Oil Risk Premia under Changing Regimes

Ilia Bouchouev, Ph.D.
Managing Partner, Pentathlon Investments; and Editorial Advisory Board Member, *Global Commodities Applied Research Digest*

Lingchao Zuo
Senior Quantitative Analyst, National Grid

Systematic commodity risk-premia strategies have been popular among asset allocators and extensively studied by researchers. It is not as widely known though that the disproportionally large share of returns in such diversified commodity portfolios is attributed to energy futures. We show that even simple signals supported by the economics of oil storage and transportation arbitrage generate superior returns when applied to oil futures alone. The challenge is to be mindful of structural regime shifts that are prevalent in oil markets.

The Evolution of Oil Regimes

The widely popular concept of risk premia suggests that over time one can systematically extract positive returns either by holding financial assets like stocks and bonds, or by investing directly in one of the common factors driving returns, such as momentum, carry, and value. Applying these concepts to commodities has turned out to be less straightforward. Long-term commodity prices tend to gravitate towards marginal costs of production while nearby futures fluctuate based on the economics of storage. Specifically in the energy markets, these primary fundamental drivers are anything but stationary.

Energy futures play a rather special role even within commodities: not only because they represent the largest tradable commodity market, but also because of their disproportionally large contribution to the performance of many systematic commodity investments, both positive and negative. Energy markets are always changing along with shifts in consumer preferences and new sources of supplies requiring additional infrastructure which, in turn, causes frequent structural changes in the dynamics of prices. Systematic traders call it a regime change, which makes backtesting over long historical lookbacks not only irrelevant, but often even misleading.

The first long-term structural oil regime started with the introduction of futures trading and lasted for over twenty years. It marked the golden age of long-only investors who were able to successfully capture the structural discount offered to them by producer hedgers, as was suggested nearly a century ago by Keynes (1930) and Hicks (1939) in their theory of "normal backwardation." In addition, oil investments brought additional inflation hedging benefits as many financial assets, valued based on their discounted future cash flows, drop when inflation unexpectedly spikes. The story of this industry is well documented, and we refer interested readers to Till and Eagleeye (2007) and Greer *et al*. (2013).

To the big disappointment of long-only investors, however, the entire Keynesian risk premia accrued for over twenty years has been lost during the following decade as their own upward pressure on futures caused energy markets to flip from backwardation to the new regime of the structural contango. The oil risk premium became negative, and investors' losses turned into the remuneration of the storage companies for providing the service of storing oil, effectively outsourced to them by financial investors. Such compensation, effectively realized by investors in the form of negative roll yield, was collected by physical traders via the carry trade which provided the ultimate linkage between physical and financial markets, the regime defined in Bouchouev (2012) as "normal contango." The structural break around 2004 was also confirmed by a more rigorous statistical analysis by Tang and Xiong (2012) and Hamilton and Wu (2014).

Figures 1a and 1b illustrate how drastically different the results of long investments on oil futures were during two regimes of normal backwardation and normal contango. As much as backwardation was behind the gains in early days, contango was responsible for the losses since 2005. All results in this paper are shown for the 3rd nearby contract rolled on the last business day of the month. The difference and the contribution of the roll yield would have been even more pronounced if futures were held closer to expiry.

Figure 1a
Cumulative WTI Excess and Roll Returns During "Normal Backwardation" (1983-2004)

