Volatility Exposure of CTA Programs and Other Hedge Fund Strategies

Marc Malek & Sergei Dobrovolsky, CFA
Conquest Capital Group LLC
August 2006

The authors of this paper are principals and employees of Conquest Capital Group LLC. No fund or investment program is offered hereby and this article is not intended to and does not constitute investment advice.

SHREK: *Example? Okay, um, ogres are like onions.*
DONKEY: *[Sniffs] They stink?*
SHREK: *Yes. No!*
DONKEY: *They make you cry?*
SHREK: *No!*
DONKEY: *You leave them out in the sun, they get all brown, start sprouting' little white hairs.*
SHREK: *No! Layers! Onions have layers! Ogres have layers! Onions have layers.*
You get it? We both have layers. [Sighs]
DONKEY: *Oh, you both have layers. Oh. [Sniffs] You know, not everybody likes onions. Cake! Everybody loves cakes! Cakes have layers.*
SHREK: *I don't care... what everyone likes. Ogres are not like cakes.*
DONKEY: *You know what else everybody likes? Parfaits. Have you ever met a person, you say, "Let's get some parfait," they say, "Hell no, I don't like no parfait"? Parfaits are delicious.*
SHREK: *No! You dense, irritating, miniature beast of burden! Ogres are like onions! End of story. Bye-bye. See ya later.*

Shrek *(©2001 DreamWorks SKG)*

1. Introduction

While it is generally believed that CTAs have a long volatility exposure, in our testing we see evidence that it is not quite true. Trend-following CTAs benefit from orderly directional trends and suffer losses in random directionless ranges. These characteristics of price behavior are not directly related to volatility levels or volatility changes.

In this article we study the dependence of CTA performance on underlying market volatility, both quantitatively and conceptually. It is a complex issue, depending in each particular case on the time scale of both the volatility and the trading system. Moreover, if one is to compare CTA performance to volatility, one immediately confronts the issue of choosing an appropriate market to compare. Since diverse markets have uncorrelated volatility profiles, one cannot expect performance of a trading system in one market to be related to the volatility of another one. Therefore, since CTAs trade diversified portfolios of systems and markets, any effects of volatility dependence or exposure exhibited by each individual system-market combination are smoothed out and the correlation of their performance to the volatility of any market is very low.

The notion of volatility exposure is sometimes confused with a dependence on volatility levels. If a CTA makes money in periods of high volatility and loses in periods of low volatility, its performance depends on the level of volatility. Volatility exposure, on the other hand, means that the CTA makes money when the volatility rises and loses money when the volatility falls. Both effects have comparable strength, are directly related and thus are quite easy to confuse. By the same token, they need to be studied together, as is done in this article.

The notion that CTA strategies are long volatility came from academic research and became widespread among traders. The first important research on the subject appeared almost ten years ago, when Fung & Hsieh showed that a trend-following CTA behaves similarly to a long lookback straddle. Due to this similarity, CTAs were called a "long volatility" strategy. That research was followed by articles of several authors, expanding the volatility exposure concept to other hedge fund styles. Trading strategies that benefit from strong market moves, like CTAs do, were thus classified as being "long volatility", while arbitrage programs that capture market mispricing were considered "short volatility".

We value and respect all the work performed on the subject and the insight and knowledge that was accumulated on various hedge fund performance issues. Yet, the notion of "volatility exposure" itself appeared as a short convenient name for the idea that CTAs act like long straddles. Used purely as a name, it is probably as good as any other expression. But when taken at its face value to mean that CTAs really have high volatility dependence or that CTAs can profit purely from volatility changes like an option strategy does, is very misleading.

To clear the confusion that has created the very notion of CTA volatility exposure, let us apply options terminology to analyze behavior of a typical CTA trend-follower in a narrow trading range. At first, the market is in a range and the CTA has no positions. Its directional market exposure – its **delta –** is zero. In order to make money, the CTA needs a breakout in either direction. Once the market starts moving, the CTA starts adding to its positions. When the market moves up, the CTA increases its long position (its delta becomes positive and increases). When the market goes down, the CTA increases its short position (its delta becomes negative and decreases). This behavior is exactly what makes the CTA similar to a long straddle. And yet, as any options trader realizes, this kind of sensitivity of delta to price has nothing to do with volatility. It is a textbook definition of being **long gamma**.

Now consider our CTA's volatility sensitivity at the time when the CTA has no positions and is waiting for a trend. Imagine that the historical volatility increases – the absolute values of market returns increase – but the market stays in a range, now widening. Or imagine that the implied volatility increases – the market stays exactly where it was, but the average volatility projections by all participants increase. Obviously in both such cases the CTA doesn't make any money and stays out of the market. The increasing range width actually causes losses in false breakouts. There is no reason to think that the CTA's vega is any different from zero.

If one wants to use options terminology, then strategies that are considered to be long volatility (CTAs and other directional traders) should be correctly termed as being **long gamma.** Strategies that are considered to be short volatility (arbitrage strategies) are actually **short gamma**.

2. Volatility Exposure as a General Concept

As any word that is used daily in various situations, the term "volatility" has come to mean different things. While they all imply increasing range or speed of price changes, the details are often unrelated and contradictory.

You can easily hear one trader speaking of being profitable due to increasing volatility, while another one is blaming increasing volatility for his significant losses. The reason for this discrepancy, of course, is that by "volatility" they mean totally different market behaviors. The first means that the market broke out of a narrow range and started a strong trend. The second means that the market broke from a quiet trend, which was edging up little by little, into a series of sharp large reversals. Or that he lost money in multiple false breakouts in an expanding price range. Both are intuitively correct. By volatility both meant strong market movements – either a strong trend or a sharp reversal. Would it be meaningful to ask them to conclude together what volatility <u>on average</u> meant to them both? That would be pretty much like asking how smart the creatures that walk on two legs are on average, and then trying to figure out whether a human and a chicken are on average smarter than two bears.

Unfortunately, this is exactly the situation in direct testing for volatility influence on trading performance – lumping together strong positive and strong negative effects and getting a weak misleading average as a result. Therefore, from a generic standpoint, the concept of volatility influence on performance is quite unintuitive and contradictory. Whether you believe that your strategy is long or short volatility, whichever way you define volatility, you can find a counter-example refuting your hypothesis.

To discuss volatility exposure in more detail, one has to consider specific definitions of volatility. We do it in Section 8 below, showing that volatility, no matter how defined, is mostly irrelevant to markets' trending behavior and thus to CTA trading profitability. Even in the case of price volatility defined on a time scale comparable to that of the trend, the relevance of such volatility definition is a trivial consequence of the definition itself and brings no additional insight or information.

Yet, if the volatility exposure concept is so confusing, how did it become so popular? The reason is that CTA strategies behave very similar to straddles, both making money in strong market moves regardless of market direction. And since straddles are long volatility, it has naturally led to the conclusion that CTAs are long volatility. Therefore, before we proceed any further, let us go to the root of the issue and consider in detail whether option strategies' being long volatility really implies the same for CTAs.

3. Volatility Exposure: Option Strategies Revisited

To understand the origin of the idea that CTAs are long volatility, let us go back to the origin of the notion of volatility exposure in the field of option trading.

As is well known, the premium of an option is sensitive to several variable parameters, most important of which are the price and the volatility of its underlying. Sensitivities are determined by calculating partial derivatives of the premium with respect to each parameter. Such derivatives, or the Greeks, are then used to estimate the changes in the premium caused by small changes in the respective parameters. Thus, delta and vega are directional and volatility exposures of an option. Gamma is the second partial derivative of its premium or the first derivative of its delta with respect to the underlying price. Adding and subtracting the Greeks of individual options, one can easily calculate exposures of complex option strategies.

If one is to expand the exposure notion to CTAs, the logic often goes like this. Everyone knows that straddles have high vegas and low deltas and thus are volatility exposure strategies. If we know that straddles make money due to being long volatility and can prove that CTA strategies make money in the same environments as straddles, CTAs must also be long volatility. And it's easy to see that both straddles and CTAs lose money when markets are quiet and make money when markets make strong moves irrespective of their direction. Instead of straddles, one can pick any other option strategy that is long volatility and has unlimited profit potential in both upside and downside moves.

The logic above is a little stretched as it is based on one important implicit assumption. It does follow that CTAs are long volatility – but only if straddles not only are long volatility, but also their being long volatility is the very reason for their making money in strong market moves. And that's not only not obvious, despite their having zero deltas and high vegas, but is actually not true, as we will show below.

