March 2024

Energy Quantamentals: Myths and Realities about Algorithmic Oil Traders

OIES ENERGY COMMENT

Ilia Bouchouev, Senior Research Fellow, OIES and Adjunct Professor, New York University

Introduction

The first article in the new series on Energy Quantamentals introduced the main participants in the market for oil derivatives[1]. The article explained how the description of these participants is full of misnomers and that the behavior of many traders is often misinterpreted by the general public, and, unfortunately, is mislabeled by regulators. For example, according to regulatory definitions, the label of producers counterintuitively applies to large physical speculators, while the trading activity of genuine oil producers is conducted over-the-counter (OTC) via swap dealers. Furthermore, many analysts tend to mistakenly refer to all quantitative funds, financial speculators, and algorithmic traders as CTAs, which stand for commodity trading advisors. Such a generalization reflects a misunderstanding of who CTAs are and their role in the oil market. The objective of this article is to shed some light on their actual trading strategies even though the term CTA itself happens to be just another misnomer.

A CTA is a regulatory designation by the U.S. Commodities Futures Trading Commission (CFTC) which requires registration with the U.S. National Futures Association (NFA).[2] It applies to anyone providing an advise, directly or indirectly, on commodity futures trading, where the term commodity somewhat bizarrely includes so-called financial or “excluded” commodities, such as interest rates and currencies. This designation is rather broad and many registered CTAs are not even involved in managing any financial assets.

In other words, not every CTA is a quantitative trader, and vice versa, not only CTAs rely on quantitative algorithms for trading in the oil market. In fact, the latter group of non-registered algorithmic traders is growing rapidly, as new trading outfits are being created in jurisdictions outside of U.S and Europe, which bypass the burdensome designation of a CTA. In addition, quantitative trading groups are also being housed under umbrellas of physical trading organizations, large family offices, and even some sovereign-controlled financial entities. In short, the amount of quantitative trading in the oil market today is significantly larger than trading which can be tracked by the subset of registered CTAs. While in this article we focus on the latter mostly because of their origins and larger market presence, we should be cognizant that similar techniques are now being adopted on a much broader scale.

The History of CTAs

By and large, traditional quantitative CTAs are descendants of technical traders which have been active in commodity markets, arguably, for as long as these markets existed. While many of their quantitative trading techniques remain largely undocumented, some, such as, for example, candlestick charts used by Japanese rice traders in the eighteenth century, are widely known. It is impossible to give credit for pioneering quantitative trading to any single person or a trading group. The sophistication of such trading grew with the development of the computing power, and by the late 1950s, several trading outfits already resembled modern-day quantitative hedge funds.

Perhaps, an honorary mention of being among the first and the most successful quantitative CTAs should go to Commodities Corp., which was founded by Helmut Weymar in 1969 in an old farmhouse near Princeton University. Seeking for practical applications of his Ph.D. thesis on fundamental cocoa trading[3], Weymar attracted interest from seed investors, including some luminaries, such as Paul Samuelson, who

---

[1] See I. Bouchouev (2024), “Energy Quantamentals: Who is Who in Financial Barrels?”, OIES Energy Comment, January 2024, https://www.oxfordenergy.org/wpcms/wp-content/uploads/2024/01/Energy-Quantamentals-Who-is-who-in-Fincancial-Barrels-Final.pdf.

[2] See https://www.nfa.futures.org/registration-membership/who-has-to-register/cta.html.

[3] See F. H. Weymar (1965), “The dynamics of the world cocoa market”, Ph.D. Thesis, MIT. In addition to the empirical study of cocoa prices using fundamental data, Weymar also developed one of the first theories of storage by linking commodity prices and inventories via coupled differential equations.

was his MIT advisor and a Nobel laureate, and Paul Cootner, another MIT professor and the author of a popular book on the random character of stock prices. The company was known for hiring analysts with the background in mathematics and computer science to study complex behavior of commodity prices. In addition, the trading roster of Commodities Corp. over the years included the storied names of Paul Tudor Jones, Bruce Kovner, and Louis Bacon, who subsequently built their own multi-billion dollar hedge funds which are still among the leaders in the commodity trading space.[4]