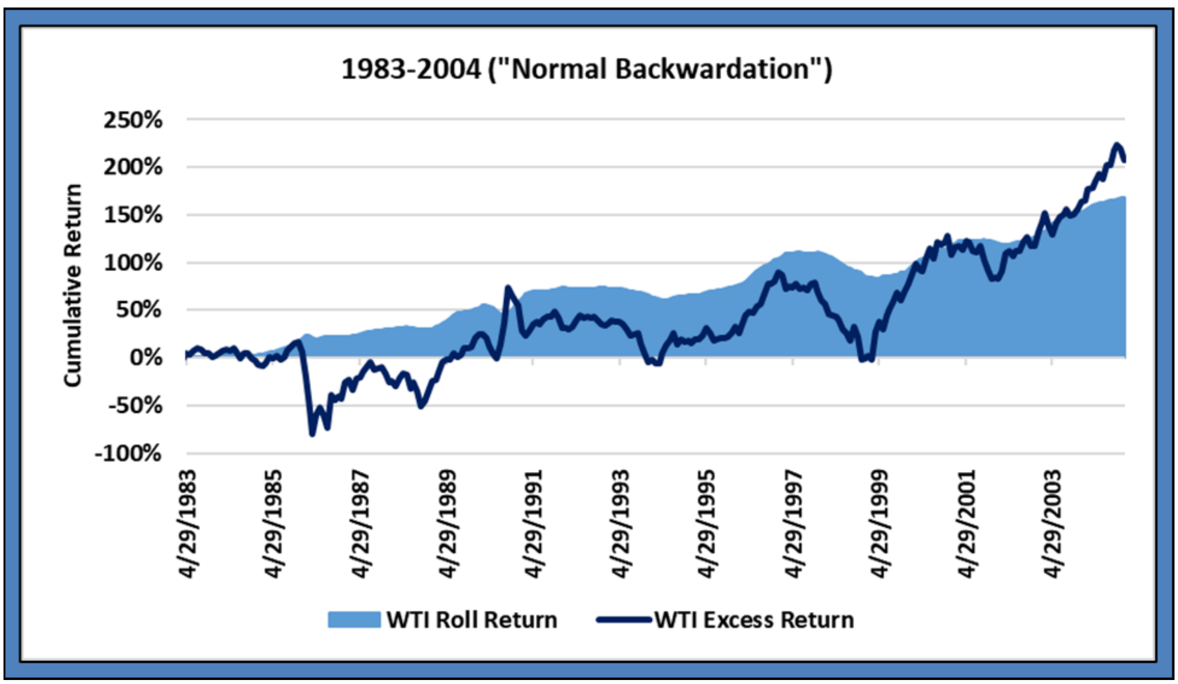

Sources: CME, EIA data, authors' calculations and graphic.

Note: Returns are for 3rd nearby WTI contract (m3) rolled on the last business day of the month.

Figure 1b
Cumulative WTI Excess and Roll Returns During "Normal Contango" (2005-6/30/2020)