Let us consider recent cross rates of the Canadian dollar vs. the Japanese yen. From 4/29/05 to 12/7/2005 this market was in a clear uptrend and after that until at least 7/10/06 was in an orderly trading range. Imagine that on

4/29/05 someone wrote us a straddle maturing on 7/10/06. Let us calculate its daily premiums from inception to expiration. We use the Black-Scholes model with a risk-free rate equal to 5% (the fact that we are considering an FX market instead of a stock or a futures contract is irrelevant to the general point of our discussion). We set the straddle strike price at 89.0 so as to have its initial delta equal to zero. The straddle premium can be calculated precisely on each trading day by using the true realized volatility from that day to the expiration date as its implied volatility input (we didn't price the straddle on the last 10 trading days to avoid small sample issues.)

Thus, on April 26 we held a zero-delta straddle with the vega of 73 and the premium of 7.40. Seemingly a pure volatility play. By December 9, the volatility of the underlying changed by 0.6%. Considering our initial vega, which was also the highest that our straddle ever had until its expiration, that meant that by December 9 our premium should have increased to 7.78. But in reality it increased to 17.96! If we really had a pure volatility play, the annualized volatility would have to rise from 10% to 25%.

To see that our position's volatility exposure was by far not our biggest concern, look at its gamma, equal to 8.6. The weekly volatility of our underlying at that point was 1.4%. Multiplying it by the price of 83.3, we can calculate that a 2-sigma weekly change in the price of the underlying was equal to 2.3 points and could cause the delta to change from zero to ±20. In other words, as is usual for straddles, we had a high **gamma** exposure.

If we look at the weekly data presented below (a week being defined as 5 business days), we can see that as the market started making a directional move, our once zero-delta straddle quickly changed its behavior. The call, now deeply in-the-money, started acting like a long position in the underlying and the put, now deeply out-of-the-money, became worthless. On 12/9/05 the delta was equal to 98 and the vega was only 4. Our "pure volatility play" straddle had practically no volatility exposure. And we see that all its gain was due to an increasing delta in a strong market trend. It had nothing to do with volatility.

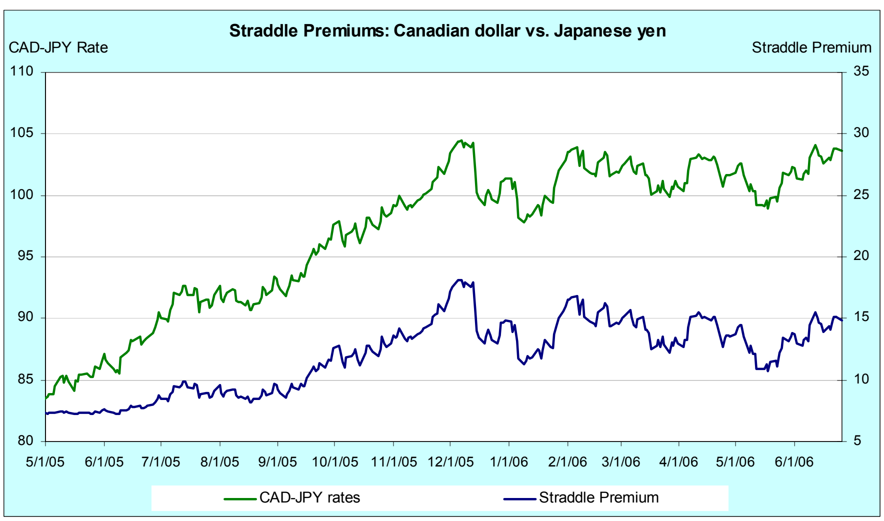

Strike price	89.0		Risk-free rate	5%				Maturity	7/10/06
Date	**Underlying**	**Years to maturity**	**Volatility to maturity**	**Straddle Premium**	**Delta**	**Gamma**	**Vega**	**Intrinsic Value**
29-Apr-05	83.27	1.20	10.1%	7.40	0	8.6	73	5.68
13-May-05	84.83	1.16	10.1%	7.32	12	8.6	72	4.12
27-May-05	86.07	1.12	10.1%	7.44	21	8.3	70	2.88
10-Jun-05	86.87	1.08	10.2%	7.53	27	8.2	68	2.08
24-Jun-05	88.48	1.04	10.2%	7.89	39	7.6	64	0.47
08-Jul-05	92.05	1.01	10.2%	9.54	62	5.8	50	3.10
22-Jul-05	91.34	0.97	10.1%	8.86	58	6.4	52	2.39
05-Aug-05	92.13	0.93	10.1%	9.15	63	6.0	48	3.18
19-Aug-05	91.15	0.89	10.2%	8.42	56	6.8	51	2.20
02-Sep-05	92.34	0.86	10.2%	8.96	63	6.1	46	3.39
16-Sep-05	94.38	0.82	10.3%	10.19	74	4.8	36	5.43
30-Sep-05	97.62	0.78	10.3%	12.61	86	2.9	23	8.67
14-Oct-05	96.17	0.74	10.2%	11.15	82	3.8	26	7.22
28-Oct-05	98.24	0.70	10.1%	12.73	90	2.6	17	9.29
11-Nov-05	99.07	0.66	10.3%	13.32	92	2.2	15	10.12
25-Nov-05	102.30	0.63	10.5%	16.20	96	1.0	7	13.35
09-Dec-05	104.29	0.59	10.7%	17.96	98	0.6	4	15.34
23-Dec-05	99.65	0.55	10.1%	13.26	94	1.7	9	10.70
06-Jan-06	98.20	0.51	10.1%	11.72	92	2.4	12	9.25
20-Jan-06	100.00	0.47	10.0%	13.22	96	1.3	6	11.05
03-Feb-06	103.74	0.44	10.0%	16.72	99	0.3	1	14.79
17-Feb-06	102.70	0.40	9.8%	15.51	99	0.4	1	13.75
03-Mar-06	102.63	0.35	9.7%	15.24	100	0.3	1	13.68
17-Mar-06	100.06	0.31	9.9%	12.52	98	0.7	2	11.11
31-Mar-06	100.71	0.28	9.6%	13.00	99	0.4	1	11.76
14-Apr-06	103.03	0.24	9.8%	15.14	100	0.1	0	14.08
28-Apr-06	101.63	0.20	9.8%	13.57	100	0.1	0	12.68
12-May-06	99.18	0.16	9.8%	10.95	100	0.2	0	10.23
26-May-06	101.88	0.12	9.5%	13.47	100	0.0	0	12.93
09-Jun-06	103.03	0.09	10.4%	14.46	100	0.0	0	14.08
23-Jun-06	103.82	0.05	10.0%	15.08	100	0.0	0	14.87

But maybe it is just semantics? Trend or not, can we still somehow attribute our gain to the straddle being long volatility? Maybe say that the rising volatility was the main reason for the trend itself? To prove that it is not so, let us modify our price data. We multiply the actual daily market returns by a linearly decreasing factor and recalculate our prices based on the new returns. The factor makes the volatility steadily go down from 10.1% to 1.1%. If we think that our being long volatility is important, and volatility goes down ten times, shouldn't the straddle premium crash down? As we see below, actually it steadily goes up, not caring about the volatility at all. It actually goes up higher, as we have to increase the initial returns to keep the overall volatility the same and thus the trend goes higher than in the original data.

On the chart below, which shows the new adjusted straddle premiums and underlying rates versus the original data, one can see how similar the premiums are to their respective underlying prices. Aside of the first several days, the lines are pretty much the same – an obvious consequence of our straddle having an almost 100-delta exposure.