Historically, quantitative trading style has been viewed as somewhat opposite to fundamental trading. The goal of fundamental trading is to estimate the price of the commodity required for balancing supply and demand. It is often associated with price mean-reversion as one expects forces of supply and demand to react to any dislocation in prices and restore some sort of a price-quantity equilibrium. In practice, however, this is easier said than done as the price elasticity of supply and demand for many commodities is rather low. This is what the star-filled roster of Commodities Corp. painfully learned in their early days of trading. While traders' original vision was to capitalize on unique fundamental insights and betting on price convergence towards some fundamentally derived fair values, the markets often defied rationality by trending further away from such values.[5]

To mitigate mounting losses, the risk management system was prudently put in place to counter mean-reverting bets of discretionary traders with a rule-based strategy which simply followed the prevailing trends and penalized traders for trading against the trend. Surprisingly, the trend-following risk-management system itself turned out to be a much more successful strategy than trading based on perceived imbalances between supply and demand. Moreover, this systematic trading system removed the need to gather expensive fundamental data and it also eliminated all human emotions and biases. It subsequently evolved into the fund's new core vision and was expanded across many other commodities. Commodities Corp. achieved a legendary success, and in 1997 it was acquired by Goldman Sachs Asset Management (GSAM).[6]

The idea of rule-based trading without any human intervention has been adopted by many other hedge funds and expanded across a much wider universe of commodities. Trading commodity futures was very appealing as similar techniques could be applied to many commodities which have low correlations among each other, generating significant diversification benefits. The business of a quantitative fund is like a casino where the edge on individual trades is rather small, and one needs to conduct them repeatedly across many commodities to ensure sustainable profits. Unlike stocks and bonds, commodities were also attractive because selling futures was as easy as buying them, allowing traders to capitalize on trends in either direction. Finally, since futures trading does not require committing much capital upfront, besides a relatively small exchange-required initial margins, commodities bets could be highly leveraged.

We will return to the composition of broader systematic portfolio in the context of spillover effects on oil in one of subsequent papers, but in this article our primary focus is on how these systematic rules are used specifically in the oil market.

From Storage to Momentum and Carry

In contrast to fundamental trading, systematic or rule-based trading does not attempt to quantify any fundamental fair value of a commodity. Instead, it merely identifies historical price patterns and speculates that such patterns will repeat themselves in the future. These price patterns are usually deemed to be caused by the behavior of certain market participants. The tendency of many markets to trend is one of

[4] Paul Tudor Jones is the founder of Tudor Investment Corporation, Bruce Kovner is the founder of Caxton Associates, and Louis Bacon is the founder of Moore Capital Management.

[5] In the early days of trading, Commodities Corp. even recruited local agents to monitor the state of cocoa trees in Africa. For more details about the history of Commodities Corp., see S. Tully (1981), "Princeton's rich commodity scholars", Fortune, February 9.

[6] See H. Lux (2003), "What becomes a legend?", Institutional Investor, February 1.

such patterns. In financial markets, trends or momentum are typically explained by various behavioral theories, such as human traders' slow reaction to news, or their propensity to take profits too quickly while holding losing positions for too long.[7]

In the oil market, however, the explanation of momentum is simpler. By and large, the roots of oil momentum lie in the theory of storage which argues for a negative and nonlinear relationship between oil prices and inventories. Since oil supply and demand are highly inelastic, if inventories are building today, then more likely they will also be building tomorrow. Thus, trends in inventories tend to be persistent and, as the consequence of the theory of storage, such persistency translates into trends in prices and time spreads.[8]

The most common way to specify a momentum strategy is to compare the current price of oil to its moving average. In such a strategy, one buys oil futures when today's price exceeds its moving average and sells futures when the current price falls below its moving average. Consider the basic strategy M(1,20), where current one-day price is compared to its 20-day moving average, which roughly corresponds to the one-month trailing price average. Somewhat surprisingly, such a basic strategy generated 9.7% annual return over the past 30 years (see Figure 1).[9] However, the strategy also experienced large drawdowns, and in four years it lost over 25%, the threshold often considered by investors as the maximum tolerance for losses. The Sharpe Ratio of the oil momentum strategy over this period is only 0.25[10], which makes it difficult to trade on a stand-alone basis.