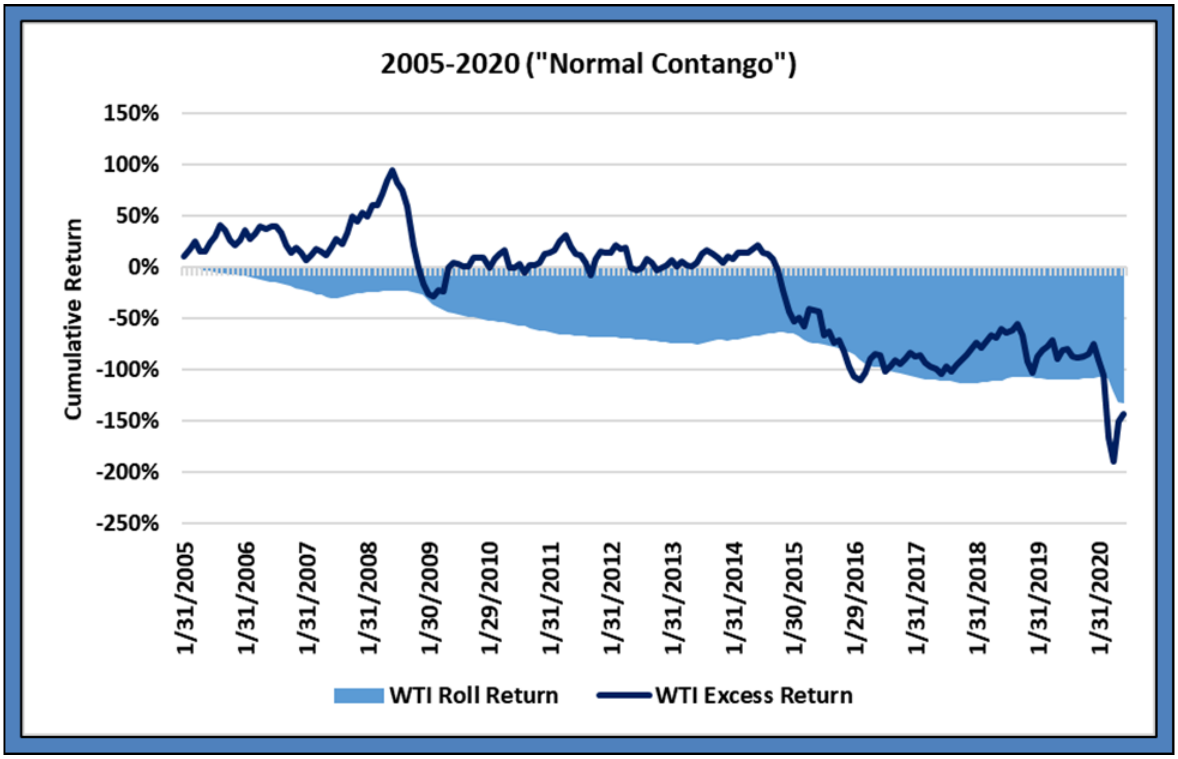

Sources: CME, EIA data, authors' calculations and graphic.

Note: Returns are for 3rd nearby WTI contract (m3) rolled on the last business day of the month.

In a recent paper, Bouchouev (2020) also suggested that the oil market will likely move into the third structural regime, sort of a Financialization 2.0. It is characterized by the two largest market participants, the inventory hedger, and the inflation hedger, reaching the structural equilibrium between their respective hedging needs and eliminating the existence of any directional risk premium. During this last phase, the speculative capital, forced to look elsewhere, has gradually transitioned away from directional investments towards harvesting so-called, "factor risk premia," embedded instead in certain trading styles such as carry, momentum, and value.

Unfortunately for investors, with only few exceptions the typical life span for factor risk premia in the energy markets is even shorter. The dynamic nature of energy markets makes historical analysis less relevant before enough data could be gathered for a robust historical backtest. By and large, any systematic strategies based on data prior to 2016 must be taken with a great amount of skepticism. While the shale revolution started gradually impacting the energy trading landscape much earlier, another structural break might have occurred around the end of 2015 when the ban on U.S. oil exports was eliminated. Not only did it fundamentally change the relationship between West Texas Intermediate (WTI) and Brent, the industry's largest futures contracts, but it also opened the door to the new regime of broader interconnectedness among energy products and significantly increased product substitutability.

In addition, many of the refined products traded today represent molecules which are very different from their predecessors. At the same time, some of these products retained the same exchange tickers creating convenient but sometimes misleading continuous time series, the fact that some quantitative researchers ignore. Today's diesel and gasoil are vastly different from what they were ten years ago. The gasoline benchmark RBOB, which stands for Reformulated Blendstock for Oxygenate Blending, is just one obscure blending component delivered in a particular location, which has very little to do with the finished gasoline product we use for driving, and equally little to do with the investment product which helped to supercharge Goldman Sachs Commodity Index (GSCI)-like oil investments during the era of normal backwardation. Historical seasonality, another material source of the historical energy risk premia, is also quickly waning, as stronger global linkages and cross-regional flows dilute established seasonality in prior regional demand patterns.

We now recap the recent performance of the main energy risk premia strategies and highlight how they evolved since the beginning of the financialization era and over the last five years which we view as the most relevant period to the current regime. It is also helpful that the cumulative change in energy spot prices during these two periods were small making our conclusions unbiased to any price directionality. This is in contrast with the full sample starting from 1993 also presented for consistency with many academic studies where backtests were conducted over the long forgotten era of normal backwardation when price dynamics and sometimes even the commodities themselves were different.

Storage and Dynamic Carry

The academic literature on the construction of systematic long-short commodity portfolios is broad. It started with strategies based on traditional price-based risk factors of carry and momentum, and subsequently expanded to incorporate non-price and commodity specific factors, including hedging pressure, inventories, congestion, and seasonality. The comprehensive review of this topic is presented by Miffre (2016). Most of the studies focus on cross-sectional properties of diversified portfolios made up of 20 to 30 different commodities, even though the energy sector alone has often been the primary driver of the portfolio performance. For many strategies, the main contribution of most non-energy commodities was in adding diversification and improving the denominator of the portfolio Information Ratios.

In this paper, we only discuss primary petroleum price-based strategies of carry, momentum, and value. The usage of non-price data such as positioning or inventory usage is arguably more interesting, but such strategies are also inherently more complex and presenting them simplistically could cause more harm than add value. While standard non-price data sources, such as weekly fundamental U.S. Energy Information Administration (EIA) reports or U.S. Commodity Futures Trading Commission "Commitments of Traders" (COT) positioning reports do contain some useful information, these signals work much better in combination with other proprietary inputs, including high-frequency storage data and over-the-counter data on hedging flows.