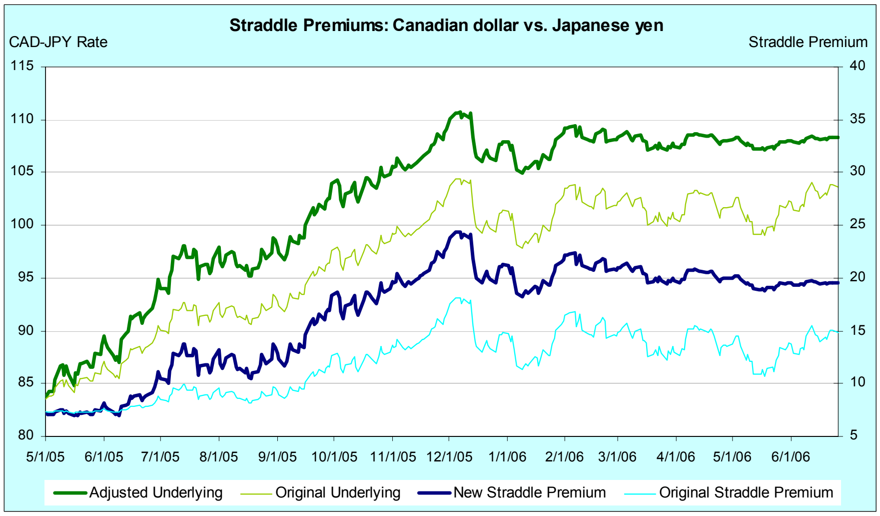

Strike price	89.0			Risk-free rate		5.0%			Maturity	7/10/06
Date	New Underlying	Return Adj. Factor	Old Underlying	Years to Maturity	Volatility to Maturity	Straddle Premium	Delta	Gamma	Vega	Intrinsic Value
29-Apr-05	83.27	1.69	83.27	1.20	10.1%	7.40	0	8.6	73	5.68
13-May-05	85.89	1.64	84.83	1.16	9.6%	7.18	22	8.6	71	3.06
27-May-05	87.91	1.58	86.07	1.12	9.5%	7.54	38	8.0	66	1.04
10-Jun-05	89.17	1.53	86.87	1.08	9.3%	7.82	47	7.6	61	0.22
24-Jun-05	91.68	1.48	88.48	1.04	9.0%	8.89	65	6.1	48	2.73
08-Jul-05	97.09	1.43	92.05	1.01	8.8%	12.91	89	2.6	21	8.14
22-Jul-05	96.06	1.38	91.34	0.97	8.2%	11.69	89	2.9	21	7.11
05-Aug-05	97.18	1.33	92.13	0.93	7.9%	12.48	93	2.1	14	8.23
19-Aug-05	95.83	1.28	91.15	0.89	7.8%	11.05	90	2.9	18	6.88
02-Sep-05	97.42	1.22	92.34	0.86	7.6%	12.33	95	1.8	11	8.47
16-Sep-05	99.98	1.17	94.38	0.82	7.4%	14.63	98	0.7	4	11.03
30-Sep-05	103.91	1.12	97.62	0.78	7.2%	18.36	100	0.1	1	14.96
14-Oct-05	102.23	1.07	96.17	0.74	6.7%	16.51	100	0.1	1	13.28
28-Oct-05	104.57	1.02	98.24	0.70	6.3%	18.68	100	0.0	0	15.62
11-Nov-05	105.47	0.97	99.07	0.66	6.2%	19.42	100	0.0	0	16.52
25-Nov-05	108.68	0.92	102.30	0.63	6.3%	22.47	100	0.0	0	19.73
09-Dec-05	110.56	0.86	104.29	0.59	6.2%	24.18	100	0.0	0	21.61
23-Dec-05	106.39	0.81	99.65	0.55	5.3%	19.84	100	0.0	0	17.44
06-Jan-06	105.26	0.76	98.20	0.51	4.9%	18.55	100	0.0	0	16.31
20-Jan-06	106.64	0.71	100.00	0.47	4.5%	19.77	100	0.0	0	17.69
03-Feb-06	109.34	0.66	103.74	0.44	4.2%	22.31	100	0.0	0	20.39
17-Feb-06	108.61	0.61	102.70	0.40	3.8%	21.41	100	0.0	0	19.66
03-Mar-06	108.56	0.56	102.63	0.35	3.4%	21.16	100	0.0	0	19.61
17-Mar-06	107.15	0.51	100.06	0.31	3.3%	19.58	100	0.0	0	18.20
31-Mar-06	107.49	0.45	100.71	0.28	2.8%	19.77	100	0.0	0	18.54
14-Apr-06	108.56	0.40	103.03	0.24	2.6%	20.66	100	0.0	0	19.61
28-Apr-06	108.00	0.35	101.63	0.20	2.3%	19.93	100	0.0	0	19.05
12-May-06	107.20	0.30	99.18	0.16	2.0%	18.96	100	0.0	0	18.25
26-May-06	107.95	0.25	101.88	0.12	1.6%	19.54	100	0.0	0	19.00
09-Jun-06	108.18	0.20	103.03	0.09	1.6%	19.61	100	0.0	0	19.23
23-Jun-06	108.32	0.15	103.82	0.05	1.1%	19.58	100	0.0	0	19.37

But what if this effect is entirely due to the specifics of our definition of volatility? Well, we use the same return volatility definition as the very option pricing models that gave rise to the idea of "volatility exposure" itself. Anyway, we can also look at other definitions of volatility. Let us also consider price volatility, defined as the standard deviation of prices. It can be also defined as the width of a price range (the highest minus the lowest price) with very similar results. This definition embodies the notion that a volatile market is the one that makes wide moves. We agree that a market making a strong move in a trend shows a relatively high volatility in this sense, but disagree that its volatility must be higher than that of any trading range – a high volatility in this sense can also be attained if the range is simply very wide.

To illustrate this point, let us replace the trending part of our data with a range whose volatility is increasing and whose width and standard deviation of prices are the same or higher than those of the trend. We split our actual returns into two equal time segments before and after 12/3/05 and replace the first (trend) segment with the second (range) segment. We then multiply each daily return in the first segment by a factor linearly increasing from 0.1 to 7.22. The trend had the standard deviation of prices equal to 5.05 and the prices varied from the highest to the lowest by 19.1 points. Our new expanding range has the same standard deviation of prices equal to 5.05 and the price range of 25.7 points. We also multiply the remaining data by the same factor to keep the volatility to expiration from significantly decreasing.

The first segment of our data became at least as volatile under the price volatility definition as the trend. Moreover, both its return volatility and its price volatility are gradually and very significantly increasing. Let us see if it helped the straddle's performance. To compare the old and the new premiums, we adjust our strike price so that our new straddle still starts at zero delta. We also multiply the premium by 0.14 to start at the same premium level as we had before and to be able to compare the strategy returns, not just absolute levels. As one would expect, the straddle's premium loses its value due to time decay, while the price moves are not significant enough for high positive returns. While the price range is becoming wider and its historical price volatility from inception is steadily increasing, the highest premium of the straddle is only 8.13 (on 01/09/06), before the call loses all its value and the put drops to the levels close to its intrinsic value.

