The Misleading Value of Measured Correlation

Babak Mahdavi Damghani
Risk Methodology, Credit Suisse, e-mail: bobby.damgani@cantab.net
Daniella Welch
Risk Methodology, Credit Suisse
Ciaran O'Malley
Risk Methodology, Credit Suisse
Stephen Knights
Risk Methodology, Credit Suisse

Abstract
Within the framework of the financial industry, when representing relationships between assets, correlation is typically used. However, academics have long since questioned this method due to the plethora of issues that plague it. Indeed, it is thought that cointegration is a natural replacement in some of the cases as it is able to represent the physical reality of these assets better. However, despite this general academic consensus, financial practitioners refuse to accept cointegration as a better tool, or even the lesser of two evils. This technical report attempts to explain this bias, specifically focusing on the various consequences of model selection considering the new and challenging regulatory environment and suggests a practical replacement hybrid alternative to both cointegration and correlation.

Keywords
correlation, cointegration, mean reversion, cointelation

Introduction
Academics have shown considerable interest in cointegration as a replacement for correlation for many years. They believe that the problems with correlation make its use, in finance, often nonsensical. The implications of this seem to be that cointegration would improve the way that the relationships between assets are assessed in these particular instances, and therefore, more profit might be produced. If this is the case, then it is difficult to understand why finance professionals have not altered their practices. In this article we will attempt to address this question. To begin with, the parameters and logical argument need to be outlined. Then, our assumptions will be analyzed, to reach the conclusion that, whilst many factors influence choice, the financial world should cast some of these factors aside, and begin to utilize cointegration as a superior alternative to correlation in at least more instances.

Logical argument:

1. Correlation is used in certain areas of the financial world, although in these particular cases, many academics have stressed their irrelevance.
2. Cointegration, although not perfect, seems to be a better candidate for modeling these particular areas of the financial world.
3. Financial professionals wish to be successful.

Therefore:

4. Finance professionals should use cointegration rather than correlation in these particular areas of finance.

If the three premises hold, then it follows that the conclusion does as well. This article will attempt to prove that this is indeed the case. Premise one and two will be dealt with, in the main, side by side in order to provide an appropriate comparison. Firstly, the variables and models for both of these methods will be demonstrated so that technical clarity is maintained.

The variables
Below is a summary of the two models, which we will discuss in the rest of this article. Model A is the traditional correlation model. Model B is a mean-reverting model inspired by the Ornstein–Uhlenbeck (OU) process, i.e., the

cointegration model. In order to explain the difference between these, we first outline the variables that will be used:

- dx_t^1 / x_t^1 denobltes the returns of asset 1.
- ρ is the correlation between two assets.
- σ_t^1 and σ_t^0 denote the volatility of asset 1 and some other volatility at time t, respectively.
- Define W_t as a symmetric error distribution (traditionally, a standard Brownian motion $W_t \sim N(0, t)$) and $W_t^{\perp}, W_t^{\perp,2}$ as two independent distributions of the same type.
- We call θ the speed of mean reversion, and μ the long-term spread between these two assets (Brent and WTI).

The different models

Model A: the two-dimensional geometric Brownian motion

$$dx_t^1 / x_t^1 = \sigma_t^1 dW_t$$

$$dx_t^2 / x_t^2 = \sigma_t^1 \rho dW_t + \sigma_t^2 \sqrt{(1 - \rho^2)} dW_t^{\perp}$$

We will use Pearson's correlation coefficient because it is the most used in the industry:

$$\rho = \frac{E\left[\left(X_1 - \mu_{X_1}\right)\left(X_2 - \mu_{X_2}\right)\right]}{\sigma_{X_1}\sigma_{X_2}}$$

Model B: the OU spread model

$$dx_t^1 / x_t^1 = \sigma_t^1 dW_t$$

$$dx_t^2 = \theta(x_t^1 - x_t^2 + \mu)dt + \sigma_t^o dW_t^{\perp}$$

Review of statistical concepts

This section outlines the key statistical theories of correlation and cointegration that will be used throughout this article.

Correlation is one of the oldest concepts in statistics and can arguably be traced to Gauss in 1823 [1]. The method was, however, formalized by Pearson [2, 3]. Despite its simplicity, the method has several limitations and the coefficient is often interpreted incorrectly. Firstly, the Pearson coefficient is only accurate if the correlation relationship is linear and concurrent; hence it cannot capture if a change in X in period 1 is correlated with a change in Y in period 2. Secondly, Anscombe [4] illustrates that the correlation result can be significantly affected by outliers in the data, which may represent incorrect measurements or extreme events. Anscombe [4] also identifies four data sets with identical correlation and other statistical properties (such as skew and mean) that can have completely different distributions of data and so demonstrates that there is a limit to how informative correlation can be. Therefore, correlation is clearly a very useful statistical method; however, it is limited by the assumptions underlying the estimation and so should only be applied when those assumptions hold.

A cointegration relationship describes a set of variables which share a common, long-run equilibrium and will revert to that equilibrium after any short-run deviations. Granger (1981)[1] defined cointegration as a linear combination of two non-stationary $I(1)$ variables which is stationary, $I(0)$.[2] Formally, if $(X_1, X_2) \sim I(1)$[3] and a linear combination of $Z = \alpha + \beta_1 X_1 + \beta_2 X_2$[4] is $Z \sim I(0)$ then the variables are cointegrated. The cointegration relationship is important because earlier econometric theory required variables to be stationary for reliable inferences [5], however, economic variables are rarely stationary and economic theory is predominantly interested in the long-run equilibriums captured by cointegration. This usefulness is emphasized by the two most important cointegration articles – both ranked in the top 10 most cited economics, finance, and econometrics articles [6]. However, there are limitations to cointegration, the main one being analyzing whether a set of variables is in fact cointegrated. Testing for cointegration has developed beyond the original Engle Granger [7] test to overcome many of its shortcomings, yet the widely used Johansen [8, 9] method is still not entirely robust. The principle problem with the test is that it tends to find cointegration when it is not present. The first cause of this is that the test cannot effectively distinguish between cointegrated relationships and near-cointegrated relationships, as Hjalmarsson and Osterholm [10] and Cheung and Lai [11] find. Second, the test result is very sensitive to the structure of the vector autoregression part of the methodology, as Gonzales [12] illustrates, and the result can be manipulated to show cointegration by changing that structure. Therefore, cointegration is a very important econometric method to model economic and financial variables, yet its effectiveness is limited by the testing procedure.

Why a measured correlation of the returns of a mean-reverting process is misleading

Definition: Misleading

When trying to model a relationship between assets, practitioners often use correlation Model A, when the correct dynamic is actually Model B. Model A has been popularized through Harry Markowitz's modern portfolio theory (MPT), but unfortunately the underlying assumptions are often forgotten, leading to the model's use in situations that conflict with the theory's conditions. As we can observe in Model B, correlation is not a relevant parameter. We illustrate how the irrelevance of that parameter can be deceptive when evaluating meaning.