Among many proposed signals, the one that stands out over a long period of time is oil carry. As Figures 1a and 1b highlight, carry drives the roll yield which has been the primary force behind the initial success and subsequent struggle of many directional commodity investments. Our most basic carry strategy goes long or short depending on the direction of the carry, measured in this paper by the spread between the third and the twelfth contracts. We trade the contract defined as the third nearby on the last business day of the month which becomes the second nearby during the following month, thereby avoiding additional noise related to futures expiry. The twelfth nearby contract was chosen as the longest maturity contract with the deepest price history.[1] Even without any optimization such a trivial carry strategy applied to a single WTI contract would have generated an impressive 17.2% annualized return with a 0.50 Information Ratio since 1993. The performance has been very robust over time, including the recent new regime.

Such long-term success of the carry signal is rooted in the theory of storage. Carry is a proxy for inventories which play the crucial role for the dynamics of a storable commodity. A contango market incentivizes storing inventories, while backwardated markets force stocks to draw. The role of inventory hedgers, who are one of the largest energy market traders, is to translate carry signals into the directional price pressure. If contango covers the cost of storage, then the storer can buy physical barrels and sell futures, therefore, putting downward pressure on the futures market. If storage becomes uneconomical, then the inventory hedger starts buying back short futures, causing upward pressure on futures. The cost of storage is directly determined by the carry and the process repeats rather mechanically.

Another popular commodity risk premium strategy, price momentum, lacks any robust foundation despite numerous supporting theories typically borrowed from the equity market but largely rejected in commodities. Perhaps the most viable explanation of the oil momentum is that it often behaves as a side-effect of storage which tends to create some persistence as supply and demand are slow to adjust. Momentum by itself no longer works well in liquid commodity markets with many quantitative hedge funds shifting towards less liquid alternative markets where momentum can still capitalize on relative illiquidity and capture some hedging imbalances. Babbedge and Kerson (2020) provide a good illustration of how the popular measure of the market, inherent trendiness, has shifted from liquid to alternative commodity futures.

Many systematic energy traders now seek the “holy grail” of risk premia in combining and overlaying multiple signals and sometimes applying more complex statistics, including machine learning techniques. One of the best blended signals for oil is so-called carry momentum. Even though price momentum no longer works well by itself, applying it to the term-structure of futures makes more fundamental sense as it adds more dynamism to the proven fundamental signal of carry. The biggest drawback of the basic carry strategy is its slowness as oil markets do not flip between contango and backwardation frequently. The application of the basic momentum signal directly to the annual time spread or to the measure of carry rather than to price forces the carry strategy to become nimbler and keep up with increasing speed of market changes. One can think of the basic carry measuring the current state of the storage while the change in carry functioning as an early signal about likely change in inventories.

Table 1 summarizes our results for traditional price driven systematic energy strategies, comparing their performance over three periods, “full sample” from 1993 until June 2020, the period of “financialization” and the "normal contango" from 2005, and the current "post-shale" regime covering approximately the last five years. The curve momentum strategy continues to achieve a remarkable success. Over the period of nearly thirty years, it generated an unprecedented 25.5% annualized return with a 0.74 Information Ratio for WTI. Moreover, during the latest "post-shale" period the Information Ratios for WTI and Brent were even higher at 1.13 and 1.30, respectively. For completeness, we also present results for refined products which do not materially contribute to the portfolio besides some basic diversification benefits. As mentioned in the endnote, our carry was defined simplistically without capturing any seasonal effects, but the contribution of refined products could potentially be materially improved with a more complicated seasonal carry model.

It has also been recently documented in a rigorous academic study by Boons and Prado (2019) that carry momentum significantly outperformed both carry and momentum even for the broader portfolio of commodity futures. For years, carry momentum has also been a popular strategy among foreign exchange traders.