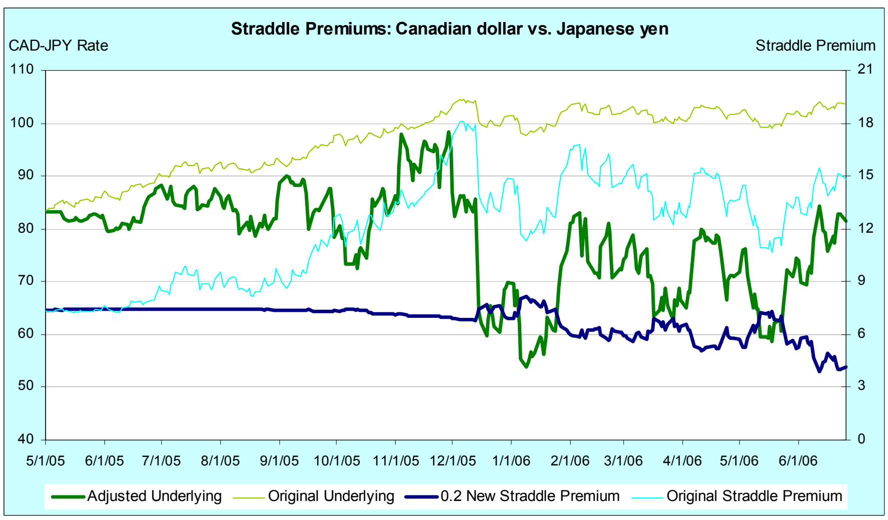

Strike price	110.65						Risk-free rate	5.0%				Maturity	7/10/06
Date	**New Underlying**	**New Returns**	**Return Adj. Factor**	**Years to Maturity**	**Vol. to Maturity**	**Price Vol. from Incep.**	**Call Prem.**	**Put Prem.**	**0.14 × Straddle Premium**	**Delta**	**Gamma**	**Vega**	**0.14 × Intr. value**
29-Apr-05	83.27	0.06%	0.10	1.20	61.2%		15.43	36.38	7.40	0	1	73	3.91
13-May-05	81.40	-0.30%	0.56	1.16	62.2%	0.75	14.45	37.48	7.42	-3	1	70	4.18
27-May-05	82.78	-0.05%	1.02	1.12	63.3%	0.72	15.07	36.92	7.42	-1	1	70	3.98
10-Jun-05	81.00	1.14%	1.48	1.08	64.3%	1.21	14.16	37.98	7.45	-4	1	67	4.24
24-Jun-05	86.12	1.06%	1.94	1.04	65.4%	1.45	16.61	35.52	7.45	3	1	70	3.50
08-Jul-05	84.47	-0.80%	2.40	1.01	66.4%	2.43	15.65	36.40	7.43	0	1	68	3.74
22-Jul-05	84.48	0.67%	2.85	0.97	67.5%	2.48	15.52	36.47	7.42	0	1	66	3.74
05-Aug-05	86.33	0.74%	3.31	0.93	68.8%	2.45	16.39	35.68	7.44	2	1	66	3.47
19-Aug-05	78.51	-2.33%	3.77	0.89	69.8%	2.52	12.43	39.75	7.45	-10	2	59	4.59
02-Sep-05	89.11	0.53%	4.23	0.86	70.8%	2.60	17.42	34.33	7.39	5	1	66	3.08
16-Sep-05	79.99	-8.08%	4.69	0.82	72.3%	2.96	12.85	39.08	7.42	-9	2	57	4.38
30-Sep-05	78.30	-6.42%	5.15	0.78	73.1%	2.93	11.77	39.90	7.38	-12	2	54	4.62
14-Oct-05	76.43	0.67%	5.61	0.74	73.7%	3.67	10.60	40.81	7.34	-17	2	51	4.89
28-Oct-05	82.67	-0.60%	6.07	0.70	74.7%	3.68	13.01	37.19	7.17	-8	2	55	4.00
11-Nov-05	92.12	3.36%	6.53	0.66	75.4%	4.20	17.25	32.17	7.06	5	1	60	2.65
25-Nov-05	92.50	5.17%	6.99	0.63	76.6%	4.84	17.07	31.82	6.98	5	1	58	2.59
09-Dec-05	85.51	2.77%	7.22	0.59	76.9%	4.98	12.97	34.91	6.84	-7	2	52	3.59
23-Dec-05	61.57	-2.75%	7.22	0.55	73.2%	6.57	3.37	49.47	7.55	-55	2	27	7.01
06-Jan-06	55.39	-11.01%	7.22	0.51	72.6%	7.73	1.85	54.32	8.02	-69	2	19	7.89
20-Jan-06	63.15	5.94%	7.22	0.47	72.1%	9.40	2.83	47.75	7.23	-60	2	24	6.78
03-Feb-06	82.31	1.32%	7.22	0.44	72.0%	9.51	7.74	33.70	5.92	-27	2	41	4.05
17-Feb-06	76.54	7.98%	7.22	0.40	70.9%	9.40	4.98	36.92	5.99	-42	2	33	4.87
03-Mar-06	76.16	1.90%	7.22	0.35	70.1%	9.29	4.08	36.64	5.82	-48	2	29	4.93
17-Mar-06	63.42	-9.46%	7.22	0.31	71.5%	9.26	1.34	46.84	6.88	-75	2	15	6.75
31-Mar-06	66.46	-3.50%	7.22	0.28	69.6%	9.50	1.26	43.93	6.45	-76	2	14	6.31
14-Apr-06	78.33	0.63%	7.22	0.24	71.0%	9.41	2.85	33.86	5.24	-57	2	22	4.62
28-Apr-06	70.97	-0.14%	7.22	0.20	70.4%	9.33	1.06	39.64	5.81	-78	2	12	5.67
12-May-06	59.51	-8.46%	7.22	0.16	70.7%	9.40	0.13	50.38	7.21	-96	1	2	7.30
26-May-06	72.24	6.26%	7.22	0.12	68.7%	9.71	0.34	38.09	5.49	-90	1	5	5.49
09-Jun-06	78.33	9.09%	7.22	0.09	75.4%	9.60	0.54	32.39	4.70	-85	2	7	4.62
23-Jun-06	82.77	-0.14%	7.22	0.05	72.1%	9.46	0.19	27.82	4.00	-92	1	3	3.98

The alternative definition of volatility does not explain the gains either. For both straddles and CTAs, profits are due to the price moving from some level A to some level B in a reasonably orderly way, not to its chopping randomly around, touching A and B and ending up somewhere in between.

The reason for all the effects that we have discussed above is simple and well-known – option premiums are most sensitive to volatility when the options are at the money. When deeply in-the-money, they become long or short positions in their underlying, when deeply out-of-the-money, they become worthless compared to any non-deeply-OTM option and can be ignored compared to other options in the strategy. Thus, a straddle acts like a textbook straddle only when it is close to being at the money. As the underlying price moves away from its strike level, the straddle's sensitivity to volatility practically disappears. And this argument is true for any volatility strategy. To have its options retain their volatility sensitivity, the strategy needs the market to stay close to their strike levels. If the market starts a trend move, all butterflies, condors and similar strategies get swept away and become simple borrowed or lent cash.

There is also a well-known fact that an option strategy can be long volatility and yet act in exactly the opposite way to what we expect from any CTA (or a straddle) – lose value when its underlying market makes a strong trend move in any direction. An example of such a strategy is a time (calendar, horizontal) spread.

We write a short-term call and buy a long-term call, both of which have the same strike. The strategy is obviously long volatility – we are long the call with a later expiration date. It is also a pure volatility play with zero intrinsic value at any time. And yet the time spread loses its value from any market move away from its current level, as both options start losing their time premiums when going either in or out of the money. In a strong market trend, this long-volatility position will quickly become worthless.

And, vice versa, a short time spread makes money in any strong trend that is profitable for a CTA, while being a pure short-volatility play.

As an illustration, the following shows the premium of a time spread for the CAD-JPY market. For simplicity, we assumed a zero interest rate. We can see how quickly our long-volatility position with a very high positive vega loses all its value in our perfect trend.