Instances of misleading correlations

There are negative consequences to using Model A when we should be using Model B, in terms of meaning. For instance, when we think of correlation we are often tempted to say that if two assets are strongly positively correlated, then their returns should, on average, move in the same direction. In a purely mathematical sense this is not necessarily true. Wilmott gives a brilliant counterexample in his latest book, *Frequently Asked Questions in Quantitative Finance* [13]. The point of Wilmott's example is to illustrate that a strong positive correlation does not necessarily imply that two stochastic processes "go in the same direction" and vice versa. When comparing Model A and Model B, we need to develop a scientific method that can compare the added value of measured correlation in assessing "meaning of relationship." With this objective in mind, we have simulated three stochastic processes – two of which are modeled through A's dynamic and two designed to follow B's dynamic. Therefore, Models A and B have one process in common. We will call this process the "driver's dynamic" ($dx_t^1 / x_t^1 = \sigma_t^1 \; dW_t$). The name arises from the distribution of the second dynamic in Model B, which is conditioned with these returns. In order to understand the "meaning of relationship," we will look at the cumulated differences between these three processes, and define the parameters of our model such that by expectation, the cumulated difference is the same between these two pairs of processes. We will then look at the measured correlation in such situations and reach a conclusion with regard to the behavior of measured correlation. Since we are in a stochastic environment, we devise a Monte Carlo method and make our conclusions with respect to the density of measured correlation. Figure 1 represents the estimated correlation density of the two simulated stochastic processes under Model A and an arbitrary $\rho = 0.92$. Figure 2 represents the correlation density estimated within the dynamics of Model B, keeping x_t^w equal to that in Figure 1 and with $\theta = 0.1$ chosen in order to map, by expectation, the cumulated returns of Model B onto those of Model A. Figure 3 represents the induced value difference with respect to choosing Model A instead of Model B. We can see that Figure 3 represents a symmetric distribution around a mean of 0, which suggests that with the same final difference, Model A has a much higher correlation than Model B (Figure 1 vs. Figure 2). Figure 4 illustrates three processes, where the blue and red

Figure 1: Histogram of a measured correlation stochastic process with $\rho = 0.92$.

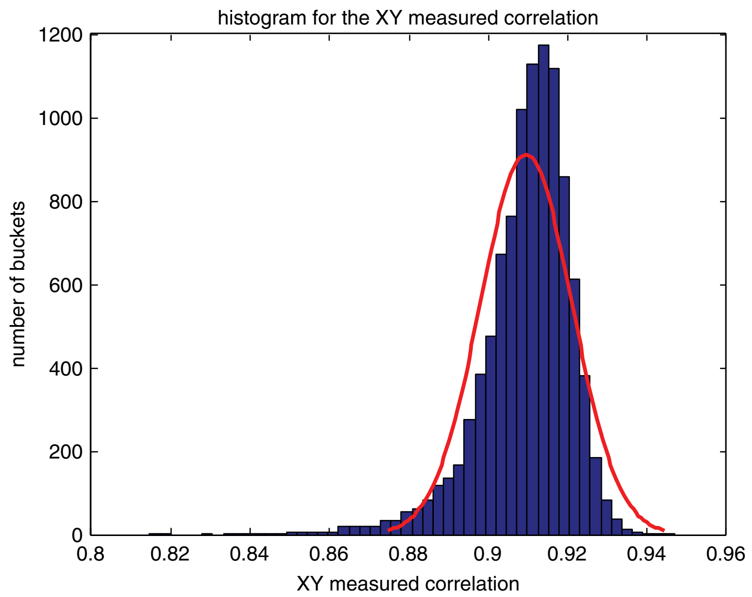

Figure 2: Histogram of a measured correlation stochastic process with the same VaR as Figure 1 but generated by Model B.

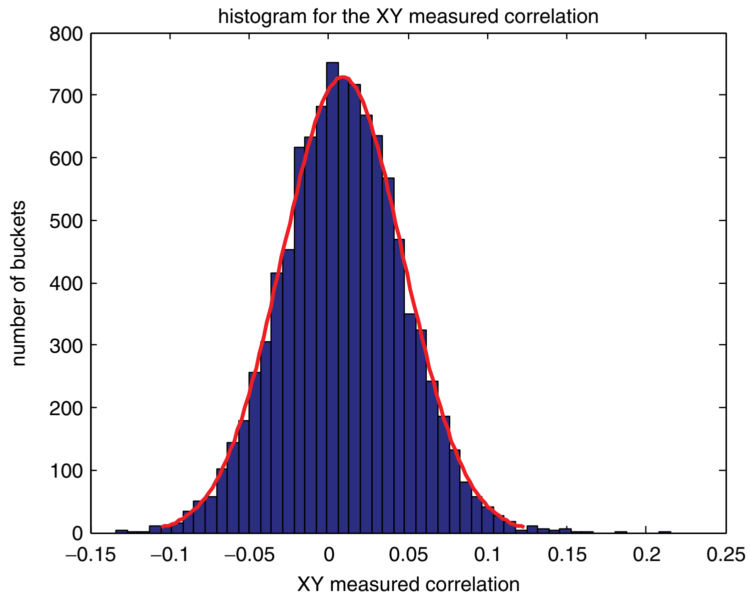

Figure 3: Symmetric distribution around a mean of 0, suggesting that with the same VaR, Model A has a much higher correlation than Model B.

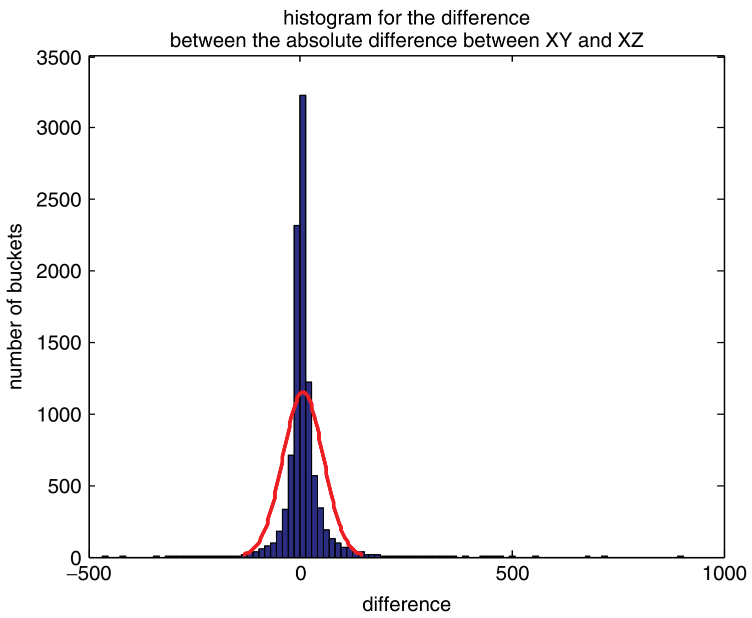

Figure 4: Three processes, with the blue and red processes being correlated, and the blue and green processes being cointegrated.