Table 1
Systematic Carry, Momentum, and Carry-Momentum Strategies for Energy Futures Under Different Regimes

		"Full Sample"			"Financialization"			"Post-Shale"		
		CAR	MOM	CAR-MOM	CAR	MOM	CAR-MOM	CAR	MOM	CAR-MOM
WTI	Return	17.15%	7.71%	25.51%	15.57%	11.81%	33.65%	8.86%	12.34%	52.61%
	Vol	34.49%	34.50%	34.47%	37.83%	37.84%	37.79%	46.69%	46.69%	46.58%
	IR	0.50	0.22	0.74	0.41	0.31	0.89	0.19	0.26	1.13
Brent	Return	16.56%	2.85%	17.99%	14.18%	3.97%	26.33%	5.01%	1.89%	52.84%
	Vol	32.39%	32.41%	32.39%	34.44%	34.45%	34.42%	40.77%	40.77%	40.64%
	IR	0.51	0.09	0.56	0.41	0.12	0.77	0.12	0.05	1.30
ULSD	Return	8.41%	2.28%	8.92%	8.61%	3.59%	7.10%	10.37%	11.59%	22.63%
	Vol	32.00%	32.00%	32.00%	32.34%	32.34%	32.34%	35.50%	35.50%	35.48%
	IR	0.26	0.07	0.28	0.27	0.11	0.22	0.29	0.33	0.64
RBOB	Return	6.90%	5.99%	5.11%	2.36%	4.65%	-3.02%	-13.57%	13.15%	7.05%
	Vol	34.98%	34.98%	34.98%	37.73%	37.73%	37.73%	44.04%	44.04%	44.05%
	IR	0.20	0.17	0.15	0.06	0.12	-0.08	-0.31	0.30	0.16
Portfolio	Return	12.26%	4.70%	14.38%	10.18%	6.00%	16.02%	2.67%	9.74%	33.78%
	Vol	26.63%	28.39%	24.30%	27.74%	30.12%	25.18%	33.16%	34.59%	32.20%
	IR	0.46	0.17	0.59	0.37	0.20	0.64	0.08	0.28	1.05

Sources: CME, ICE, EIA data, authors' analysis, and calculations.

Note: The data set is from January 1993 through June 2020. Log-returns are for 3rd nearby (m3) contract rolled on the last business day of the month. Information Ratio (IR) is Return / Annualized Volatility. Carry (CAR) is measured by m3-m12. Heating Oil and Unleaded Gasoline were used as predecessors for Ultra Low Sulfur Diesel (ULSD) and RBOB, respectively. Momentum (MOM) and Carry-Momentum (CAR-MOM) compare the current price and carry, respectively, to their 20-day moving averages.

As good as the recent abnormal performance of the carry momentum is, it also highlights the problem that many statistical studies hide. Energy futures have much larger tails which make cumulative statistics sensitive to whether the strategy was right or wrong during the extreme event like the one that happened in the spring of 2020. Such disproportionally large contributions from extreme events often makes strategies based on historical moving averages easy to overfit. Figure 2 shows that about half of five-year profits from our simple carry momentum strategy came from correctly capturing both the large move down in prices in Mar-Apr 2020 along with the equally fast recovery.

Figure 2
Cumulative Return of Carry-Momentum Strategies for WTI and Brent Futures