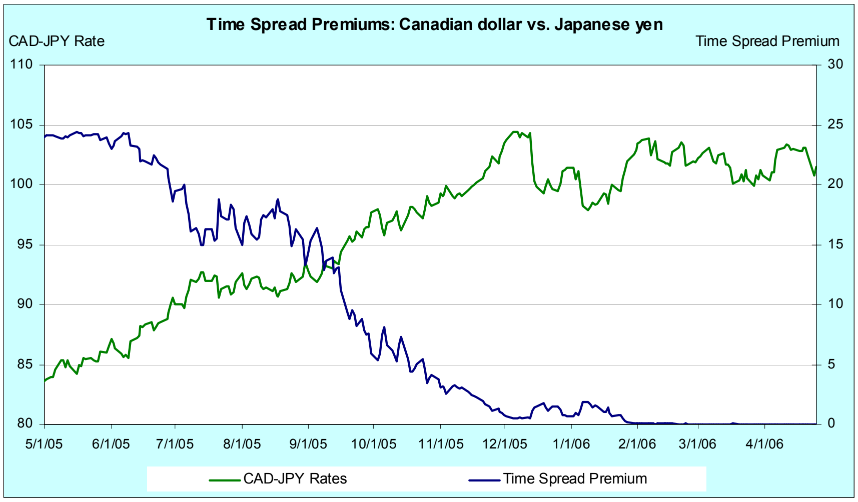

Strike price	83.27		Risk-free rate	0%			Maturity S		5/10/06		Maturity L	7/10/06
Date	Underlying	Returns	Years to Maturity S	Maturity L	Volatility S	Volatility L	Call S	Call L	100 × Time Spread	Delta	Gamma	Vega
29-Apr-05	83.27	-1.76%	1.03	1.20	10.2%	10.1%	3.45	3.68	23.7	0	-0.3	261
13-May-05	84.83	-0.56%	0.99	1.16	10.1%	10.1%	4.22	4.46	24.1	0	-0.3	272
27-May-05	86.07	0.97%	0.95	1.12	10.2%	10.1%	4.95	5.19	23.7	-1	-0.3	286
10-Jun-05	86.87	1.58%	0.92	1.08	10.3%	10.2%	5.44	5.68	23.3	-1	-0.2	298
24-Jun-05	88.48	0.11%	0.88	1.04	10.3%	10.2%	6.54	6.75	21.7	-1	-0.2	320
08-Jul-05	92.05	0.94%	0.84	1.01	10.3%	10.2%	9.40	9.56	16.1	-2	0.0	342
22-Jul-05	91.34	0.90%	0.80	0.97	10.2%	10.1%	8.72	8.89	17.4	-2	0.0	354
05-Aug-05	92.13	0.45%	0.76	0.93	10.2%	10.1%	9.36	9.52	15.9	-2	0.0	360
19-Aug-05	91.15	0.55%	0.73	0.89	10.3%	10.2%	8.48	8.66	17.8	-2	0.0	366
02-Sep-05	92.34	-0.48%	0.69	0.86	10.4%	10.2%	9.49	9.64	15.3	-2	0.1	367
16-Sep-05	94.38	1.04%	0.65	0.82	10.5%	10.3%	11.34	11.45	11.2	-2	0.2	342
30-Sep-05	97.62	1.24%	0.61	0.78	10.5%	10.3%	14.43	14.49	5.9	-1	0.2	256
14-Oct-05	96.17	-0.54%	0.57	0.74	10.4%	10.2%	12.99	13.07	7.3	-2	0.2	294
28-Oct-05	98.24	-0.23%	0.53	0.70	10.3%	10.1%	15.00	15.04	4.1	-1	0.2	215
11-Nov-05	99.07	-0.19%	0.50	0.66	10.5%	10.3%	15.82	15.85	3.1	-1	0.2	180
25-Nov-05	102.30	0.80%	0.46	0.63	10.8%	10.5%	19.04	19.05	1.2	0	0.1	87
09-Dec-05	104.29	0.38%	0.42	0.59	11.0%	10.7%	21.02	21.03	0.5	0	0.1	46
23-Dec-05	99.65	-0.38%	0.38	0.55	10.4%	10.1%	16.38	16.40	1.5	-1	0.2	109
06-Jan-06	98.20	-1.53%	0.34	0.51	10.3%	10.1%	14.94	14.95	1.9	-1	0.2	135
20-Jan-06	100.00	0.82%	0.31	0.47	10.2%	10.0%	16.73	16.74	0.7	0	0.1	61
03-Feb-06	103.74	0.18%	0.27	0.44	10.2%	10.0%	20.47	20.47	0.1	0	0.0	9
17-Feb-06	102.70	1.11%	0.23	0.40	10.0%	9.8%	19.43	19.43	0.1	0	0.0	7
03-Mar-06	102.63	0.26%	0.19	0.35	9.8%	9.7%	19.36	19.36	0.0	0	0.0	3
17-Mar-06	100.06	-1.31%	0.15	0.31	10.3%	9.9%	16.79	16.79	0.1	0	0.0	8
31-Mar-06	100.71	-0.49%	0.11	0.28	9.8%	9.6%	17.44	17.44	0.0	0	0.0	2
14-Apr-06	103.03	0.09%	0.07	0.24	10.2%	9.8%	19.76	19.76	0.0	0	0.0	0
28-Apr-06	101.63	-0.02%	0.03	0.20	10.4%	9.8%	18.36	18.36	0.0	0	0.0	0

Now let us return to CTAs. We have established that for a straddle to remain sensitive to volatility it has to be at the money, i.e. the underlying price has to stay close to its strike level. But this means that the market should not move, which in turn means that a CTA does nothing, either staying out of the market or seeing no gain or loss on its positions as the price stays the same. For a CTA to enter new positions and to start showing profits or losses, the market has to make a move. But in this case, the straddle's volatility sensitivity becomes negligible compared to its directional exposure. Due to its high initial gamma, the straddle stops acting like a typical ATM straddle, loses its volatility exposure and starts behaving like a simple long or short position.

To summarize our point, one cannot claim that CTAs are long volatility by showing that they act like straddles. CTAs only act like those straddles whose volatility exposure became negligible compared to their directional exposure. Trend-following CTAs make their profits in trends – orderly directional price moves away from equilibrium. Whether return volatility is increasing or decreasing is not important. Considering option strategies, it is quite clear that volatility exposure and exposure to strong market moves are totally different things, as one can easily create a strategy that is long volatility and at the same time completely loses its value in any directional moves, and vice versa.

The feature that is common to CTAs and straddles is increasing their position deltas in moves up and decreasing in moves down. CTAs do it by buying in uptrends and selling in downtrends, straddles simply have their put or call deltas increase as they get away from their ATM levels. In options this behavior is due to their having a long gamma exposure. If one insists on applying option terminology to CTAs, one has to conclude that CTAs are long gamma.

4. Average Long Term Volatility

Before proceeding further, let us make a side note regarding CTA performance independence from long-term market volatility levels.

While in more volatile markets all trends and ranges become wider, it does not matter for trading profits or losses of a typical CTA trading FX or futures. All the CTA has to do is size its positions inversely proportionately to the market volatility. One can trade natural gas and eurodollars in the same portfolio without ever worrying about the fact that a daily percentage change of the former can be orders of magnitude higher than that of the latter. The reasons for this are very low margin requirements in futures and FX trading, giving one the freedom to base position sizing purely on risk management considerations.

Stock traders do seem to care about the overall level of stock volatility, separating stocks into blue chips and risky stocks. Yet an excessive risk of any financial instrument that offers high enough returns can be ameliorated very easily. Imagine being offered an investment whose possible loss over some time period is 100%, while you can only stand a 1% loss. If you like the returns, the solution is to invest only 1% of your money and keep the remaining 99% in cash. The only thing important in a risky stock is not its regular long-term volatility, but its high vulnerability to catastrophic events and often a precarious competitive market position of its underlying company.

5. Volatility Level Dependence vs. Volatility Exposure

Let us clarify one more issue regarding the notion of volatility exposure. One needs to take care not to confuse volatility exposure with dependence on the level of volatility.

Used consistently with the terminology of options and of trading in general, exposure to some parameter means sensitivity to its changes. Being long an instrument means that we benefit from its becoming more expensive and lose due to its becoming cheaper. A positive vega of an option strategy means that the option premium goes up when the volatility increases and goes down when it decreases. To calculate a strategy's volatility exposure, one thus has to compare the strategy's P&L with changes in volatility, not with its levels.

Yet volatility exposure can be easily confused with dependence on the level of volatility. A statement that traders perform better when the volatility is high has nothing to do with their volatility exposure. For example, VIX is the most popular S&P volatility indicator. If one wants to use it to check whether a strategy is long volatility, it seems very tempting to check the correlation of the strategy returns to VIX itself. The correct approach, of course, is to calculate the correlation of the returns to changes in VIX. Another example is the ability to hedge one's volatility exposure. One can't hedge a trader's dependence on high volatility levels by going short a long-volatility instrument.

Therefore, in addition to showing that CTAs are not long volatility we also show that they are not very sensitive to volatility levels.

6. Exposures of Other Hedge Fund Strategies

Let us now generalize our observations to other hedge fund strategies, also commonly viewed as being long or short volatility.

Hedge funds differentiate themselves from buy-and-hold strategies by being directionally-neutral to their underlying markets, attempting to profit from downside moves as well as from upside ones. Yet, their strategies are based on two profoundly opposite approaches.

Consider a market whose price level is, according to some measure, in a state of equilibrium. That is, it is correctly priced with respect to all available information, with respect to expected major events, or relative to similar markets. Then at some moment the price level changes and the initial equilibrium thus gets disturbed. The likely outcomes of this event can be classified into two scenarios. First, the price returns to the equilibrium. Second, the disturbance increases further, and the price continues to move away from the equilibrium. Multiplying the likelihoods of the disturbance disappearing or increasing by the amounts of profits or losses to be made in each case, a hedge fund decides whether its strategy is to fade or to follow such movements away from the equilibrium, whichever direction it views as more profitable in the long run.

Since in everyday language volatility often means disturbance, strategies that capitalize on expectations that small disturbances will grow are called "long volatility". The ones that expect everything to return back to normal are called "short volatility". For example, a trend-following CTA views a breakout from a stable range as a sign that the move away from the range will continue further. An arbitrage fund, on the contrary, views a widening spread between two related instruments as a sign that the spread must narrow back. As we showed above, volatility actually has nothing to do with it. In options, it would be called gamma. In many other fields even this analogy breaks down. It is simply what it is – positive or negative exposure to moves away from the current state of equilibrium. As many researches have already suggested, it is most similar to buying or selling insurance on the underlying equilibrium.