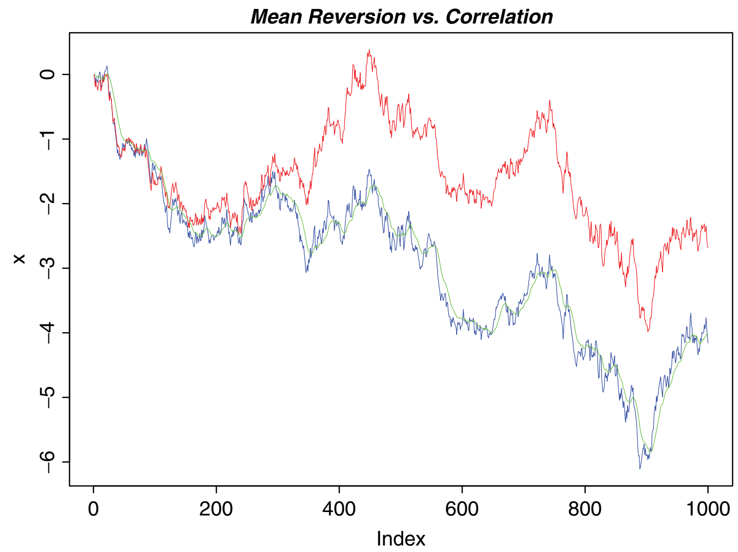

processes are correlated, and the green process is cointegrated with the blue one. The measured correlation between the red and blue returns is 0.92, whereas the measure correlation between the green and blue returns is only 0.21, which is not very intuitive given how much closer the blue and green processes stay together and do not depart. Note that where, by chance, the blue process suddenly increases or decreases significantly, the induced measured correlation for the returns of the green and blue processes increases drastically. This is because in Model B, dx_t^2 tries to always "catch up" to a process that is going up by chance already, so one gets a high measured correlation for the returns by accident: the stochastic side of the differential equation never lowers the correlation by adding random variability in the returns. Given that a high number is expected because a strong relationship is suspected (in terms of high correlation in Model A), it becomes very easy to disregard Model B in such a situation (we will explore the biological reason for this phenomenon later in this article). Similarly, imagine the induced x_t^1 happens to be neither going up nor down; in this situation in Model B, the deterministic part of the differential equation is never significant and most of the dynamic is explained by the random noise which induces a small correlation in the returns. As a consequence, one may think that there is no relationship between these two underliers. It is therefore extremely dangerous to use measured correlation for model calibration without taking careful steps to assess which model is the more accurate.

Figures 5 and 6 are examples of the correlation "function of the slope of x_t^1." In Figure 5, we get a correlation of 0.92 between the blue and the red process returns and a correlation of 0.41 between the green and the blue process returns (x_t^w is "somewhat going up"). In Figure 6, x_t^1 (the blue process) remains near its mean. The induced correlation on the returns is twice as small at 0.18. These are illustrations of how the correlation behaves with respect to the degree of "catching up" phenomenon just introduced. Figure 2 is the correlation density of the green and the blue process returns in Model B given random samples of x_t^1. The density is much wider than in Figure 1, which accounts for its non-irrelevance in terms of evaluating VaR.

Figure 5: An example of Model B yielding a high correlation with the leading stochastic process based on the latter going up by chance.

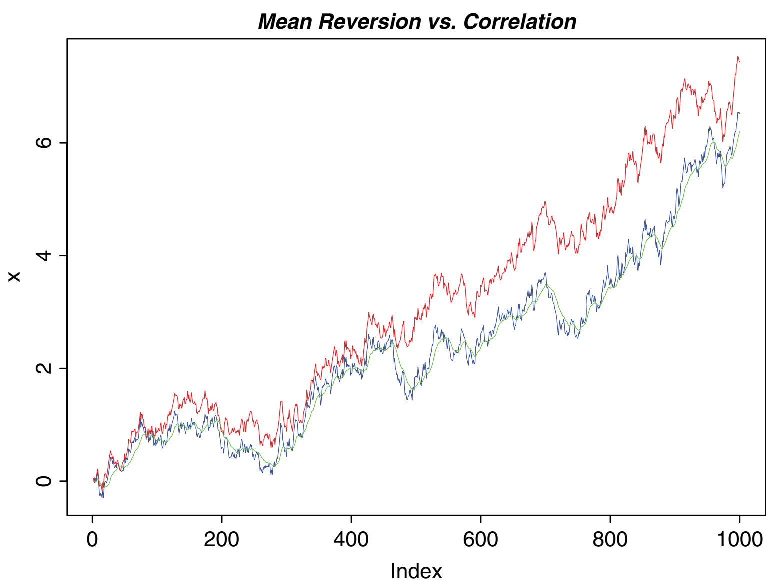

Figure 6: An example of Model B yielding a low correlation with the leading stochastic process based on the latter neither "going up nor down."

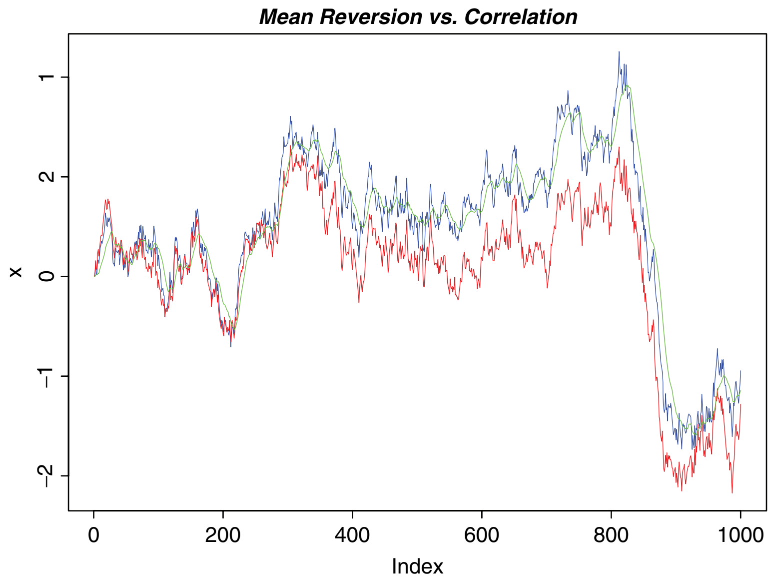

Application in the markets – how correlation fails

Commodities modeling

Consider: there is a basket between WTI and Brent – two commodities that are highly correlated in the sense of Model A. This is assumed by most quantitative models. However, in the case of WTI and Brent, this model on its own does not accurately capture the physical reality which depends on geopolitical factors, sea routing, pipelines, and storage. Because of these physical factors, the spread between these two underliers could be modeled through some sort of mean-reverting process (i.e., Model B). The concept of mean reversion is not limited to the energy market. For instance, mean reversion regularly occurs on the commodities market (e.g., soybean-related products). The whole futures curve is a mean-reverting process: the curve never stays indefinitely in contango or in backwardation. Thus, modeling the different points on the curve with a correlation parameter is very dangerous.

Equities modeling

In theory, there should be many stationary signals in the equities market. Pairs trading is a consequence of popularity induced by the concept of mean reversion. For example, in a theoretical sense, because Walmart and Target are in the same industry (one could, equally, also use Nike and Reebok or Dell and Apple in this instance), whenever one has its shares rise, the other should have its shares fall, simply because they are competitors. Note that if it were a strongly bullish market in the sector in general, then when one had its stocks rise, so would the other. However, given that they are competitors, their combined accumulated returns (linear combination of these two signals) should, in theory, be stationary.

