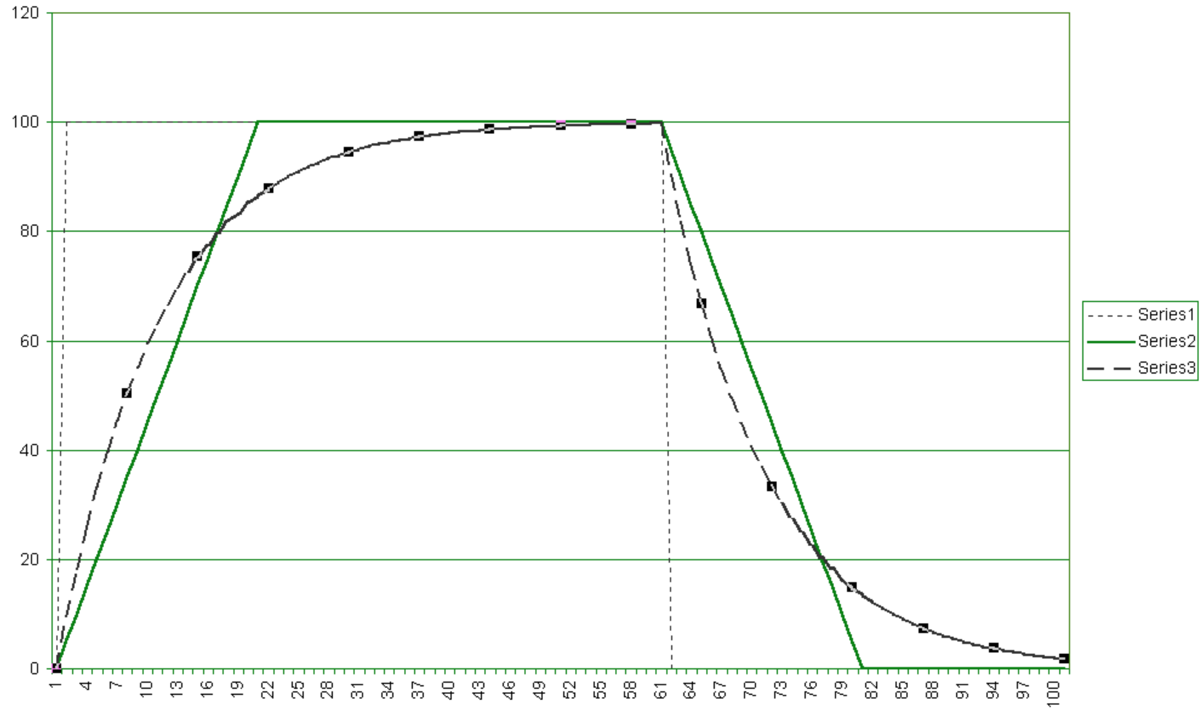

The EMA20 rises more quickly then falls off to exponentially to asymptotically converge on the stable output. ***The two outputs cross over at the 80% mark, and this a reference to be used when comparing a myriad of other responses.***

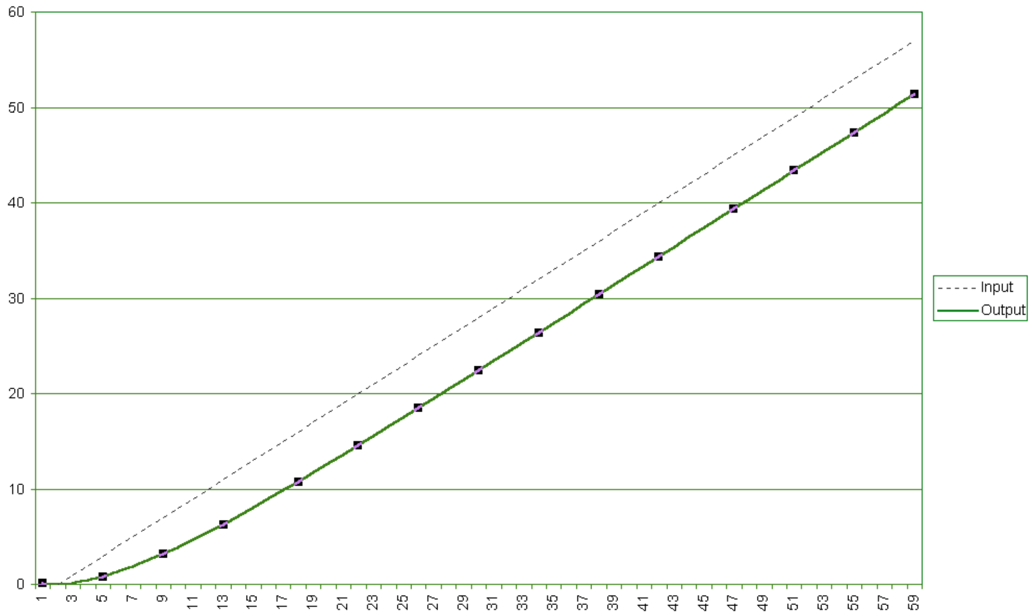

The graph above shows an IIR filter response to a unit ramp (one vertical position per horizontal step). (This could be looked at as say 1 cent per day.) This time k = 0.15 so the Periods = 12.33333, and the Time Constant (T) is therefore 6.166667 Periods.

The Unit Ramp is the straight dotted thin positive sloping line and under that is the thick line output response to the ramp, which also takes off and becomes asymptotically parallel to the ramp. The vertical distance between these two is the (amplitude) Error.