Let us consider insurance itself as an example. An insurance buyer pays a fixed price and benefits in case of an unusual detrimental event (a move away from equilibrium), while an insurance seller is paid a fixed price and benefits from stable (equilibrium) conditions. Imagine, for example, a small-town taxi company owner who pays upfront to buy a one-year car accident insurance on his large park of cars from a local bank. Both sides estimate the chances of car accidents and after some negotiations sign the contract. The situation is obviously reminiscent of a hedge fund bet. If they were hedge funds, the bank would be said to be "short volatility" and the owner "long volatility". Yet not only "market volatility" does not matter here, it practically does not exist. Each accident is a random event depending on a myriad of external influences. There is not enough data to calculate changes of the car accident probability in a small town on a time scale of several months, not to mention its volatility. Moreover, its volatility does not matter for the ultimate payout, as both sides have purely directional bets. Even if they decide to break the contract and figure out how the fair premium has changed, it will be the accident probability itself that will matter, not its incalculable volatility.

The reason that volatility, contrary to the situation with options, is not anywhere in the picture is that options give their buyer a chance, and the chance is priced based on volatility. Yet, as anyone who has ever calculated the Greeks for a hedged option position knows, the underlying instrument always has a 100 delta and all its other Greeks equal exactly zero. A directional position is a directional position, it does not have other exposures.

To clarify volatility versus directional exposure of hedge fund strategies, let us break a decision-making process by a hedge fund into consecutive stages. At Stage One, the fund makes a general decision about its strategy. At the risk of oversimplification, it decides whether in case of a market move it will follow or fade it. That gives it a kind of mental positive (follow) or negative (fade) "gamma exposure". At this stage, before having taken any positions, its volatility exposure is zero. If the volatility increases or decreases, the fund will simply decrease or increase the size of its future position. At Stage Two, the market starts to move, and the fund takes a position. Its exposure at this point is best understood by making the future decisions from two standpoints: one regarding the position already taken and the other one ignoring the current position. From the standpoint of the current position, the trader is 100-delta directional. The volatility can cause him to increase or decrease the size, but the P&L is purely directional. From the latter standpoint, the trader is back to Stage One, having a mental gamma exposure and waiting to follow or fade another move.

Thus, our logic applies to any hedge fund bet. Hedge fund positions at any point are purely directional, most probably with volatility-dependent sizing, but no direct volatility exposure. Hedge fund strategies considered as following or fading market moves are best described as, respectively, being **long or short gamma**. Or, in other words, they either, correspondingly, benefit from **price moves away from equilibrium** or from **stable equilibrium prices**.

These changes in the exposure definitions don't in any way negate the most important distinction of long and short gamma managers – the profound difference in their return profiles which is not reflected in common risk measures like Sharpe ratio. These risk measures systematically underestimate the risk of short gamma programs and overestimate the risk of long gamma programs. Managers that are short gamma collect small stable returns when markets return to equilibrium and risk large losses during unusually large moves. They can show years of positive returns whose volatility does not reflect the true risk of the strategy, until one extraordinary event causes a tremendous loss. Long gamma managers, on the contrary, by the nature of their strategy pay small costs of exposure to the market during choppy directionless periods, while expecting large profits during strong directional moves. Barring unreasonable drastic actions by an investment manager, long gamma strategies that suffer demise normally do so not from a shock, like short gamma strategies, but from a large prolonged drawdown, where ordinary expected losses mount up without any gains to counter them.

7. Hedging Implications: "Long Volatility" or "Long Volatilities"?

There is one more point that needs to be made and that is not taken care of by realizing that a "volatility" exposure is actually an exposure to gamma or to movements away from equilibrium.

The most useful consequence of knowing that your investments have long or short exposure to something is the ability to hedge one exposure by the other. Every once in a while it is suggested that one can actually hedge a "volatility" exposure by combining long and short exposure strategies. Sort of like combining several options into one zero-delta or zero-vega strategy and saying that we have, respectively, no total directional or volatility exposure. Unfortunately, for hedge fund strategies this is not true. While it is definitely beneficial to combine long and short gamma strategies in one portfolio as they usually have low to negative correlation to each other, it is not correct to speak of such strategies being able to hedge each other.

The reason for this is that combining long and short gamma hedge fund strategies is pretty much like combining options based on different and often totally unrelated underlying instruments. Moreover, unlike options, finding long and short exposure strategies based on the same underlying instrument is practically impossible.

Consider two investment instruments A and B. The investment A is long gamma and the investment B is short gamma. Can we hedge our total gamma ("volatility") exposure by investing in both of them? Most likely no. As an obvious example, imagine that the investment A is a commodity trend-following trader that specializes in trading soybeans (a long gamma strategy), while the investment B is a mortgage arbitrage fund (a short gamma strategy). Can we say that we are hedged in our total gamma ("volatility") exposure? For a soybean strategy, our being long gamma means that we risk losing money when the soybeans are in a random non-trending mode. For a mortgage fund, our risk is a significant destabilizing event in the mortgage market. Our two exposures don't hedge each other in any way whatsoever.

More interestingly, in special cases long and short gamma strategies can actually have the same risks. Consider a mutual fund (or a mostly long equity hedge fund) as a long-gamma Strategy A and a fixed income arbitrage fund as a short-gamma Strategy B. The biggest risk for both of them is a major equity market crash, obviously devastating for Strategy A and causing bond margin expansion and thus major losses for Strategy B.

Again, the above does not in any way detract from the obvious truth that it is beneficial to combine long and short gamma strategies in one portfolio, because their correlations are most likely close to zero. But in choosing such strategies, one has to check individual strategy characteristics and correlations. Their gamma ("volatility") exposure is not enough and is not really important.

8. Volatility Definitions vs. Price Behavior

Trend-following CTA strategies perform very well in orderly trending markets and lose money in choppy random ones. Such properties of a market are usually obvious from looking at a price chart.

In this Section, we consider conceptually different definitions of volatility and show that they don't define the price behavior in a way that influences CTA profitability. Same levels of volatility can accompany both range and trend behavior. Moreover, changing the volatility of a range doesn't turn it into a trend, and vice versa. This Section is applicable both to dependence on volatility levels and to volatility exposure, showing that neither volatility, nor its changes make a difference in the market behavior important for profitability of CTAs.

Now let us consider possible volatility definitions. Using various parameters of prices and their changes, one can devise a large number of expressions that somehow express the generic idea of the strength/speed/range of market moves and thus be considered a definition of volatility. Yet, conceptually there are **three categories** of volatility definitions which are fundamentally different.

The definition that is used in Black-Scholes-type option pricing models is **return volatility**, calculated as the standard deviation of returns, which in turn are defined as natural logarithms of consecutive price ratios. The other popular definition is that of **price volatility**. It can be calculated in a variety of ways, the most straightforward being (i) the standard deviation of prices; (ii) the highest high minus the lowest low of prices and (iii) the highest minus the lowest closing prices. The third definition that is sometimes used is **daily range volatility**, calculated either as daily ranges or as average true ranges (i.e. daily ranges including previous closes).

While various modifications of these three definition categories can be created, one still measures one of the same three things about the same market. It is hard to imagine any important effect being observed due to defining the same fundamental thing in technically different ways. For example, if a price range of the same market over the same period increases due to one definition and decreased due to another one from the same category, the amounts of such increase and such decrease are probably not significant. Most probably, in such case one can say that the range stayed pretty much the same. If some effect is observed under one definition but not under another one from the same category, such effect is most likely not robust enough to be used anyway.

Now let us consider actual price data. To intuitively demonstrate that volatility is not necessarily beneficial (or detrimental) for supposedly long-volatility CTAs (or any other strategies that are based on price action), we show how volatility can easily be higher in a range than in a trend. For this purpose, we take an obviously very profitable recent rally in the copper and compare its volatility to that of the trading range that followed. To avoid any effects of longer or shorter time periods, we chose the data from Feb. 8 to Aug. 11, 2006 as a period which can be divided into two equal halves representing, respectively, a trend and a range. The range itself is quite difficult for trading – it gets narrower with unexpected spikes and large up and down days.