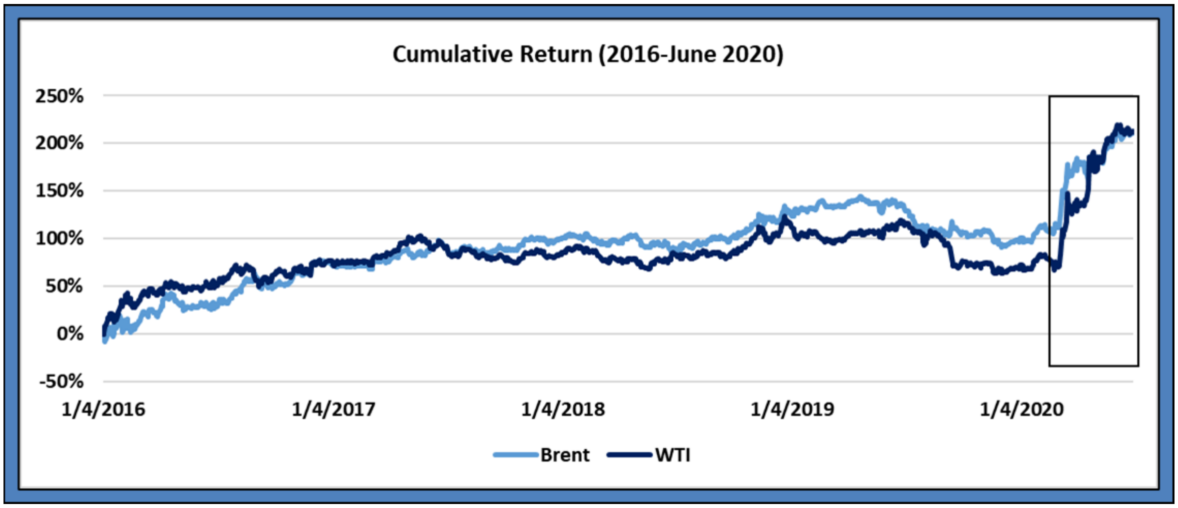

Sources: CME, ICE, EIA data, authors' calculations and graphic.

Nevertheless, even if one discounts the contribution from this abnormal period, the carry momentum strategy continues to stand out among many permutations of popular signals. This signal has also been successfully used for over a decade by one of the authors in practice as an overlay to fundamental trading strategies.

Value and Mean-Reversion

In simple terms, the value strategy for a single asset can roughly be defined as the opposite of the time series momentum. In other words, value is just the contrarian signal, which sells an appreciated asset and buys the depreciated one. Therefore, when momentum works then the value signal typically does not, and vice versa. One exception is a blended signal where momentum trades up to a certain threshold beyond which the strategy flips to the contrarian signal. Another way to combine momentum with value is to blend trading frequencies as oil momentum tends to work better on shorter frequencies which are followed by some mean-reversion.

Unlike momentum, value does have a fundamental rationale as prolonged high and low prices force not only change in a physical supply and demand, but also incentivize financial hedgers to lock in increasingly better economics of production and consumption, pressuring prices to return back to their normal range. The challenge is that such fundamental adjustments often take time and one is typically facing the headwind of the negative carry while waiting for prices to mean-revert. High spot prices tend to be associated with backwardated markets when the contrarian value signal indicates selling futures often below the spot price. Likewise, low prices tend to occur when the market is in contango, so the contrarian value signal would lead to buying futures typically above spot prices. In both cases, fading the previous market moves must overcome the pressure from the negative carry.

Where the value style works much better is in trading closely related energy pairs, sometimes referred to as a commodity statistical arbitrage. The rationale makes perfect sense as many energy commodities are linked via strong economic relationships. After the U.S. lifted the ban on oil exports, WTI and Brent became much tighter linked by the economics of the shipping arbitrage, incentivizing flows in both directions depending on the level of the spread. Likewise, the spread between refined products and crude oil is largely driven by the profit margin of a refiner, and the spread between different grades of crude oil is dictated by the economics of the oil blender. Cross-commodity spreads are often driven by the optionality owned by the asset owner. Asset owners are then incentivized to monetize their optionality leading to the price dynamics akin to delta-hedging of the long option on the spread. For example, the refinery will increase (decrease) run rates as the spread between the basket of refined products and crude oil widens (narrows). The behavior of these primary market participants creates statistical boundaries which approximate the economics of the option's owner.

Despite its popularity and success among physical traders, statistical arbitrage in commodities has largely stayed below the radar of systematic traders. Historically, quantitative backtests generated by such contrarian rules did not look very appealing. Contrary to momentum, the distribution of returns from mean-reversion strategies tends to have a negative left tail, much like being short an option which makes them difficult to market to potential investors. These strategies often generate steady profits but incur infrequent but large losses during extreme events when one of the legs in the pair disconnects because of short-term fundamental squeezes or logistical bottlenecks. In the past, physical traders with access to some other fundamental data were able to forecast such events with reasonable probabilities which was sufficient to turn these strategies from neutral to very profitable.

We now posit that such pair strategies rooted in the economics of asset owners could become the hallmark of the new oil regime, the regime of cross-product integration. With the shale revolution, enormous amounts of new energy infrastructure have been built. Not only is WTI now better connected with Brent, but the same applies to refined products and even natural gas. Shale also increased product substitutability as many refined products can now be used for multiple competing usages. In addition, the proliferation of many new pricing points connected with recently built pipelines allows one to potentially create a sufficiently diversified statistical arbitrage portfolio just within the energy sector alone. More unique fundamental data, such as local inventories, pipeline flows, or cargo tracking are now available to support and backtest the performance of these "quantamental" statistical arbitrage strategies.