Figure 7 illustrates the historical price changes between Walmart and Target, two pairs that ought to be cointegrated for physical reasons. The statistical test as described in Table 1 suggests that Walmart and Target are cointegrated, so, assuming a similar market environment, one could take

Figure 7: Walmart and Target time series between 2005 and 2012.

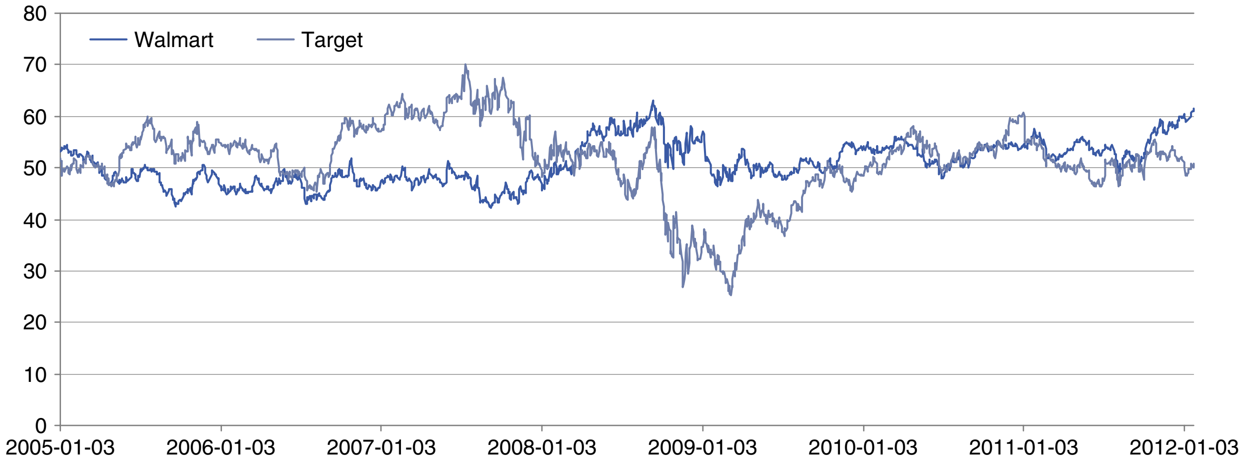

Table 1: Cointegration tests performed on Walmart and Target, two stochastic processes whose combination ought to be cointegrated for physical reasons.

Cointegration test	Test statistic	Critical values		
Johansson maximum eigenvalue test:		10%	5%	1%
with constant	15.98	13.75	15.67	20.2
with linear trend	20.92	16.85	18.96	23.7

long-term exposure on Target via Walmart, for example, despite a measured correlation of −0.2 between them.

Note that from Figure 7, the negative correlation is due to the generally flat path taken by Walmart (the driver's dynamic according to a simple lag test such as the Granger test) induces a very small correlation. Note that as developed in "Instances of misleading correlations" above, had Walmart been steadily going up or down we would have had a higher measured correlation because Target would have had to catch up with a process going up by chance.

Credit modeling

Modeling correlation between the different tranches using Model A is extremely dangerous, as observed during the latest credit crisis. Model A does not take into account the domino effect that Model B could capture. A fixed correlation across every tranche and around 0.4 is generally very suspicious. My understanding is that these base correlations have been indirectly calibrated through a type of historical stationary time series and suffer from the misleading value of measured correlation phenomenon we have discussed in this technical report. Figure 2 demonstrates that a correlation of around 0.3 can have significant risk. Thus, as 0.4 is very much within the range of highly likely measured correlation within cointegrated pairs, so a correlation of 0.4 should be worrying as well. Remember, in Figure 2 the expected correlation for mean-reverting signals is more toward 0.4 than 0.9.

Hybrid modeling (equity, commodities, and FX)

There are many commodities that have equity mirrors (oil and BP, for example). Pairs trading between commodities and commodities mirrored via equities is considered one of the simplest pairs trades one could create. Similarly, as in the previous subsection on credit modeling, a correlation of around 0.4 between a commodity and its equity mirror should be seen as very suspicious.

Likewise, if one examines the FX carry strategy (FXcs) and the emerging markets vs. developed country equity strategy (EMvsDCs) then one can see that the FXcs works as a great overlay to the EMvsDCs, which suggests that there perhaps exists very subtle cointegration-like relationships between these two asset classes. Thus, using correlation could prove to be dangerous when modeling FX with equities.

Basis modeling

Every model that uses correlation for basis spread as an input is in danger of being modeled under the wrong assumptions and exposed to the misleading value of measured correlation. The commodities market futures curve is one of the many different potential areas one should be careful of when modeling. On the equities market, the index skew basis has similar issues. On the credit market, the basis between the loan and the CDS mirror is another example of an area with potential danger when using correlation modeling.

Bias toward correlation, despite cointegration's plausibility

Correlation should be applied to short-term relationships, cointegration to longer-term, but this is not implemented in the markets

Traders are interested in always knowing the market. Given that correlation is universally used in the market for pricing options, it makes sense to use correlation when working out competitors' option pricing, rather than attempting to correctly model how an option should be priced using cointegration. Usually, implied correlation is defined on a one-year rolling basis, independent from the time to expiry of the option. This is contrary to our earlier assumption that correlation is useful in the short term but

>

that cointegration is a better model in the long term. The correlation factor for highly corrected assets should diminish as the time to expiry increases, which does not happen here. What is interesting is that traders (especially in the commodities asset class) feel that correlation is a poor measure of the true relationship, but rather than adjusting the implied correlation in their prices, they prefer to increase their commission in order to compensate for the increased risk associated with the rolling measured correlation.

It is sometimes difficult to change a model that has defined an industry for decades. However, in the same way people adjust volatility in order to account for the limitations of the Black–Scholes model, we could adjust the correlation skew in order to reconcile the Black–Scholes model with what happens on the market. What does seem odd is that people really do not take the time to mark the correlation with respect to the spread. This is what they ought to do if they do not question the Black–Scholes model in its entirety. Instead, because it is "industry practice," the correlation is marked on a yearly rolling basis. At the same time, traders understand that correlation becomes unsafe during times when historical spreads are being stretched dangerously. In these situations, because the correlation is marked according to rolling historical data independent from the spreads, and because the traders' work within banks is that of a market maker (whose job entails *always* providing a market), traders create very wide markets in order to avoid a position in which risk has not been accounted for. Thus, liquidity decreases in these situations, which is unfavorable.

Burden of proof

As we have seen, correlation has many problems, and does not work in many situations. Therefore, one could turn to cointegration. But the burden of proof seems not to be on those who propound cointegration to prove it to be correct, which they cannot, as it is impossible to prove that cointegration is worth using just because it is not correlation. As Russell [14] famously explained, one could tell everyone that there is a teapot orbiting the sun: just because we cannot prove that he is wrong, does not mean that he is right. However, unlike Russell's teapot, cointegration is not so unbelievable. There have been situations, as mentioned previously, that show it to be a better methodology than correlation. Thus, there must be some other reason that financial professionals are disinclined to utilize cointegration when it could, potentially, create more profit. Tradition, reputation, and implementation appear to be the main factors involved, as will be explained.