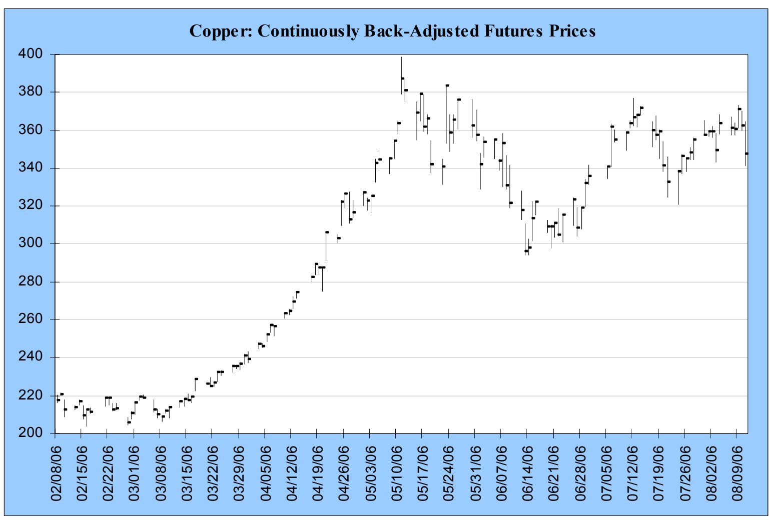

a) Return Volatility

Both in option pricing and in other theoretical finance models volatility is usually defined as the **standard deviation of returns**. If markets were random, their daily returns would be normally distributed and their average and standard deviation would allow one to calculate the probability of observing any particular return in the future. Real markets are leptokurtic with large moves happening much more often.

In the case of copper, the **trend return volatility equals 35%** and the **range return volatility 57%**.

This definition of volatility is actually unrelated or even inversely related to the notion of orderly market behavior. In fact, trends often have stable returns with low volatility, while choppy market conditions with gaps have high volatility. Therefore, CTA strategies actually often prefer periods with low return volatility.

One reason why this definition is irrelevant to the markets' trending behavior and thus to CTA performance is that it doesn't change if one adds some fixed value to each daily return. It shows only variations of returns from the average and does not care what that average is, quite contrary to the very root of the trend-range trading philosophy. So, for example, in the chart below we subtracted 1% from each daily return, thus turning the trend into a range and the range into a trend. Yet the volatility remained identically the same for any time period on the chart!

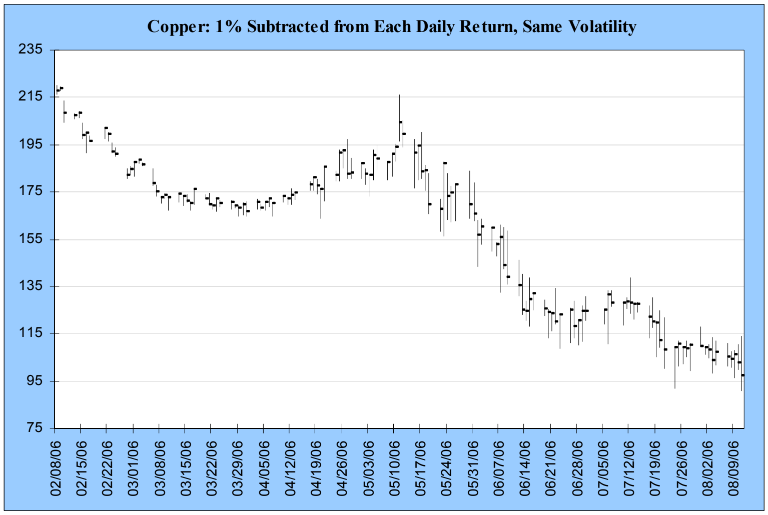

b) Daily Range Volatility

Another possible definition of volatility is the **average daily range**, i.e. the high minus the low of the same daily bar, possibly averaged over some time period.

The fact that for our copper chart this volatility is higher in the range than in the trend is obvious from the chart. The **trend volatility equals 5.7 points** and the **range volatility 13.7 points**.

Daily range volatility definitions based on intraday ranges look even less useful for a CTA trend-follower than do return volatilities and there is no reason to suspect CTAs to benefit from this kind of volatility whatsoever. Average trade lengths for practically all CTAs are much longer than one day. Therefore, wider daily ranges do nothing but increase risks and intraday drawdowns in holding positions, requiring wider stops. Higher highs and lower lows without changes in closing prices mean more false breakouts. Obviously, intraday ranges fail to capture daily trending action completely.

Moreover, daily range volatility is just a special case of price volatility for intraday data scales. Therefore, we limit our discussion to price and return volatilities.

c) Price Volatility

The most common way to define price volatility is to calculate the **standard deviation of closing prices**. This definition captures the notion of volatility as a wide-range price movement. It is used, for example, as volatility in Bollinger bands.

While a trend normally travels further simply due to its nature, a range can be just as wide. In the figure below, we multiplied the returns of the copper price range to make it have the same standard deviation of prices as the trend:

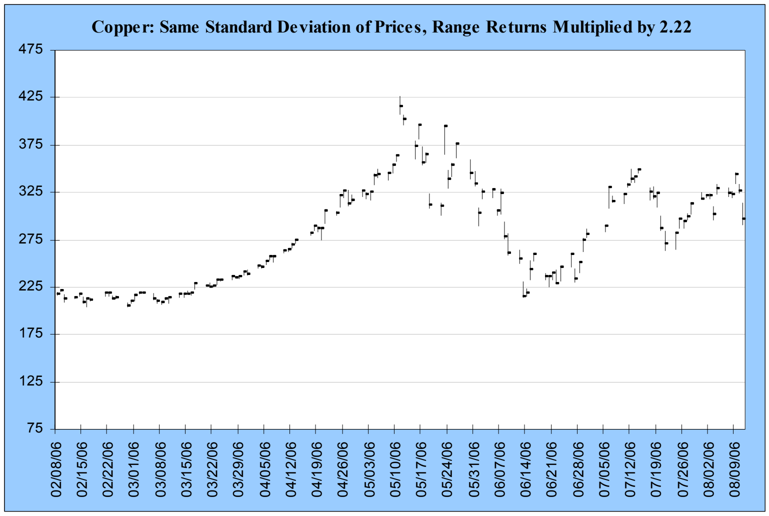

This kind of volatility does not reflect the order in which the prices occur in a sample, so if we randomly mix our prices in time, making them look totally unlike real prices, the volatility will remain the same. Also, the chart above clearly shows again that it is the way the prices move that makes the trend. One can compare the orderly development of the trend in the first half of the chart with the choppy slide down of the first leg of the range. While the range of both moves is the same, the move down is very unstable and a trader getting short has to set much wider stop-loss levels and thus will be substantially less profitable for the same amount of risk than a long trader in the trend.

The difference of this volatility definition from the other two is that it **does imply a higher likelihood of trends**. If returns do not change dramatically, an average trend will definitely have a wider price range than an average trading range. For example, to get the same price "volatility" on the chart above, we had to increase the return volatility in the range by 2.4 times, which is unusual and gave the chart above its unnatural look.

Yet, unlike return volatility, the usefulness of this definition as a measure **separate** from directional price exposure is quite questionable. Returns and their volatility are a measure of, so to speak, the speed of the market. That's why their behavior is often so different from that of prices themselves. Price volatility, on the other hand, is simply the range of prices themselves. Price volatility is most likely to increase in two cases. First, in a trend, when the price volatility pretty much repeats the price move itself, the changes in price volatility being the same as market returns. In this case CTAs make money and their returns obviously correlate with price moves and thus with price volatility changes. Second, price volatility can increase in an expanding range. In this case CTAs mostly lose money due to false breakouts. One can definitely lump these two cases together and call them periods of high price volatility, but the practical usefulness of such exercise is quite questionable.

Anyway, as will be shown below, testing shows that the trend cases of high price volatility win over the expanding range cases. If one is careful about the time scale of the system and of the volatility, then a trend-following strategy performance is correlated both to the absolute values and changes of price volatility. Yet, for a whole CTA such effects are averaged away by using different markets and time scales.

9. Market Testing

In this Section we test dependencies of CTA performance on volatility levels and CTA volatility exposures directly.

The best way to test such dependencies is to calculate correlations. Although one can also compare, for example, average performances under various market conditions, such testing is most often inconclusive as its statistical accuracy is low. Correlations, on the other hand, are very accurate, especially when we use over 4000 data points in our 15 years of daily data. They are especially useful in checking that two data streams are unrelated. While a high correlation can be due to spurious effects and does not prove any dependence without further study, a correlation close to zero is practically always a definite sign that there is no influence of one data stream on the other. Exceptions to this rule imply some very unusual dependencies and are not relevant to our studies.

As we discussed above, due to some confusion in understanding of what volatility exposure actually is, we need to test for two different phenomena:

1) The first phenomenon is dependence of CTA performance on ***volatility levels***. For example, higher levels of price volatility are under some conditions more beneficial to CTA performance, as they imply wider ranges of price movements and are more likely to be due to a trend than to a trading range.