Figure 1: WTI Momentum Strategy (1994-2023)

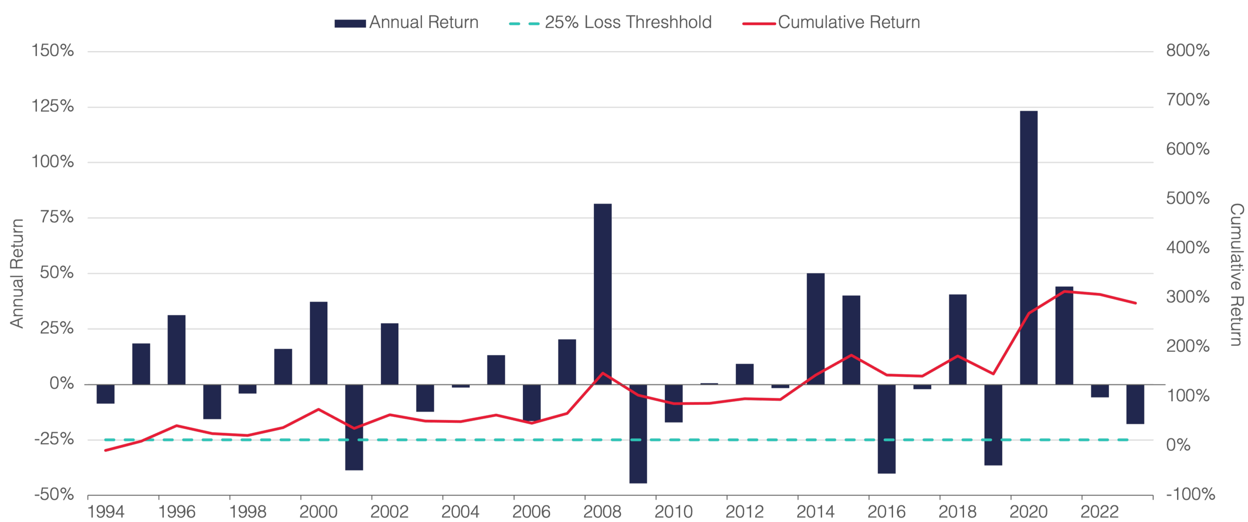

Source: Author's own calculations.

[7] Here, we use the terms momentum and trend-following synonymously. In general, trend following refers to a time series momentum where the trading signal is generated by the asset's own price history. The term momentum is more often used in cross-section settings for multi-commodity portfolios which rank assets based on the strength of their momentum signals, and then buy higher ranked commodities and sell lower-ranked ones.

[8] For more on the theory of storage, and how the boundary of zero inventories and the boundary on the maximum storage capacity impact futures prices and time spreads, see Bouchouev (2023), "Virtual Barrels", https://link.springer.com/book/10.1007/978-3-031-36151-7.

[9] In all strategy examples in this article, the investment is held in the nearby futures contract and rolled to the second-month futures contract on the last business day of the calendar month.

[10] Sharpe Ratio is traditionally defined as the ratio of annualized excess return over the risk-free rate to the annualized volatility of returns. Since in the oil market, it is more common to measure profit-and-loss in dollars per barrel and the interest rate plays only a minor role in the performance of the strategy, here we define Sharpe ratio simply as the ratio of annual profit-and-loss to the annualized standard deviation of daily profit-and-loss.

There are numerous variations of momentum. For example, instead of reacting to a price move on a single day which could produce a false signal, traders may want to see some confirmation of the trend before initiating the trade. In this case, they often use the crossover of moving averages, where for example, buy (sell) signals are defined when the 5-day weekly moving average exceeds (falls below) the 20-day monthly moving average. We denote such a signal as M(5,20). Multiple momentum signals on different frequencies are often combined to form an aggregate momentum signal. For example, one can equally weight short-term momentum of daily versus weekly moving averages M(1,5), medium-term momentum of weekly versus monthly M(5,20), and long-term momentum of monthly versus annual M(20,250). One can also require the signal to remain intact for several days before taking on a position. Another popular variation of momentum is a class of breakout strategies where the position is taken only when the current price exceeds the previous maximum price or falls below the previous minimum price, calculated over some historical lookback period.