Tradition

Correlation has always been the main concept expounded within foundational financial theory and MPT, and is taught from the earliest level of finance. On the other hand, cointegration is barely mentioned until Masters or PhD level. Thus, many people within the financial industry will not have given any attention to cointegration because they have had no need to. In reality, it may be that they do not even realize there is another option (at least a feasible option) to correlation, let alone that the alternative is cointegration. As Herrnstein observes: "Most people, even most academics, do not have the time, training, or occasion to work through the technical literature on a controversial topic, and so, they must rely on professionals for a disinterested evaluation" [15].

Although cointegration *is* prevalent in academia, this does not mean that it has filtered down to the people who actually work in the financial industry. We are inclined to take the word of authority figures (e.g., teachers and parents) as indisputable for evolutionary reasons. Until very recently life spans have been relatively short, making scientific research in everyday tasks impossible. For instance, rather than conducting a scientific test to discover whether jumping off a cliff would result in injury, many would just take the word of a parent or teacher. Those who did not listen to authority figures often ended up perishing, meaning that only the compliant survived, wiping out the overly inquisitive from our gene pool. Since the early 1960s, experiments [16] have been conducted to show that this level of compliance with authority figures is still ubiquitous.

Reputation

Reputation is everything in finance. Whole markets can crash from loss of confidence in either an asset or a company. When this happens, the institutions involved can lose billions of pounds or even fold completely. Therefore, if the rest of the financial world does not believe that a move to cointegration would be advantageous, then if the change is made, that institution or company could, in the worst case, become bankrupt. Evidently, most institutions would not be willing to take this risk, particularly as correlation seems to be adequate. Even if cointegration did turn out to be better, there may be such a huge loss to begin with that recovery could be impossible.

As I have said already, there are very few companies who would be willing to take this risk, mainly due to a sort of peer pressure. It has been shown by Solomon Asch [17] that many of our choices are influenced by what other people think. This seems obvious to most of us, but perhaps we would not believe this to be the case in the financial world. However, because what other people think matters so much to prices, and therefore business, what other people believe is very important. If there is too much doubt about cointegration, when a company decides to switch to using it, it is likely that those with their shares will try to sell them off as quickly as they can, in which case the company could suffer a loss from which it is unable to recover.

Implementation

Let us imagine that correlation has been conclusively proven to be wrong, and the problems with cointegration have been resolved so that it is now considered to be correct. Even so, there are still problems.

On a practical matter, the swapping from using correlation to cointegration would use a huge amount of resources. Correlation is the basis for most of MPT, and the foundation of finance. Thus, to start using cointegration in the financial world would require everyone who works in finance to learn how to use cointegration in the place of correlation. Correlation is a relatively simple idea to grasp and to use, however, cointegration is not that straightforward. To create the mathematics (e.g., models) needed when using cointegration is much more complex than doing the same for correlation. Indeed, considering what we have said earlier about correlation being taught from an early stage in financial/economic theory, and the fact that cointegration is often not taught until Masters or even PhD level, means that many members of the financial world will have to learn a whole new theory from scratch. This is bound to be very expensive and time-consuming, not to mention the fact that many will be unwilling to learn something difficult and new when they already know a system that works (perhaps not perfectly, but at least well).

Equally, it is not necessarily true that cointegration could ever make up for the cost of implementation. Even if it is a better system than correlation, the increase in profit may not be a particularly large one. Of course, another problem with this is that any mistakes that might arise will be slower to solve than mistakes that arise with correlation. We have had many years to practice using correlation, so in theory, the system is streamlined. This would not be the case with cointegration, as every problem would be "new," i.e., one that has not been observed in the practical world before.

A hybrid model between correlation and cointegration

We have shown that WTI and Brent should perhaps be represented by the dynamics of Model B, for physical reasons. We have also shown that the induced measured correlation between mean-reverting signals was a function of the driving underlier (in terms of Model B, what we call the driving underlier is dx_t^w in $dx_t^w/x_t^w = \sigma_t^w dW_t$ because dx_t^b depends on it: $dx_t^b = \theta(x_t^w - x_t^b + \mu)dt + \sigma_t^o dW_t^\perp$). We have also shown that the more the driving underlier "steadily goes up" (or down) as opposed to "oscillating around its current value," the higher the induced "deceptive" correlation. We can notice this phenomenon in the market as well. Figure 8 is a graph of the historical values of WTI and Brent and the realized 10-day measured correlation for the returns. We can see that the hypothesis that the driving underlier impacts the induced correlation is verified. Indeed, we can see that the rolling correlation can significantly decrease in periods of relatively stable price movements or in periods where the ratio (the spread) is abnormal. The marked correlation is on the order of 99 percent historically, yet there are periods when the correlation moves to 100 percent (from +1 to 0).

One possible way to find a more realistic value for implied correlation is to find a value of a possible instantaneous correlation and try to integrate this over the time to expiry of the option. Intuitively we have shown that whenever the spread is large, the induced realized correlation should be negative; as time goes forward and the spread is progressively absorbed, the correlation reaches its equilibrium state. This can be modeled through the OU process that we have explored extensively in this document and would be the best choice in terms of reconciling the consequences of a wrong assumption with what happens in reality.

A new model

Due to the physical reality behind the fluctuation in price of commodities and because of speculation, the markets perhaps use a mixture of Models A and B. Model A on its own cannot adequately model the mean reversions required (with respect to the physical reality inherent to the relationship of mean-reverting assets) and Model B on its own cannot highlight the very high historical measured correlation on pairs that are theoretically cointegrated. This really means that perhaps Model A is appropriate most of the time but it becomes totally irrelevant when spreads reach their breaking points and where Model B becomes relevant: we return to this duality between the relevance of the models in the short and long term. So, the induced correlation may be higher than it should be if the model is purely mean reverting in nature. As a result, a mixture model between mean reversion and correlation ought to be temporarily chosen to represent these dynamics until the misconceptions are fully understood in the markets and the traders and structurers have taken the appropriate steps to correct their models. With this in mind, perhaps the following trivial candidate mixture model would be more realistic despite its calibration difficulties:

Figure 8: Brent and WTI time series between 2003 and 2011 and their 10-day rolling correlation.

