CTAs and Alpha Generation: Is Efficient Implementation the Answer?

Andrew Beer, Managing Partner – DBi

By Andrew Beer – DBi

After a decade of studying CTAs, we have drawn three conclusions about the nature of their alpha generation. At the strategy level we believe:

1. CTA alpha is generated by early, contrarian bets on big price moves across commodities, rates, currencies and equities.

2. CTA alpha is lumpy, because these types of "regime shifts" are rare.

3. CTA alpha is concentrated in major liquid markets simply because regime shifts play out in, and more money is made in, those markets.

The obvious question for allocators is how best to access this alpha generation. The standard institutional playbook has been to diversify risk by investing in multiple managers (or QIS models), which is sensible given wide dispersion and no evidence of persistence of alpha. Further, many allocators have successfully driven down fees, which shifts alpha from managers (banks) to investors.

The framework above, though, raises a third issue: whether CTAs should focus more on implementation efficiency. Why should allocators care? Because despite constant model evolution and expansion into more esoteric instruments, strategy alpha has been in decline. The SG CTA and BarclayHedge BTOP50 indices have delivered Sharpe ratios of 0.2-0.3 over the past decade, roughly half that generated during the 2000s. One theory is that implementation costs have risen faster than incremental alpha generation, and hence model and portfolio "evolution" is partly responsible for the decline in strategy-wide alpha.

To take a step back, implementation efficiency can be thought of as the difference between the theoretical "signal" and the actual realized performance (after transaction costs but before fees). Some managers refer to this as "slippage." Generally, implementation costs are broken down between explicit and implicit trading costs. The former are straightforward: effectively, the commissions and fees incurred in trading. More complicated are implicit costs: the market impact of buying and selling a given instrument. Simplistically, if buying drives up the

> "Regardless of the precise calculation methodology, it is axiomatic that, relative to many types of investing, the trading costs of a typical CTA are very, very high."

www.hedgenordic.com – June 2026
HEDGENORDIC

price of a future by 10 bps over an hour or two, this is a real economic cost of implementation. Industry experts estimate these latter costs at 3-5x the former.

A fair calculation of implicit implementation costs is admittedly tricky. The analysis requires numerous assumptions and a tick-by-tick analysis of price moves for each contract over an extended period. But several parameters are beyond dispute. Relative to most strategies, CTA turnover is remarkably high: due to leverage, false positives and constant readjustment, a dollar invested today might turn over twenty times during the coming year. Trading costs rise nonlinearly from the most liquid to least liquid futures markets; plus, gap risk in the latter is nearly unquantifiable when multiple CTAs simultaneously sell (or cover). Lastly, short term models and risk management tools tend to increase turnover, often when liquidity is scarce.

The implementation cost question has never been a secret. A decade ago, several brand name managers claimed that heavy investments in trading technology had reduced "slippage" from 400 bps to 200 bps per annum. In our deep dive into the space back then, numerous sophisticated allocators shared their own calculations of 200-300 bps per annum. And yet in recent years, the industry-wide push to roll out more complex models and push into "alternative" markets arguably has reversed those efficiency gains. Regardless of the precise calculation methodology, it is axiomatic that, relative to many types of investing, the trading costs of a typical CTA are very, very high.

The next and obvious question is whether these costs can (or should) be reduced. Many strategies separate signal generation from implementation. A macro trader, for instance, might develop a view from dozens of esoteric markets and data points but then implement in the most liquid and efficient way possible. By contrast, a CTA manager might evaluate two hundred markets, then reflexively buy or sell instruments in two hundred markets. Of course, positions generally will be sized to reflect market liquidity or accessibility (e.g. swaps not futures). Still, complex implementation can sometimes seem like a badge of honor, like when a mutual fund holdings report stretches to two dozen pages.

"The evidence today is strong that the efficiency gains in a well designed replication model can translate into a meaningfully higher Sharpe ratio for investors."

To detect the CTA "signal", as described above, some complexity is necessary. We know this because the Sharpe ratios of trend following models rise with more instruments (at least up to a point). And, of course, the model will identify some truly idiosyncratic opportunities. Most of the time, though, the models appear to identify "clusters" of positions that are driven by deeper, seismic factors. This raises the question as to whether CTAs should "compress" the majority of the portfolio into a smaller number of more liquid, more efficient exposures.

A decade ago, this discussion would have been entirely theoretical. Since then, factor-based replication of CTA returns has been a real-world test of the efficiency question. For the uninitiated, factor replication models seek to infer the broad positioning of the CTA space using only reported daily returns of a diversified portfolio of funds – whether hedge fund, mutual fund, UCITS or, now, ETFs. Replication in effect seeks to compress a highly complex portfolio into one of a dozen or fewer major markets -- detecting not just directional positions but, importantly, more subtle cross asset relationships. Controlled rebalancing in the most liquid markets has the potential to drive down implementation costs by 90% or more. The evidence today is strong that the efficiency gains in a well designed replication model can translate into a meaningfully higher Sharpe ratio for investors. This suggests that efficiency gains may be more important than any theoretical "lost alpha" from esoteric instruments, risk controls and other features.

Managers understandably express frustration about the lack of AUM growth across the space. Perhaps the unspoken reason is the recent low Sharpe ratio of the strategy. The statistically inclined tend to overlook this issue: after all, there are few strategies with low correlation to both stocks and bonds, a tendency to perform best during market crises, modest drawdowns relative to other diversifiers, and accessibility with relatively low fees and liquidity. Those statistical features mean that, even with a lower Sharpe ratio, alpha generation is more than respectable. That said, for most allocators, "at the end of the day, you can't eat alpha." Perhaps an industry wide focus on efficiency will bring Sharpe ratios back in line with other asset classes and, finally, open the floodgates.