Obviously, the more parameters the strategy uses the more susceptible it becomes to overfitting, and momentum strategies are particularly notorious in this regard. This is an important and often overlooked drawbacks of momentum strategies described by the author in a much greater detail in his recently published book on quantitative oil trading.[11] In particular, it explains the root cause of extremely high sensitivity of the momentum strategy to the choice of model parameters, such as the selection of lookback periods for the moving averages. Relatively small changes in optimized parameters of the momentum strategy could lead to drastically different performance. As the result, many systematic traders often learn the harsh reality that oil momentum works until it does not.

A much stronger systematic signal in the oil market, which also has its roots in the theory of storage is known as carry. Systematic traders simply define carry as the slope of the futures curve. In a backwardated market, the carry is positive, and in a contango market it is negative. A basic carry strategy then buys futures when the market is backward and sells futures when the market is in contango. The strategy is not particularly sensitive to the choice of specific futures as most of the futures time spreads have the same sign. If one wants to incorporate seasonality, then it is common to measure carry as the annual spread between 1-month and 13-month futures contracts. The performance of such a simple strategy generated 18.5% annual return over 30 years with a much better Sharpe ratio of 0.48.

Figure 2: WTI Carry Strategy (1994-2023)

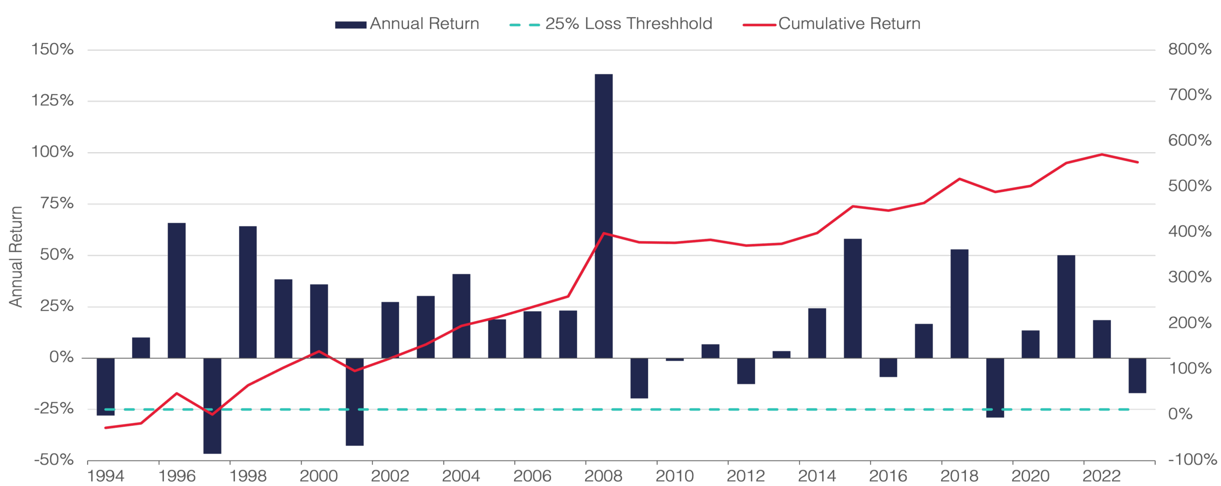

Source: Author's own calculations.

[11] See I. Bouchouev (2023), "Virtual Barrels", https://link.springer.com/book/10.1007/978-3-031-36151-7.

