Technical analysis as robust and cheap statistical filtering

05 May 2015

Technical analysts have been much derided over the years, particularly by the mainstream economics and finance professions. Perhaps the most fun account of chartists is still provided by Ed Malkiel's classic 'A Random Walk Down Wall Street'.

"On close examination, technicians are often seen with holes in their shoes and frayed shirt collars. I, personally, have never known a successful technician, but I have seen the wrecks of several unsuccessful ones. (This is, of course, in terms of following their own technical advice. Commissions from urging customers to act on their recommendations are very lucrative.) Curiously, however, the broke technician is never apologetic. If you commit the social error of asking him why he is broke, he will tell you quite ingenuously that he made the all-too-human error of not believing his own charts. To my great embarrassment, I once choked conspicuously at the dinner table of a chartist friend of mine when he made such a comment. I have since made it a rule never to eat with a chartist. It's bad for digestion." Malkiel (1973).

To the extent technical analysis has received attention it has mainly been in studies of whether it 'works', and here the jury is still out. To summarise somewhat glibly, some studies claim it does – while others claim it doesn't. In the interests of balance Andrew Lo is a rare example of an economist who has taken technical analysis very seriously, see Lo and Hasanhodzic (2010).

But regardless of whether academics think technical analysis works, the fact is that a huge investment management industry (widely referred to as 'managed futures' which includes commodity-trading-advisors or CTAs) has grown to manage hundreds of billions of dollars, including increasing amounts of pension and mutual money in the US. Many of these managers share a common DNA based on technical analysis, Burghardt and Walls (2011).

It is therefore worthwhile thinking about why technical analysis has been so enduring.

I believe one answer is that a lot of technical indicators can be interpreted as robust and cheap statistical filters, and as such are a very smart and quick way of learning about underlying market dynamics. This can be seen very clearly with an example: comparing a moving average crossover signal with a hidden Markov model switching model.

Moving average indicators

The moving average crossover rule is one of the oldest and most heavily used techniques by chartists, see Kaufman (2005) – probably today's bible of technical trading methods. In its simplest form this rule involves drawing one line on a price chart: an average of the price level over a rolling window of some fixed number of days, for example 14 days. For decades chartists have used this charting method as a signal of when to buy or to sell in a market. When the current price is above the average the chartist would be long the market, and when the current price is below the average the chartist would be short the market. Very simple, very naïve?

Hidden Markov Switching Models

Now let's think about how a sophisticated statistician, or, and I hesitate to use the ever popular term – a data scientist – might approach the problem of wanting to decide whether to be long or short. He may start by admitting no fundamental knowledge of the market in question, or even of markets in general. In fact to get a job in some CTAs today he would definitely *not* want to confess to any previous study of economics or finance! The scientist these days is probably of a Bayesian persuasion, i.e. they will approach the problem with fairly minimal prior assumptions about the underlying price process and attempt to learn from the data. They admit that their conclusions about the process may evolve over time as more data arrives.

One starting point would be to assume that there are two states of the market: when prices are going up, and when prices are going down. This sounds sensible enough, after all people frequently talk of bull and bear markets.

Assuming bull and bear markets exist, a next step could be to try and formalise this by assuming a probabilistic model for the underlying price process, and then to try and fit this model to the data. A plausible construction would be that of a hidden Markov model, a type of model that has been used extensively in other learning contexts such as speech recognition. The model could work as follows. In the "bull" regime, price returns (changes from one day to the next) are pulled from a random distribution with a positive mean. In the "bear" regime returns have a negative mean. What determines the regime at any point in time is a random process itself (the Markov process). Because we tend to think about markets being in bull or bear regimes for persistent periods of time (months or years) we want the regimes to be persistent, i.e. at any point in time we are more likely to stay in the regime we are in than to switch out of it. But from time to time we will switch, and that creates up and down trends in prices.

So far I have described the scientist's approach in words, but of course in reality he is going to have to implement the idea and that's where maths and computation come in. In the appendix below I have set out the scientist's assumptions, and also the equations that determine how he calculates the probabilities (in real-time) of being in one regime or the other. The overall exercise is one of statistical filtering, and many other machine and statistical learning type methods share similarities with this approach. The model certainly looks sophisticated and one would want a decent grounding in statistics and computation to be able to implement it confidently. Many papers have been published in highly rated journals applying this type of model to financial and economic series.

