Theory will only take you so far...

Let us imagine you are a trend follower in a CTA circa 2010. Your trading system is really a fancy name for the moving average crossover (MACD) that despite being ridiculously simple, has been making tons of money for the last 20 years, without a single negative year.

You have a big research department to impress your clients, but do you allocate to any other trading signal? Like hell you do. Trend following is probably the nicest form of alpha out there: It is simple and robust, does not require lots of data or computation power, completely explainable, risk-manages itself, no correlation to your clients' long-only portfolio, positive skew, consistent historic behaviour over hundreds of years and all asset classes. You would be an idiot to trade anything else.

The real issue is how to keep your research department occupied, knowing fully well you will not deallocate away from momentum.

As it turns out, there is one job to give them: in 2010 the cleverest of PhDs from the very best universities were all doing (essentially) one thing in most CTAs: fitting a non-linear response function (a "spline") to the moving average crossover.

Fitting the spline: some maths

The MACD is an average of historic returns. FAST MACD will average last week's returns, the MEDIUM speed one will look at two weeks, or a month. The SLOW MACD perhaps averages the recent quarter's returns. We tend to normalize MACD by price volatility so that its underlying distribution is roughly Normal(0,1).

The question is: if the MACD we observe today is X, what is the most likely Y, the price move we see tomorrow? Y is important as the best position to take in the market depends on Y: If you think the market is most likely to go up (i.e., Y>0), we buy, and if Y<0, we sell.

We can look at historic returns and decide. It is, after all, a simple X-Y plot. We average over all historic X and Y, and we can also average it over all assets we trade (and even assets we don't).

You can try and plot it yourself. What you will find is that a linear fit Y = const X is a pretty good fit. Unfortunately, the constant is very small: around 1% to 2% (at best) of the Y variance is explained by X.

But that is all you need. Imagine you are running a casino. Your roulette edge over even-odds is not much higher. But with the many spins of the wheel comes utmost certainty: Central Limit Theorem dictates that over time, the casino will make money. If you have not gone bankrupt by allowing a $100m bet on a single spin.

That is precisely the job of the MACD spline: to ensure you have not allocated too much risk to a single market at any one time. So your position isn't just linear in X but a response function that is capped.

There are other reasons to cap your risk exposure. The first is real life expectations: Suppose the markets just jumped because the Fed cut rates unexpectedly. Even though that caused a big move,

do we really expect the unexpected rate cut to happen again tomorrow? The second is uncertainty: The MACD X is N(0,1). This means that we don't see X=3 often. And we rarely see X=5. So when we do, we are in uncharted waters. Do we want to take 5 times the risk on an event we only observed a handful of times historically?

So now we know we want to cap the spline, the response function, but how?

The council of Elrond

Determining the exact spline turns out to be a great occupation for scores of researchers over years and years of pseudo research. It is a problem that keeps on giving... every six month you can refit the spline with just a little bit more data. Each sector (Equities, Fixed Income, Commodities, FX) could fit a spline particular to the sector's particular set of markets. Unfortunately, by and large, it is a futile exercise in rear-view mirror driving (remember how great the uncertainty is in the first place?). Many years ago, I hired Energetic Mike to my Fixed Income team, and we decided to end the discussion by not looking at historic data at all. We came up with xN4, as the one spline to settle the dispute. It had three amazing properties:

- xN4 was theoretically "correct" in a Mathematical sense (expectation over sqrt(uncertainty)) so intellectually satisfying.
- xN4 is a catchy name (for geeks).
- Most importantly, xN4 was similar enough to each of the splines proposed by each of the other sectors so that none of the other sectors would object too vociferously if forced to use it.

It was a mathematical (and political) success: xN4 became the standard internally and xN4 and travelled to other CTAs: everywhere I went, xN4 seemed to be ubiquitous.

When I moved to another CTA, another good colleague, Gentle Mike, was just putting the finishing touches on a big spline fitting exercise, optimizing the "firm spline". Gentle Mike's Spline was a piece-wise linear function purely data driven but if you plotted it next to xN4 you could hardly see the difference. It was very satisfying to see real life following theory.

So xN4 is the one spline to rule them all.

But they were all of them deceived...

You cannot unscramble an omelette

Suppose my fast predictor FAST has taken three values: 0, 5 or -5. Feeding FAST through the spline, xN4(0), xN4(-5) and xN4(5) are all very similar and are essentially zero. Yet they represent very different realities:

- We have high certainty of no fast trend.
- We have expectations of a strong negative trend but with high uncertainty.
- We have expectations of a strong positive trend but with high uncertainty.

If you are only trading FAST, the risk you take *is* equal: in all these cases we should be out of the market!

However, we have our good friends MEDIUM and SLOW and they make different predictions. And the way we should combine FAST with them is now inherently different. It makes a great deal of difference to know that FAST is very certain there is no trend at all vs. FAST indicating a very strong, though uncertain, trend.

This is true not just of xN4 but of any spline: The CTA industry has fallen into the trap of early over-simplification: Any spline will combine two quantities (expectation and uncertainty) early in your process, in a sensible way, but we are then unable to unscramble these later.

This has real implications, e.g., when it comes to machine learning. When we want to build a non-linear ML driven aggregator of FAST, MEDIUM and SLOW, we should always have in mind the underlying distributions of each. Specifically: regions where we have little data to train the ML, should be monitored and we should be out of the market (and ignore ML) in these regions. Interestingly, most ML researchers are completely oblivious to this consideration: ML is a powerful hammer you can swing at any problem, and we all tend to forget that we must understand the underlying statistical properties of our variables and thinking long and hard about what to do in atypical situations. I have seen high profile ML trend following attempts failing precisely for this reason: ML can capture non-linear relationships beautifully and then fail to cap its risk, not realising its own data limitations.

It is all about the journey.

Even if we apply FAST on its own, it may not make sense to use xN4, especially when we consider costs. xN4 may represent the optimal uncertainty-adjusted *position* but we must think about how FAST journeys through time before making each *trade*. Suppose FAST increases from 1 to 1.1. How much incremental risk should we add?

FAST is a mean-reverting stochastic process (which is why it has a stationary distribution). That means that when FAST starts an excursion from 1 to 1.1, it will spend a reasonable amount of time over 1.1. FAST will then revert back and hit 1.0 and we will undo the incremental trade we are about to execute. If FAST is at 3 and moves to 3.1 it will have only a very short excursion above 3.1 before reverting back to 3. What it means is that alpha-decay, *for the same predictor*, depends on where we are. We know that when we trade with costs, we should adjust our trading based on alpha decay so our incremental risk adjustment as FAST moves from 1 to 1.1 is very different to the incremental risk we should apply as FAST moves from 3 to 3.1. This is especially important if you happen to trade less liquid, more expensive, markets like I do these days.

Theory will only take you so far.

One of the benefits of working in a small firm is how closely we work with the traders (as I am writing this article, I am sitting across Ashley, our interest-rates swaps specialist trader) and talking to them about how they see the book's behaviour. It is only when you hear comments like "it is a shame, we opened a position only to close it again" that you realise you must start incorporating cost implications much earlier in your trade construction. xN4 has been a good friend for over ten years and is theoretically exact but we must part ways: (cost-free) theory will only take you so far.

Yoav Git
Managing Director, Quant Research
SAFI Model Owner

Opinion piece. See disclosures