How can one explain such a strong performance of what, at the first glance, may appear to be a rather naïve trading strategy? The answer lies in the behavior of market participants, specifically the role of the storage hedger. In the previous article, we showed that the category of inventory hedgers indeed represents the largest set of participants in the oil futures market.[12]

The business model of a storage operator hinges on the arbitrage relationship between spot and futures price.[13] If the market is in contango and the spot price falls below futures by more than the cost of storage, including the cost of borrowed money to buy oil, then the storage trader buys a physical barrel of oil, puts it in storage, and sells futures to compete the arbitrage. The futures market does not directly see the physical leg of this transaction, but the financial selling of futures generates a powerful and persistent force on futures which often overpowers other factors and pushes the futures price down. The opposite occurs when the market is in backwardation and the spot price exceeds futures. The storage hedger then pulls oil out of storage, sells it in the physical market, and buys back short futures hedges. This creates an upward pressure on price. The dynamics is illustrated in Figure 3.

Figure 3: Carry and Inventory Hedging

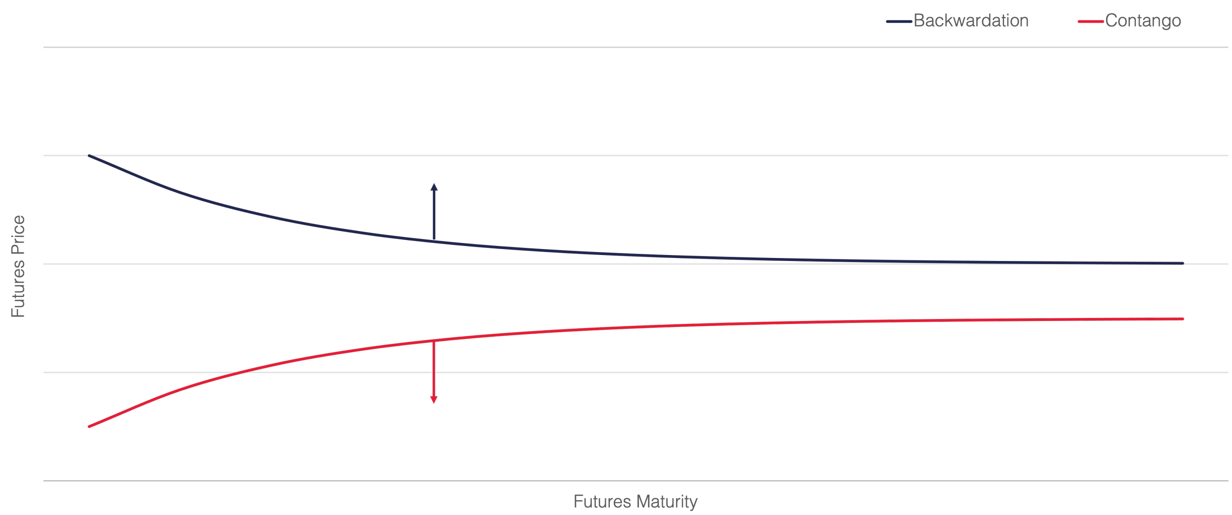

Source: Author's illustration.

In other words, when inventories are building, storage hedgers sell futures, and when the market is fundamentally strong, they are buying short futures back, and this is precisely what the systematic carry strategy does. Effectively, the role of the storage trader is to convey vital fundamental information about supply and demand to the futures market. The financial systematic trader simply trades ahead of the anticipated behavior of the largest market participant. The advantage of the carry strategy is that unlike momentum, carry is essentially model free. Carry is a forward-looking signal, and there are no lookbacks to choose and no parameters to estimate. The signal is fully determined by the current shape of the futures curve, and the strategy is much more robust.

One challenge with the carry strategy is its relative slowness. The market simply does not flip between contango and backwardation too frequently, which results in the same futures position being held for too long, even when market fundamentals start changing. To speed up this strategy, quantitative traders often

[12] See I. Bouchouev (2024), "Energy Quantamentals: Who is Who in Financial Barrels?", Oxford Institute for Energy Studies, January.
[13] Technically, in the oil market the spot price with an instantaneous delivery does not exist as oil always trades for forward delivery. Here, we use the term spot only for a pedagogical convenience, which in practice, refers to the futures contract with the shortest maturity.

apply the above-mentioned momentum signal directly to carry, or to the slope of the futures curve, instead of applying it to the price of oil. In such a strategy, which we call carry-momentum, one buys and sells futures when short-term moving average of the futures time spread exceeds its long-term moving average. In other words, one buys futures when backwardation accelerates or contango decelerates, and vice versa, one sells futures when backwardation decreases or contango increases.