So now we know that this simple IIR filter has an exponential first order response, that has a zero error to a stable input value and a known constant error to a ramp input.

The formula for the Ramp error is Error = R/k – 1, where R is the rate of slope of the input. Substituting k = 0.15 into this equation gives an infinite error of 5.66666 and that is exactly what the graph shows.

The EMA in Practice

The above section has just described the inner workings of the most simple recursive filter, (IIR filter) which just happens to be the identical workings of an Exponential Moving Average (EMA) and virtually nothing is changed – apart from some naming!

For example a 20-day EMA is really an IIR filter with k = 0.095238 and that should be no surprise. We now also know that the Time Constant for a 20 day EMA filter is therefore 10 days and that the ramp error factor is 9.5 (assuming one cent per day ramp rate).

The above graph (taken from MarketTools Chart) shows the response difference between a SMA20 (Green) and an EMA20 (Blue). As the Close price begins to ramp the EMA initially tracks closer and wavers around while the SMA20 slides in slower (rounder) and forms a virtually straight line. This should be no surprise, as we know that the SMA is much less reactive to recent changes than an EMA.

You can clearly see the error that they have to a ramp in prices and this can be used to an advantage when doing technical analysis! This graph also shows the Moving Averages tracking the prices but with a very similar price offset (error) caused by the virtually constant rate of change in price over a limited time (in this case).

The problem with prices is that there is a feedback system that regulates the price variations and this feedback is human managed that works like this: For some reason, someone sees that they would like to purchase a particular stock, but the price is marginally higher than the earlier trading price.

When they purchase the stock the new price is now higher. Others see that price as either too high, correct or still cheap. With this thought in mind, other traders use the earlier prices as a reference and tend to correct that price back towards the reference price that each of them have. This causes the price to fluctuate in an oscillatory manner that tends to stabilise with time.

All is not lost as this is important to be understood that the Moving Average technology is a 1st order system, for now it can be used in the knowledge that if the prices are in general below the Moving Average, then the prices are actually falling with time, and if the prices are above the Moving Average, then the prices are in general rising with time.

It therefore makes lots of sense to know this very basic rule, as it means that the only shares to be involved in are those with the prices above the moving average line. But what time constant should be used for the moving average and why?

Virtually no technical analysis packages come anywhere near this depth, and they all treat SMA and EMA with a real lack of understanding. The problem is almost self explanatory in that virtually all data is EOD based and because of that, crossing over moving averages can resolve most ‘buy-sell’ signals!

In other words, the advancement of technical analysis stopped like a bus hitting a cliff when moving averages were resolved with EOD data. It works; profits from technical based sales can be realised; stop development!

One Moving Average

Having firmly established the fact that an SMA and an EMA are both 1st order systems, and that both of these effectively minimise the noise of trade variations, notably the close values based on EOD data; it comes as no surprise that these averages have a use as a buy or not buy indication for securities that have any form of trend. Their use is a simple application in that the error between the actual close price and the moving average when positive indicated that the security should be held and the converse.

This indicator is the most primitive of all technical indicators, and it is light years beyond using any form of financially generated indication to show if a security price is rising or falling in a trend.

The indicator really shines when the security is in a trend, but when the price hovers or flattens out; it has a problem of indecision. The graph below indicates this situation, and it is exemplified by including a switch function to show what can happen.

The switch function is shown the price moving average graphs. In the upper chart it is a EMA12, and as the close price fluctuates, the switch becomes very indecisive when the price trend level out or changes direction. One way around the problem is to use a slower moving average like the EMA21 as shown directly below.

The number of indecision points are reduced, meaning that the number of useless trades would be significantly reduced, but look closer and considerable profit runs are lost because the moving average is too late in switching over. In the background there is a positive in that the 12 and 21 EOD moving averages are smoother than the EOD close and that in itself can be used to advantage.

Two Moving Averages

By comparing two moving averages (which in themselves are already smoothed by their own attributes), a cleaner indication can be obtained and it can offer some advantages. The graphs below show some examples on the same security for direct comparison.

The above graph has the same switch function based on two moving averages EMA12 and EMA26 and see that the indecision is virtually nil.

This is a positive step, but a closer look at the actual switch over points shows that it is very conservative and in many cases considerable gains are lost before the decision is made to pull out. If it was not for this then this could be an ideal hold/sell indicator – purely based on close prices from EOD figures.

The next graph above (taken from OmniTrader) shows a six-month view of a stock and there are two exponential moving averages (EMAs) also on the graph. In this particular case the moving average that hugs the share prices is an EMA8 and the other one that slowly converges in the share price is an EMA35.

This chart is a good example as the faster EMA has the range of the EOD values of the share price intersecting it on several occasions. The slower EMA barely reaches the EOD price ranges.

OmniTrader has a very nice feature in that each testing indicator can be set to self-optimise itself for each security over a specified history (eg 250 trading days). This gives the indicators a good chance to provide a much better ‘hit-rate’ than you would normally get by simply setting the indicator parameters yourself. In this case they started at EMA12 and EMA40 and settled on EMA8 and EMA35 for an optimum result.

The problem is that of uncertainty as both moving averages converge on each other and do not have a clean crossover. This is not a major issue as we know that both SMA and EMA both are 1st order systems and because of that they asymptotically converge on a constant input, so if a price remains constant, then the two moving averages will both converge on that constant value, but at different rates.

One of the real problems is one of noise (actually price fluctuation about a constant value) and this can cause the faster moving average to whipsaw over the more stable slower (longer) moving average. There are several solutions to this problem, and each has their merits.

Comments and Corrections are welcome