Applying chartism and learning to the S&P

Let's take a simple example of the S&P over the last few years.

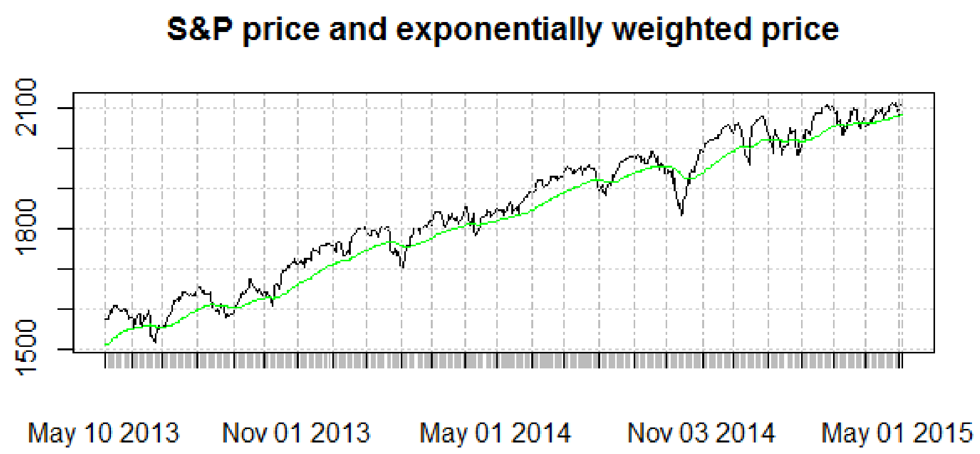

Over the period shown May 2013 to May 2015 the S&P trended up but there have been a number of short-lived periods (no more than a few weeks at a time) where the price dropped, most notably in October 2014.

What would our chartist and our scientist end up looking at? The chartist in our example is going to look at the difference between the price and his moving average of the price, often called an oscillator. In this case I assume he is using an exponentially-weighted moving average with a persistence of 0.96, a popular smoothing parameter in practice. The scientist, being much more sophisticated, thinks in terms of probabilities, and he would likely look at what his model tells him the probability of being in the bull regime is at each point in time.

The chartist's oscillator

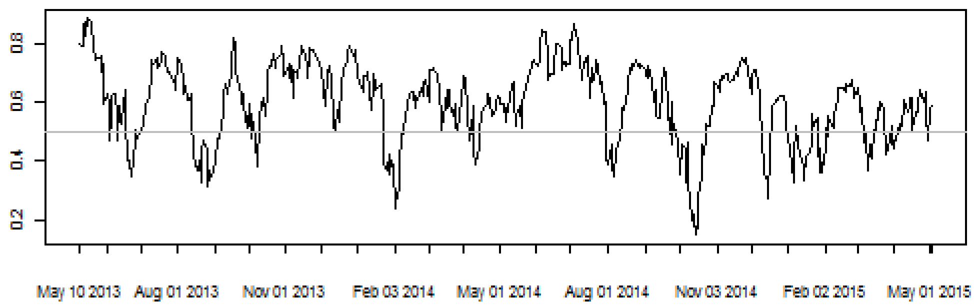

The scientist's probability of being in a bull regime

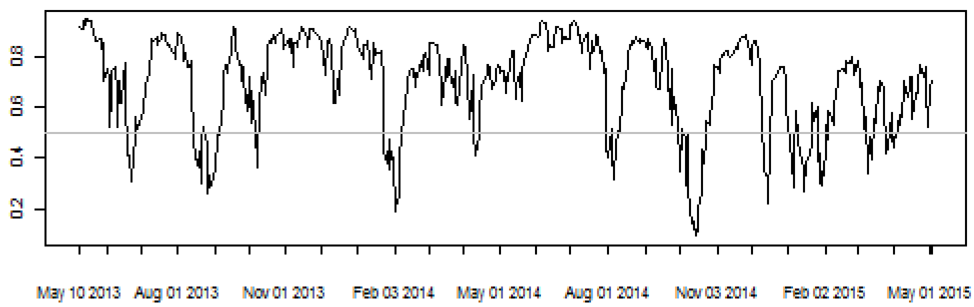

As will no doubt have been anticipated by the reader, the two look almost identical. The chartist who has simply drawn a line on a chart (all be it with some help from a computer), and the scientist who has employed a probabilistic model and implemented a real-time learning algorithm both end up looking at essentially the same thing. Both the chartist and the scientist will have different stories to explain how they got there and what it all means, but ultimately they have arrived at a similar point – and now need to make a decision about whether to trade.