The carry-momentum strategy forms the core of many CTAs' trading systems. Even basic carry-momentum strategy based on the crossover of the one-day and the one-month moving averages generated truly spectacular 24.3% annual return over past 30 years with the Sharpe Ratio of 0.63 (see Figure 4). One can think about carry reflecting the current state of inventories and carry-momentum reflecting market expectations of future changes in inventories. Since the market trades based on forward expectations, applying some momentum property to the term-structure of futures serves as a good proxy for such expectations.

Figure 4: WTI Carry-Momentum Strategy (1994-2023)

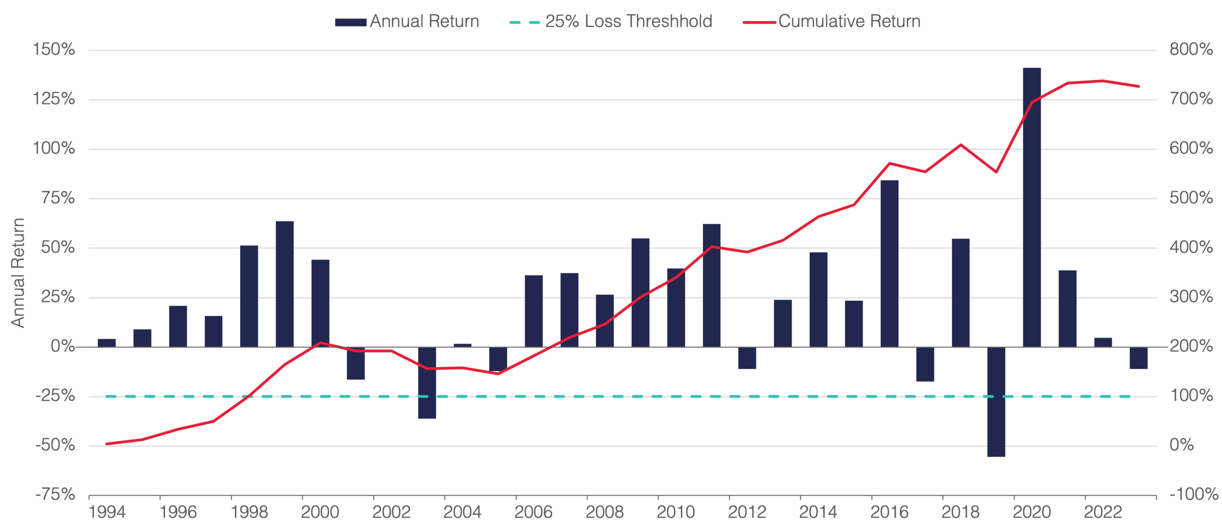

Source: Author's own calculations.

The predictive power of systematic signals is enhanced by the skill of combining multiple signals in a smart way. The application of momentum signal directly to carry is one example of how different signals can be blended. Many other alternative approaches, such as so-called signal-integration techniques, have been developed but they are more suitable for broader cross-commodity portfolios which are outside the scope of this paper. However, it was also shown that such techniques can also be used to construct a systematic long-short portfolio within the sector of energy futures.[14]

[14] See A. Fernandez-Perez, A.-M. Fuertes, and J. Miffre (2021), "The risk premia of energy futures", Energy Economics, 102, pp.1-13. The authors use seven common systematic signals to construct long-short portfolio made up of ten energy futures, including several grades of oil, refined products, natural gas, coal, and electricity. However, the dominant factors in the portfolio are still carry and momentum. Another important factor is speculative and hedging pressure, which we will address in detail in one of our future articles.

The Reaction Function

Another critical element in constructing systematic trading signals is position sizing. While momentum and carry tend to buy the highs and sell the lows in anticipation of the trend to continue, it pays to know the limits. If the signal is already too strong and the price moves too far outside of its normal range, then the likelihood that supply and demand will adjust increases along with the possibility of a rapid correction. Similarly, in the carry strategy when contango is too steep and storage becomes economical even for the most expensive storage tiers, then most of futures selling has already been completed. Likewise, if backwardation is too strong, then most of futures shorts have already been covered, and the pressure on price starts waning.