As an example, we show the results of a convergence strategy for the spread between WTI and Brent. The strategy sells and buys the spread when its current level exceeds or falls under its moving average by a certain threshold. Figure 3 shows the volatility-adjusted profit/loss (Information Ratio) and its robustness with respect to the chosen threshold on the x-axis and to the lookback period on the y-axis. The strategy is not optimized and can be further improved by making the sizing and trading threshold more dynamic, or by conditioning on various fundamental inputs. Our intent here is not to show the best-looking statistics, but rather to highlight the important and often overlooked concept with a solid fundamental rationale, which we expect to work well in the new regime of energy interconnectedness. Similar strategies could also be constructed for other linked energy pairs, including refined products to further improve the performance of such an energy pairs portfolio. We should also note that all of our strategies are based on daily settlements, and given the existence of liquid TAS (trading at settlement) products for WTI and Brent, the slippage and its impact on the performance is minimal, reducing Information Ratios by less than 0.1.

Figure 3

Information Ratios for a WTI-Brent Pairs Convergence Strategy (Jan 2016-Jun 2020) for Different Trading Thresholds and Lookback Periods

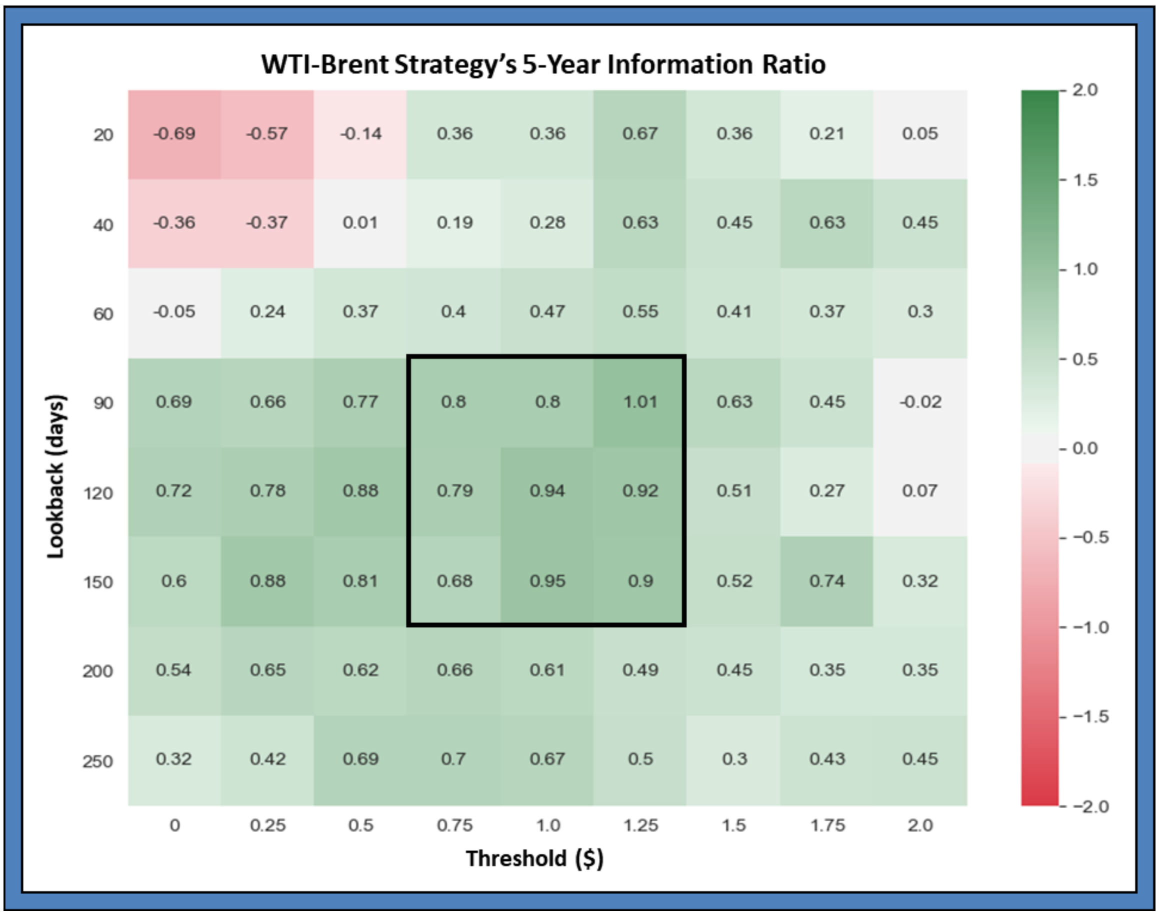

Sources: CME, ICE, EIA data, authors' calculations.