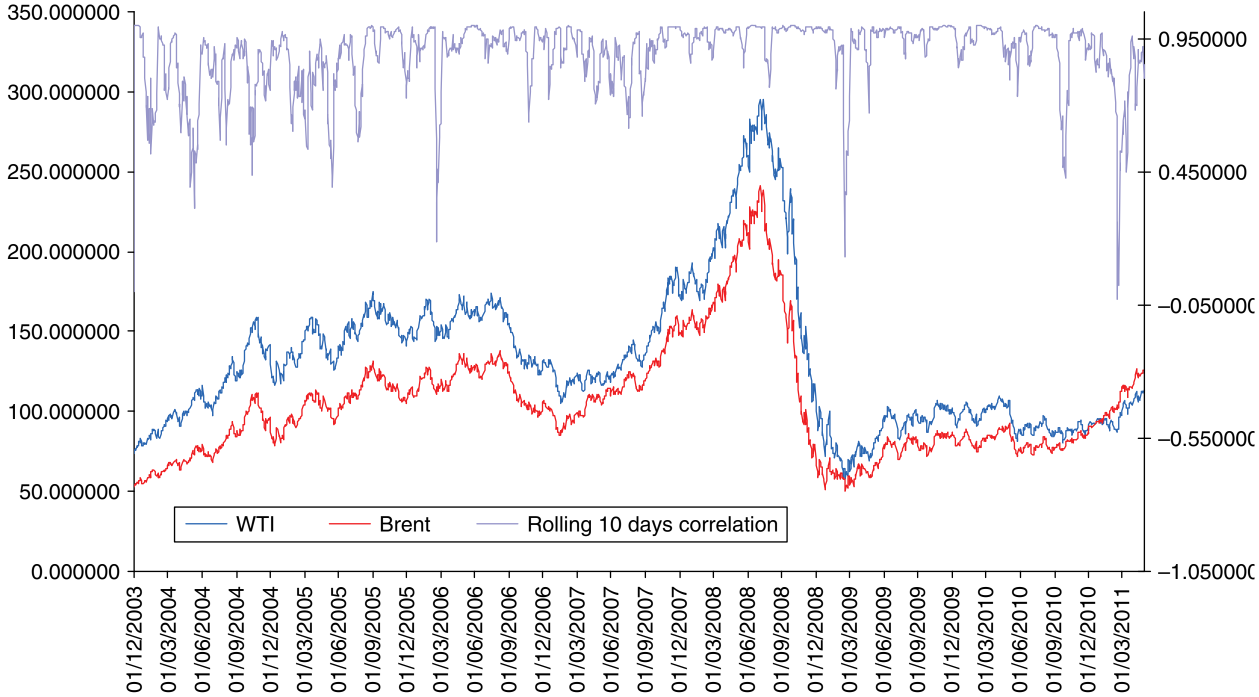

$$dx_t^w/x_t^w = \sigma_t^w dW_t$$
$$dx_t^b = \theta\left[\mu - \left(x_t^w - x_t^b\right)\right]dt + x_t^b\sigma_t^b\left(\rho dW_t + \sqrt{(1-\rho^2)}dW_t^{\perp}\right)$$

In this model note that the driving dynamic given by x_t^w follows the log-normal model but x_t^b follows a mixture between a cointegrated model, the OU, and a classic correlation model. We propose to call this stochastic process "cointelation" (for COINTEgration mixed with correLATION) because the dynamic is still dominated by the mean-reversion model in the long run (cointegration comes first) and because phonetically cointelation sounds better than correlgration, corretion, cointegralation, or correlointegration. The idea of the model is that when $\mu - \left(x_t^w - x_t^b\right)$ is close to 0, that is when the spread between these two assets is at its average value, then the correlation model takes the leading role in the dynamic of x_t^b's move. In a situation where the spread is big, the deterministic part of the SDE takes the ascendant, otherwise some sort of weighted average of these two models takes part in determining the move. Note that the model does not prevent x_t^b going negative. Including $\sqrt{x_t^b}$ in front of the stochastic part of the SDE, similarly to the CIR model, would not prevent x_t^b from becoming negative if there is a big drop in x_t^w. Another incentive to keep it simple, in the current method, is that doing otherwise would alter the dynamic in Model A, critical in correctly capturing the dynamics of our measured correlation in the short run. There are obviously ways to enforce this non-negativity constraint,[5] but we will address these possible models in another article and try to focus for now on this simplest version. We will assume that the movement of the driving dynamic x_t^w is such that x_t^b stays positive. To keep the naming consistent with the usual financial jargon, we will call ρ the correlation of the cointelation. Note that because the drift part of the SDE in the cointelation equation penalizes the measured correlation of these cointelated pairs (in non-extreme cases), the correlation of the cointelation, in these cases, must necessarily be bigger than the measured correlation (this assumption is not necessarily true depending on very specific paths of the driving dynamic). Empirically, a rough estimate for the correlation of the cointelation can be the average of the measured historical correlation with the maximum 10-day rolling historical correlation. Typically, in the case of representing the cointelation for our Brent–WTI example, a correlation of the cointelation around 98 percent when the measured historical correlation is around 93 percent is a good approximation. The other parameters can be estimated sequentially by rearranging the dx_t^b and assuming first that $E\left[x_t^b\sigma_t^b\left(\rho dW_t + \sqrt{(1-\rho^2)}dW_t^{\perp}\right)\right] = 0$. The only challenging parameter estimation is perhaps θ, infamous for being very slow to converge. We suggest a conditional sampling variance reduction technique [18] to address this particular practical issue. The idea of conditional sampling comes from noticing that in our dynamic, the estimation of θ is made difficult in instances where $\mu - \left(x_t^w - x_t^b\right)$ is close or equal to 0. Conditional sampling only samples θ in regions away from this noisy central part. However, by doing this conditional sampling we do get less noisy data but we also lose sample points. Therefore, there exists an optimal way to define these barriers that we will explore in a following article. In the meantime, one simple way to define these regions is via a couple of barriers which, for "normal" financial market data, can be approximated by $b_+ = \frac{\hat{\mu}+\mu_{\max}}{2}$ and $b_- = \frac{\hat{\mu}+\mu_{\min}}{2}$, where $\mu_{\max}$ is the maximum recorded historical spread and $\mu_{\min}$ the minimum recorded historical spread. θ will be sampled when the current spread is bigger than b_+ or smaller than b_-. Note that WTI is a higher-graded type of crude compared with Brent because it is both "lighter" and "sweeter," the main reason that WTI was historically traded at a premium compared with Brent. However, looking at Figure 8, we can see that the historical spread between WTI and Brent became negative early in 2011, which could be an argument against the new model. This negative spread appeared for a mixture of physical reasons, involving, for example, pipelines which created an oversupply of WTI and an undersupply of Brent: so mathematically there was a jump in the μ parameter (something we assumed was constant for the sake of introducing the model in an intuitive way). Note that in the limits, if μ follows the same kind of dynamic as x_t^w, then we are back to Model A. However, these changes of μ for cointegrated pairs are rare enough for the new model to beat, in these specific constant μ intervals, Model A in the long run or Model B in the short run while doing equally well as Model A in the short run and Model B in the long run. As is unfortunately often the case, the assumptions behind the model should not be forgotten by practitioners, the regime change in μ being one of many examples.