To capture this phenomenon, systematic traders use the concept of the signal transformation, or the reaction function. These functions allow position to increase up to a certain so-called inflection point, where it reaches it maximum, but if the signal strengthens beyond the inflection point, then the position is gradually reduced. Figure 5 shows two examples of such functions where the size of the trading position is graphed versus some normalized value of a systematic signal, for example, the standard deviation of a signal. One non-linear exponential reaction function is chosen in such a way that it reaches its maximum and minimum when the signal is exactly one standard deviation strong.[15] The other function takes a more practical piece-wise linear form, as it directs the position to be held intact for the range of signal values, also reducing transaction costs.

Figure 5: Reaction Functions (RFs)

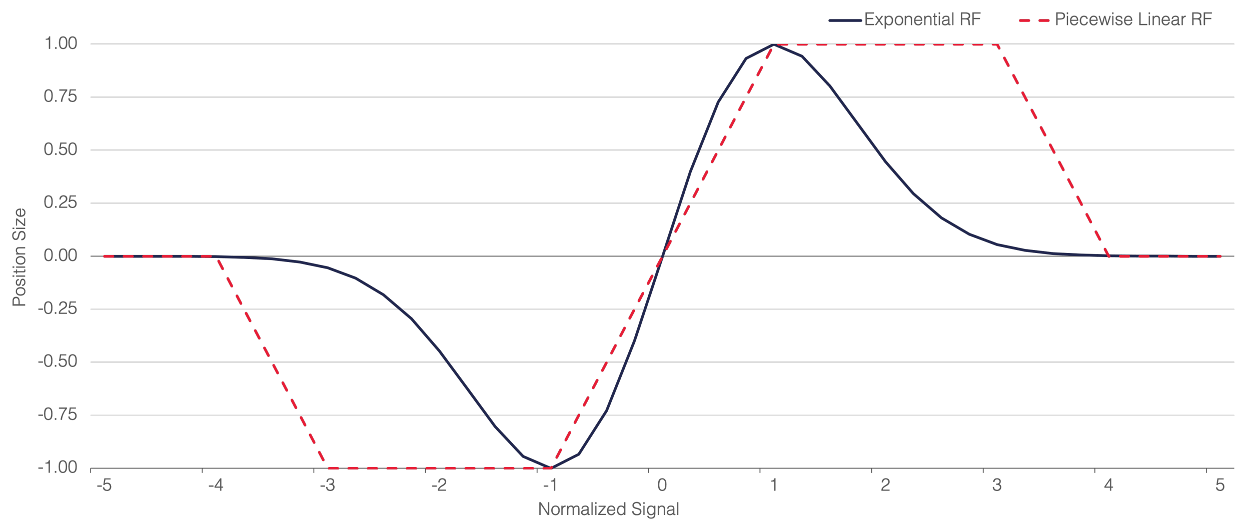

Source: Author's illustration.

Conclusion

This article presented the main systematic signals which are commonly used by algorithmic traders in the oil market. While initially these signals were used primarily by CTAs, nowadays they are often utilized by fundamental traders as well. In contrast to similar technical signals in financial markets, these indicators are particularly powerful in the oil market where they have their roots in fundamentals, and, in particular, in the theory of storage.

[15] This nonlinear reaction function is specified by $R(x) = xe^{0.5(1-x^2)}$ where x is a systematic trading signal scaled by its volatility.

We conclude by highlighting that many algorithmic strategies are managed across multiple asset classes. These strategies tend to have long-optionality profiles where one can tolerate relatively frequent small losses during range bound markets but reap abnormal positive returns when the market makes a large move in either direction. The manager of such a diversified systematic portfolio does not necessarily expect all assets to perform well during the same environment, as a few large winners are usually sufficient to pay for many small losers. For example, while the performance of systematic strategies in the oil market last year was poor, broader systematic portfolios were still profitable driven by stellar momentum in fixed income and equity markets. In the next article, we describe one of the most popular of such portfolios, known as a risk-parity fund, and why oil futures play an important role in its construction.