Note: Information Ratio is Profit (Loss) / Annualized Volatility, shown for a WTI-Brent convergence strategy for different lookback periods and trading thresholds. The strategy buys/sells 3rd nearby WTI-Brent spread when the current spread value is below/above its moving average by more than a given threshold.

We conclude by emphasizing that for systematic energy trading to succeed, it is so much more important to spend time on identifying relevant regimes than on generating the best looking backtest over any fixed lookback period. Energies are always on the move, and this is the structural property of the sector, where nothing will work for a long time and regimes will continue to change frequently.

Endnote

1 To properly capture seasonality of refined products, it would have been more accurate to replace the twelfth nearby futures with the fifteenth nearby so that the carry is defined for the same calendar month one year apart. However, availability of such data is more limited requiring more complex seasonality models which are beyond the scope of this paper and left for future studies. Our primary focus here is on the crude oil market where our definition of carry is meaningful and widely used by traders.

References

Babbedge, T. and J. S. Kerson, 2020, "Part 2: Trend's not Dead: (It's Just Moved to a Trendier Neighborhood)," *Global Commodities Applied Research Digest*, Vol. 5, No. 1, Summer, pp. 141-148.

Boons, M. and M. P. Prado, 2019, "Basis-Momentum," *The Journal of Finance*, Vol. 74, No. 1, pp. 239-279.

Bouchouev, I., 2012, "Inconvenience Yield, or the Theory of Normal Contango," *Quantitative Finance*, Vol. 12, No. 12, pp. 1773-1777.

Bouchouev, I., 2020, "From Risk Bearing to Propheteering," *Quantitative Finance*, Vol. 20, No. 6, pp. 887-894.

Greer, R., Johnson, N. and M. Norah, 2013, Intelligent Commodity Indexing, New York: McGraw Hill.

Hamilton, J. and J. Wu, 2014, "Risk Premia in Crude Oil Futures Prices," *Journal of International Money and Finance*, Vol. 42, pp. 9-37.

Hicks, J. R., 1939, Value and Capital, Oxford: Oxford University Press.

Keynes, J., 1930, A Treatise on Money, Vol. II, London: MacMillan.

Miffre, J., 2016, "Long-Short Commodity Investing: A Review of the Literature," *Journal of Commodity Markets*, Vol. 1, pp. 3-13.

Tang, K. and W. Xiong, 2012, "Index Investment and the Financialization of Commodities," *Financial Analysts Journal*, Vol. 68, No. 6, pp. 54-74.

Till, H. and J. Eagleeye (Editors), 2007, Intelligent Commodity Investing, London: Risk Books.

Author Biographies

ILIA BOUCHOUEV, Ph.D.
Managing Partner, Pentathlon Investments

Dr. Ilia Bouchouev is the former president of Koch Global Partners where he managed the global derivatives trading business for over twenty years until his retirement in 2019. He is currently the managing partner for Pentathlon Investments. Dr. Bouchouev is an Adjunct Faculty at New York University where he teaches Trading Energy Derivatives for the Mathematics in Finance Program at Courant Institute of Mathematical Sciences, and a Research Associate at Oxford Institute for Energy Studies.

He has published in various academic and professional journals on the economics of energy trading and derivatives pricing. Dr. Bouchouev has a Ph.D. in Applied Mathematics. He is a frequent commentator on commodities on Twitter @IliaBouchouev and on LinkedIn.

LINGCHAO ZUO
Senior Quantitative Analyst, National Grid

Mr. Lingchao Zuo has worked in the energy industry since 2011. He is currently a senior quantitative analyst at National Grid. His interests are in quantitative analysis, derivatives pricing, risk management and trading in the markets of natural gas, power, crude oil and refined products. He holds an M.S. degree in Mathematics in Finance from Courant Institute of Mathematical Sciences, New York University where he studied systematic energy strategies.