Three observations for researchers and investors

The first point is that as the example above suggests, one can make the case that for decades CTAs have been using technical methods that can be interpreted as approximations to statistical learning algorithms. Researchers and traders hoping to apply machine learning should think carefully about what existing practitioner techniques are already providing. Likewise, technical analysts might benefit from thinking in more probabilistic terms (some work on this is forthcoming).

Secondly, heuristic and popular trading rules should not be dismissed as naïve, useless, or irrational. Why not start from the basis that as such methods seem to have survived multiple decades, it is worthwhile trying to reverse-engineer what type of assumptions and learning algorithms those agents might implicitly be using. Recent work by Barberis and co-workers (2015) on expectations formation is along these lines and reflects a growing interest among economists to understand how individuals process data.

The third and final observation is that rule-of-thumb techniques may often be as good as machine-precision computational techniques. Partly this is trivially because the theoretical maximum level of predictability is very low – markets are very close to efficient most of the time. Complexity in terms of modelling technique cannot be expected to produce earth-shattering improvements.

But also I suspect many apparently complex models, derived after pages of algebra, can often be approximated by much simpler "reduced form" equations, or convex combinations of simple equations. The endurance of exponential smoothing, despite decades of more sophisticated filtering (state-space, nonlinear etc) supports this observation.

Selected References

'Managed Futures for Institutional Investors', Galen Burghardt and Brian Walls, Bloomberg 2011.

"X-CAPM: An Extrapolative Capital Asset Pricing Model", Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer,**Journal of Financial Economics** 115, 1-24, January 2015
"Regime-Switching Models" , James Hamilton http://econweb.ucsd.edu/~jhamilto/palgrav1.pdf

"New Trading Systems and Methods", Perry J. Kaufman, Wiley 2005.

"The Evolution of Technical Analysis", Andrew W. Lo and Jasmina Hasanhodzic, Bloomberg, 2010.

"A Random Walk Down Wall Street", Ed Malkiel, Norton, 1973.

Appendix – algorithms to produce trading indicators

This appendix is provided to demonstrate as clearly as possible how an old popular technical indicator can produce something observationally equivalent to the output of a statistical learning tool.

The hidden Markov model is very widely studied but a good reference is by James Hamilton who popularised these models amongst econometricians: "Regime-Switching Models" http://econweb.ucsd.edu/~jhamilto/palgrav1.pdf

The chartist's computations

$$\text{Oscillator}(t) = P(t) - \text{EWMA}(P(t))$$

$$\text{EWMA}(P(t)) = 0.96*\text{EWMA}(P(t-1)) + 0.04*P(t)$$

The scientist's computations

Volatility adjusted returns, $r(t)$ are defined as $(P(t) - P(t-1))/\text{Volatility}(t)$

I have used a simple Riskmetrics style EWMA volatility model. The standard-deviation of these vol-adjusted returns is approximately 1.

There are two possible states $s(t) = 1$ or 2

In state i $r(t) \sim N(\mu_i , \sigma)$

The probabilities of switching between states are governed by a 2×2 transition matrix

$Pr(s(t)=i \mid s(t-1)=j) = p_{ij}$

It is possible to estimate the parameters from the data by numerically solving for those that fit the data best, by for example maximising the log-likelihood. In practice this can be quite hard to do because often many different sets of parameters appear to fit the data equally well. Partly this is because the data sample isn't long enough, but mainly it is because any proposed model specification is only an approximation.

In our S&P example we have set $p_{11} = p_{22} = 0.99$. Therefore $p_{12} = p_{21} = 0.01$, or in other words there is a 1% chance of switching between regimes each day. We have chosen a mean parameter for each regime that produces an intuitive segmentation of the data and then applied the updating algorithm below to generate the probabilities.

Given the closeness of the two indicators observed above, another interesting exercise is to start from an oscillator of one's choice (e.g. a 2 day and 10 day crossover) and then use a Markov switching framework to ask the question, 'If chartists are implicitly trying to back out probabilities of being in bull or bear markets, then what underlying assumptions about the strength and durations of these regimes might they implicitly be making?'