Conclusion

In this technical report we have shown that model selection can have very important consequences: both in terms of evaluating risk and abiding by the rules of the challenging new regulatory environment. More specifically, we have shown that the traditional correlation model ought to be substituted with a mean-reversion model when appropriate. This is due to the fact that the induced measured correlation in a mean-reverting dynamic is misleading with respect to its dictionary meaning. This is particularly important because it could impact whether the marketing frameworks employed by regulated firms are in line with the FSA's rules of proper market conduct. What is interesting in the mean-reverting process is that in instances where, by chance, the driving dynamic goes up very much on average or down very much on average, the induced measured correlation for the lagging process returns increases drastically. This is because, in a mean-reverting process, the second process always tries to "catch up" to a process that is by chance significantly going up (or down, as opposed to hesitatingly switching direction). In these situations one accidentally gets a high measured correlation for returns. In the second equation of Model B, the stochastic side of the differential equation never gets the opportunity to lower the correlation by adding randomness in the returns. Given that one might expect this high number because one suspects a strong relationship and understands a "strong relationship" in terms of high correlation in Model A, it becomes very easy to disregard a model that uses mean reversion in such a situation. This is because of how the mind is genetically hardwired. Indeed, because the benefice of seeing a true pattern far outweighs the cost of being mistaken for seeing a false pattern, the mind is naturally set up to see patterns fast at the cost of being sometimes mistaken. Thus the mistake in model selection becomes very apparent in situations where the spreads are stretched significantly more than their historical mean and where the market is hesitating between bullish or bearish. Within the framework of the job of a market maker then, using correlation would make sense for the sake of being consistent with the market prices only because everybody is wrong. But given that traders use a bigger margin for longer expiries because they "feel" the model cannot really model the risk efficiently, why not come clean and admit the model is wrong to begin with? This is unconstructive for the future of the financial industry. This is particularly so if the underlying hypothesis in model selection is conveniently chosen to be that of the correlation model when a more realistic model should be a mean-reversion model, as we have illustrated with the commodities and equities markets.

Acknowledgments

The views expressed in this article are the authors' personal views, and unless specifically stated are not those of Credit Suisse.

Appendix

When we think of correlation we are tempted to say that if two assets are strongly positively correlated then they should, on average, travel together in the same direction. However, this is not necessarily true in a purely mathematical sense. Wilmott gives a brilliant counterexample in his last book. Figure A1 represents an intuitive example of two processes for which the returns are perfectly negatively correlated. Figures A2, A3, and A4 on the other hand represent counterintuitive representations of two perfectly correlated time series returns, two perfectly negatively correlated time series returns and, finally, two time series for which the returns have zero correlations, respectively. Correlation is a mathematical formula appropriate to the most "infinitesimal" level, and perhaps not relevant if we are unable to generalize it at the VaR significance level.

The systems driving the dynamics of Figures A1–A4 can be summarized by equations (E1)–(E4), respectively.

$$\begin{aligned} dx_t^r &= 2_{t \text{ is odd}} - 1_{t \text{ is even}}, \quad x_1^r = 0 \\ dx_t^b &= -2_{t \text{ is odd}} + 1_{t \text{ is even}}, \quad x_1^b = 0 \end{aligned} \tag{E1}$$

Figure A1: An example of intuitive negative correlation described by equation (E1).

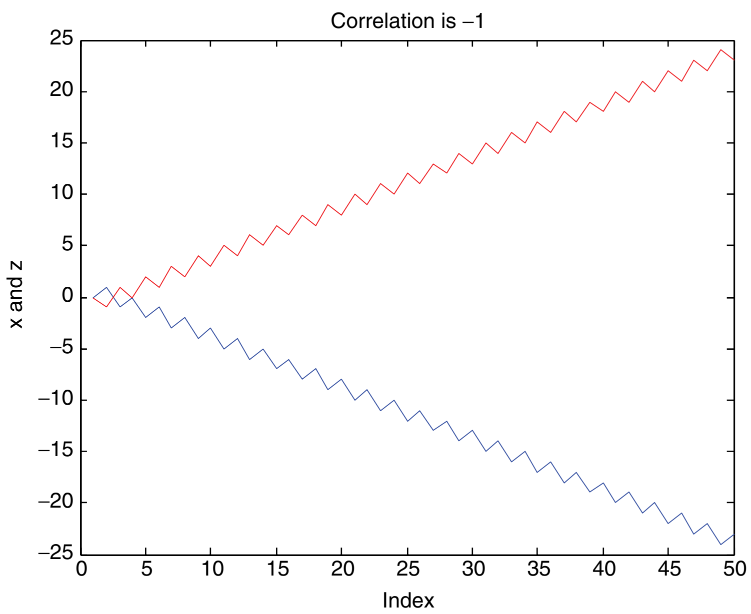

$$\begin{aligned} dx_t^r &= -1_{t \text{ is odd}} + 2_{t \text{ is even}}, \quad x_1^r = 0 \\ dx_t^b &= -2_{t \text{ is odd}} + 1_{t \text{ is even}}, \quad x_1^b = 0 \end{aligned} \tag{E2}$$

$$\begin{aligned} dx_t^r &= -1_{t \text{ is odd}} + 2_{t \text{ is even}}, \quad x_1^r = 0 \\ dx_t^b &= +2_{t \text{ is odd}} - 1_{t \text{ is even}}, \quad x_1^b = 0 \end{aligned} \tag{E3}$$

Figure A2: An example of unintuitive positive correlation described by equation (E2).

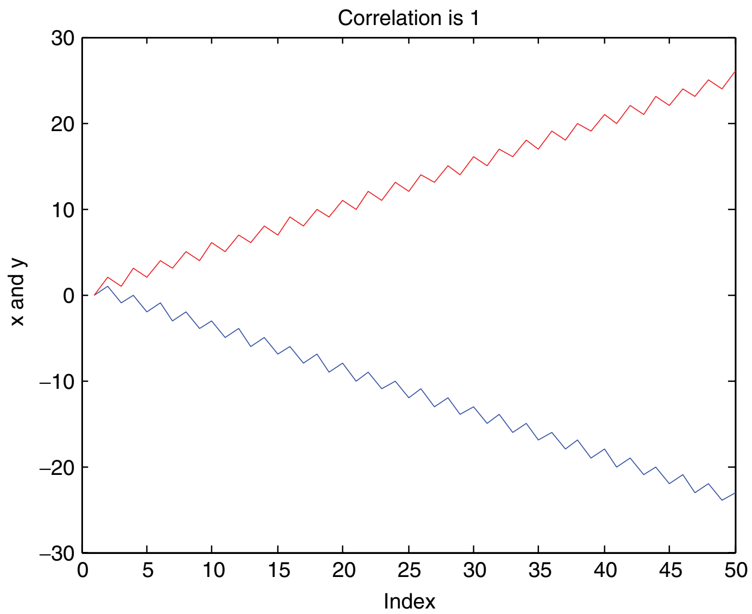

Figure A3: An example of unintuitive negative correlation described by equation (E3).

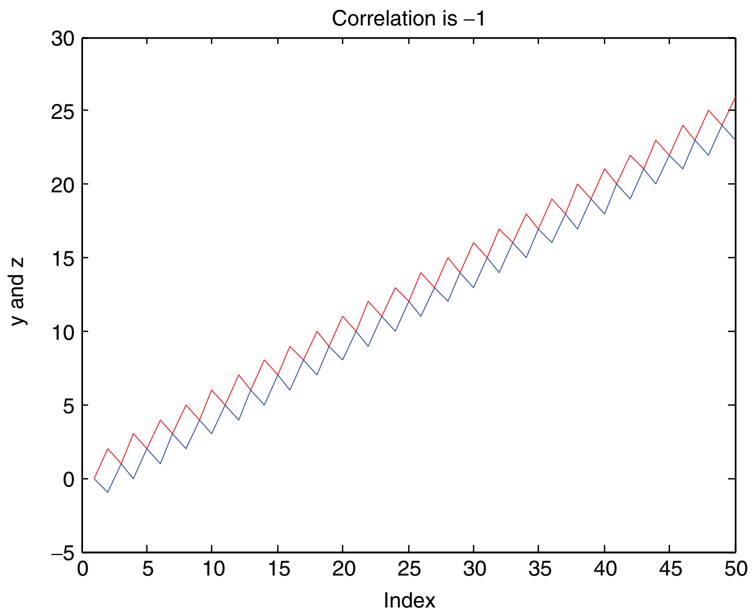

Figure A4: An example of unintuitive zero correlation described by equation (E4).