Defining

$\Omega(t)$ is a vector containing all the observations of returns up to and including time t

θ is a vector of the parameters $\{ \mu_1 , \mu_2 , \sigma , p_{11} , p_{22} \}$

What we want to get to is the probability of being in state j at time t, based on data we can observe up to and including time t (but not beyond), that is:

$Pr(s(t)=j \mid , \ \Omega(t-1) ; \theta) =$

$$\frac{\sum_{i=1}^{2} p_i p_j Pr\,(s(t-1) = i \mid \Omega(t-1)\,;\,\theta) \quad f\,(r(t) \mid s(t) \,=\, j,\ \Omega(t-1);\ \theta)}{f\,(r(t)\,|\Omega(t-1);\ \theta)}$$

We need to calculate this iteratively starting at the beginning of the data and updating our estimates as we go along. This is why the algorithm is a form of learning.

First calculate the density of each data point conditional on belonging to each regime

For each i =1,2

$$f(r(t) \mid s(t) = i,\ \Omega(t-1)\ ;\ \theta) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\ \frac{\left(r(t)-\mu i\right)^2}{2\sigma^2}\right)$$

Next calculate the density for each data point (the denominator above) as

$$\sum_{i=1}^{2} \sum_{j=1}^{2} p_i p_j \Pr\left(s(t-1) = i \mid \Omega(t-1)\ ;\ \theta\right) \quad f\left(r(t) \mid s(t) = j, \Omega(t-1);\ \theta\right)$$

The whole process is iterative and in practise requires some starting values for the conditional probabilities of being in either regime at time 0, often one can begin with 50% chance of being in each and then as the computation continues the data will quickly push the probabilities to ones that fit the data.

Disclaimer

This research note is published by Neuron Advisers LLP ('Neuron'). You have been provided with this material only upon your acceptance of these Terms and Conditions herein.

Neuron Advisers LLP is authorised and regulated by the Financial Conduct Authority, with firm reference number 563919. Neuron Advisers LLP is registered with the Commodities Futures Trading Commission under NFA ID 0439462.

Investments or investment services mentioned in this research note are not available for investment and are not being marketed in any jurisdiction. The contents of this research note are not intended to be read by any persons in any jurisdiction other than the United Kingdom.

This material is presented for information purposes only. It is intended for your personal, non-commercial use. No information or opinions contained in this material constitute a solicitation or offer by Neuron to buy or sell securities or to furnish any investment advice or service. Neuron does not provide investment advice, tax advice or legal advice through this material and you agree that this material will not be used by you for such purposes.

This material is intended as a general introduction to Neuron and the Neuron research blog by which means Neuron can express ideas and opinions. The material contained herein is the sole opinion of Neuron.

The information provided in this research note is intended for institutional investors and Professional Clients and Eligible Counterparties as defined by The Financial Conduct Authority and for those who are considered as qualified eligible persons as defined by Commodities Futures Trading Commission Regulation 4.7. It is not intended for retail investors.

The contents of this research are not intended for distribution to, or use by, any individual or entity in any jurisdiction where their distribution or use would be contrary to local law or regulation or which would subject Neuron Advisers LLP to registration with the jurisdiction. You should be aware that any rules and/or regulations applicable to providing financial services (and the resultant investor protections that may be available), may not apply to persons who obtain information from the internet and its various applications, of which this material forms part.

Neuron Advisers LLP assumes no responsibility for access to this material by any person located within a country or jurisdiction where such access would be contrary to any law or regulation in that country.

We try to ensure that the information in this research note is correct, but we do not give any express or implied warranty as to its accuracy, timeliness or completeness, nor is Neuron under any obligation to update such information. Any data supplied has not been audited and is provided for information purposes only.

We are not liable for any damages (including, without limitation, damages for loss of business or loss of profits) arising in contract, tort or otherwise from the use of or inability to use this research note, or any material contained in it, or from any action or decision taken as a result of using this research or any such material.

This research note may provide links to other research and websites. We do not control the linked sites or research and we are not responsible for the contents of any linked site or research, any link in a linked site or research, or any changes or updates to such sites or research. We provide these links to you only as a convenience, and the inclusion of any link does not imply our endorsement of the site or research.

Neuron Advisers LLP is a limited liability partnership registered in England & Wales with number OC367248 and registered office address at Challoner House, 2nd Floor, 19 Clerkenwell Close, London, EC1R 0RR, United Kingdom.

Any questions about the contents of this material should be directed to: enquiries@neuronadvisers.com.