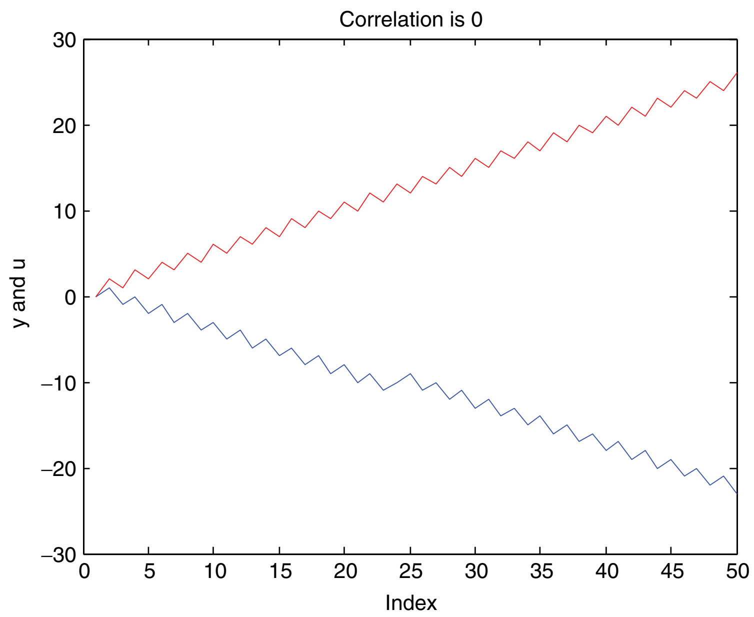

$$\begin{aligned} dx_t^r &= 2_{t \text{ is odd}} - 1_{t \text{ is even}}, & x_1^r &= 0, & t &< T/2 \\ dx_t^b &= -2_{t \text{ is odd}} + 1_{t \text{ is even}}, & x_1^b &= 0, & t &< T/2 \\ dx_t^r &= -1_{t \text{ is odd}} + 2_{t \text{ is even}}, & x_1^r &= 0, & t &\geq T/2 \\ dx_t^b &= -2_{t \text{ is odd}} + 1_{t \text{ is even}}, & x_1^b &= 0, & t &\geq T/2 \end{aligned} \tag{E4}$$

Babak Mahdavi Damghani is an FSA-approved person who has worked at Socgen, Cantab Capital Partners, Citigroup, and Credit Suisse in, respectively, Equities Derivatives High Frequency Statistical Arbitrage, FX Quantitative Strategies, Exotics Derivatives Commodities Trading, Hybrid Structuring, and Cross Asset Quantitative Risk Methodology. He has post-graduate degrees from Oxford, École Polytechnique, and Cambridge in various mathematical and computational sciences

Dani Welch is a graduate in philosophy from the University of Cambridge. She is currently working as part of the Risk Methodology team at Credit Suisse.

Ciaran O'Malley is a graduate in philosophy, politics, and economics from Durham University. He has been working as a cross-asset Junior Quantitative Analyst in Risk Methodology at Credit Suisse.

Stephen Knights has a DPhil in econometrics from the University of Oxford. He has been working as a senior cross-asset Quantitative Analyst in Risk Methodology at Credit Suisse.

ENDNOTES

1. Granger, C. W. J., 1981. "Some properties of time series data and their use in econometric model specification," *Journal of Econometrics*, **16**(1), pp. 121–130.
2. The order of integration $I(p)$ is defined as the number of times a series must be first-differenced for it to become stationary, where stationarity is defined as $I(0)$. Hence, an $I(1)$ series must be first-differenced once to become stationary, $I(0)$. First-differencing is defined as finding the change from one data point in a series to the next, therefore in times series X, the first difference is $\Delta X_t = X_t - X_{t-1}$.
3. $X \sim I(p)$ is defined as the variable X integrated of order p.
4. The linear combination is frequently defined as the residuals of an ordinary least squares regression of X_1 on X_2.
5. For example, $dx_t^b = 1_{x_t^b > 0}\theta\left[\mu - (x_t^w - x_t^b)\right]dt + x_t^b\sigma_t^b\left(\rho dW_t + \sqrt{(1-\rho^2)}dW_t^\perp\right)$.

REFERENCES

Rodgers, J. L. and Nicewander, A. W. (1988) "Thirteen ways to look at the correlation coefficient," *The American Statistician*, **42**(1), pp. 56–66.

Pearson, K. (1895) *Royal Society Proceedings*, **58**, p. 241.

Pearson, K. (1920) "Notes on the history of correlation," *Biometrika*, **13**, pp. 25–45.

Anscombe, F. J. (1973) "Graphs in statistical analysis," *The American Statistician*, **27**(1), pp. 17–21.

Granger, C. W. J. and Newbold, P. 1974. "Spurious regressions in econometrics," *Journal of Econometrics*, **2**(2), pp. 111–120.

According to the RePEc database. http://ideas.repec.org/top/top.item.nbcites.html. Accessed November 3, 2011.

Engle, R. F. and Granger, C. W. J. 1987. "Co-integration and error correction: Representation, estimation, and testing," *Econometrica*, **55**(2), pp. 251–276.

Johansen, S. 1988. "Statistical analysis of cointegration vectors," *Journal of Economic Dynamics and Control*, **12**(2&3), pp. 231–254.

Johansen, S. 1991. "Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models," *Econometrica*, **59**(6), pp. 1551–1580.

Hjalmarsson, E. and Osterholm, P. 2007. "A residual-based cointegration test for near unit root variables," International Finance Discussion Papers 907, Board of Governors of the Federal Reserve System (USA).

Cheung, Y-W. and Lai, K. S. 1993. "Finite-sample sizes of Johansen's likelihood ratio tests for cointegration," *Oxford Bulletin of Economics and Statistics*, **55**(3), pp. 313–328.

Gonzales, J. 1994. "Five alternative methods of estimating long-run equilibrium relationships," *Journal of Econometrics*, pp. 203–233.

Wilmott, P. 2009. *Frequently Asked Questions in Quantitative Finance*, 2nd edn. New York: John Wiley; pp. 336–337.

Russell, B. 1952. "Is there a God?"

Herrnstein, R. J. 1973. *I.Q. in the Meritocracy*. Boston: Atlantic Monthly Press.

Milgram, S. 1974. *Obedience to Authority: An Experimental View*. London: Tavistock Publications.

Asch, S. E. 1955. "Opinions and social pressure," *Scientific American*, 193.

Mahdavi Damghani, B. and Kos, A. 2012. "Dearbitraging with a weak smile," *Wilmott*, pp. 16–18.