2) The second phenomenon is a CTA volatility exposure per se, i.e. its ability to generate positive or negative returns due to ***volatility changes***.

As previously discussed, we test CTA dependence on price and return volatilities. We define the price volatility as the standard deviation of prices and the return volatility as the standard deviation of returns.

To represent CTA strategies, we use the following data series:

1) S&P Managed Futures Index ("S&P MFI"), constructed and managed by Standard & Poor's. Consists of 12-15 managed futures funds which are open to new investors. The funds are selected and equally weighted once a year. Daily data are available from Jan. 2003. Monthly pro forma data are provided from 1998.

2) Calyon Financial Barclay Index, constructed and managed by Calyon Financial. Currently consists of 14 CTAs which, unlike those in the S&P MFI, represent a wide variety of futures and FX strategies. The funds are equally weighted and rebalanced once a year. Data are available from Jan. 2000.

3) Conquest Managed Futures Select Program ("MFS") is a systematic trend-following program designed and managed by Conquest Capital Group. The program currently consists of 20 systems trading 55 diversified futures and currency markets. All systems are based on N-day breakouts. MFS is 87% correlated to S&P MFI and is a traded benchmark of a typical trend-following CTA strategy. In this testing, we use pro forma MFS results from Jan. 1990 to Aug. 2006 with pro forma slippage and commissions based on actual Conquest Capital Group trading.

4) Individual MFS systems (B5, B6, …, B165, B200), which are the 20 systems that, as described above, comprise the MFS program. As each system is based on N-day breakouts, we denote each of them by the letter "B" with the value of N next to it.

a) VIX Index

We start with calculating correlations of CTA index returns to the most popular volatility measure – the VIX index. Since VIX is an implied monthly volatility of S&P 500 options, we also calculate correlations of returns to S&P 500 monthly volatility directly. The historical monthly return volatility for S&P has an 84% correlation to VIX.

To test volatility dependence, we proceed as follows. To check whether strategy performance depends on the **level of volatility**, we have to select a specific length of the volatility time period. Say, we want to test it for monthly volatility, defining the month as 22 business days. Then we need to calculate one-month rolling returns and calculate their correlations to one-month rolling volatility, the time periods for the returns and volatilities being the same.

To estimate a strategy's **volatility exposure**, we need to test whether the strategy returns depend on changes in volatility levels. This can be done in a variety of ways, but the results should be similar to those under the following two definitions. First, the correlation of the strategy's daily returns to one-day changes in monthly volatility taken on the same day. We compare, for example, the daily CTA return today to the change of the market's monthly volatility between today and yesterday. Second, the correlation of the monthly returns to one-month changes in monthly volatility. We thus compare the monthly return this month to the change in the monthly volatility in this month from the previous month.

The table below shows daily correlations of VIX closing values to daily returns of two CTA indices since their inceptions and to MFS since 1990. The middle two rows show correlations to MFS returns taken since the inception of each of the other indices to avoid any effects due to time period differences. The columns *VIX* and *VIX Changes* show two effects discussed previously – the dependence of performance on VIX itself and on VIX daily changes. In the columns *S&P Monthly Vol.* and *S&P Vol. Changes* we compare daily returns to monthly historical volatility of S&P 500 and its daily changes.

Correlations of VIX to CTA Index Returns

						Monthly CTA Return Correlations			
Index	**Data Since**	**VIX**	**VIX Change**	S&P Monthly Vol.	S&P Vol. Change	Ending Date VIX	Starting Date VIX	S&P Vol.	1 Month S&P Vol. Diff.
Conquest MFS	**1990**	**2.8%**	**25.4%**	-1.9%	-6.0%	13.0%	-2.5%	4.8%	7.8%
Calyon Barclay CTA	**2000**	**3.6%**	**20.1%**	-0.2%	-3.2%	17.9%	1.4%	8.7%	5.4%
S&P Managed Futures	**2003**	**-0.8%**	**8.5%**	-1.5%	-3.7%	3.9%	0.0%	-14.4%	-26.2%
Conquest MFS	***2000***	***5.9%***	***27.5%***	*-0.8%*	*-7.2%*	*24.2%*	*4.5%*	*-0.8%*	*-2.9%*
Conquest MFS	***2003***	***0.4%***	***15.6%***	*-1.2%*	*-5.5%*	*17.9%*	*3.9%*	*-1.2%*	*-9.4%*

Comparison of daily VIX values to daily CTA returns is made problematic by VIX being based on monthly volatility. To make the comparison more meaningful, we thus have to compare VIX values to ***monthly*** CTA returns. This is done in the next column and does show higher correlations than those of daily returns. Since VIX values are calculated daily and are supposed to be related to the future, we can match them with monthly data applying various time shifts. We do it in two opposite ways: in the first column of the rightmost part of the table we compare VIX data to previous month returns, in the second column – to next month returns. As expected, the next month correlations are zero. Even though VIX is based on estimates of *future* volatility, it is not a crystal ball. It is based on historical volatility and is most closely related to historical performance. The third column compares monthly returns with monthly S&P volatility over the same month. The last column compares monthly returns with the difference between this month and last month volatilities.

As can be seen in the table, ***all correlations are too low to imply any significant relationship between CTA returns and VIX volatility***.

The only correlation in the table that is significantly different from zero for all three indices is the correlation between returns and VIX changes. Interestingly, the correlation of returns to S&P volatility changes is zero, even though the two volatility measures differ only by their historical versus implied nature. One possible reason for this effect is the following. On days of strong market moves, when CTAs have large positive or negative returns, put or call premiums (and thus implied volatility levels) can increase more significantly due to an increased likelihood that a strong market trend or a countertrend move will drive prices much further than a random walk model would imply. The fact that this effect is due to extreme moves can be checked directly by considering a scatterplot diagram of MFS returns and VIX changes.

On this diagram, we marked with a different color those observations that exceed the 2 standard deviation level for at least one of the two parameters. While the overall correlation equals 25%, the correlation for the data lying within the 2-sigma box is only 8%, while the correlation for the data lying outside is 42%. The correlation for the data lying outside the 3-sigma box is 56%!

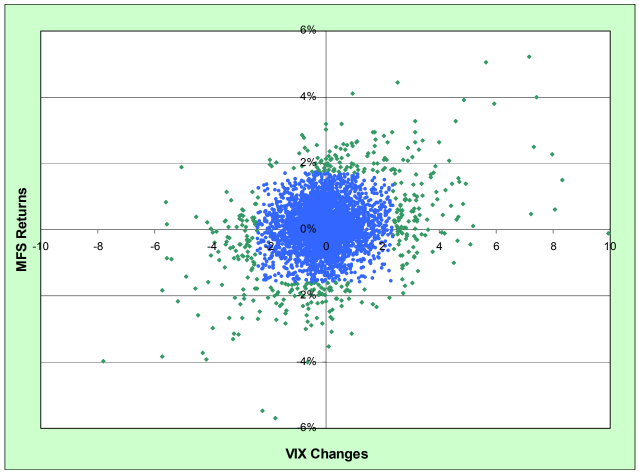

As an implied volatility measure, VIX reflects overall trader perception of the future volatility of financial markets and thus has broader implications than a simple S&P 500 volatility measure would have. Yet, as we saw above, neither its level nor its changes show significant correlations with CTA returns. We tested several other volatility measures, which gave similar results. The other measures are even more market- or sector-specific and have less chance than VIX to influence a diversified market portfolio of a typical CTA.

b) Historical Volatility: Market and Time Scale Dependence

Conquest MFS system, in addition to allowing us to test assumptions based on over 15 years of data, allows us to overcome the biggest obstacle in testing for volatility exposure. As described above, if we do assume that CTAs have an exposure to volatility, we need to specify the market whose volatility we consider. Yet each CTA trades its own portfolio of markets. Moreover, the underlying markets are extremely diverse and the allocations are often not public and can vary in time. Thus, if we were to test the dependence of total returns of a specific CTA on its underlying market volatility, we would have to construct some kind of a weighted average, combining the volatility of its individual markets into one index. The accuracy and even the procedure of such index construction would in any case be quite questionable. The same goes for testing volatility exposures of CTA indices.

The chart below shows 15-year correlations of daily data for rolling monthly volatilities of five markets from different sectors. The table shows both correlations between volatility levels and between daily volatility changes. As one would expect, volatilities of markets from different sectors are not closely related:

Correlations of Return Volatilities: Daily Returns, Jan. 1990 - Jul. 2006										
	Volatility Values					Volatility Changes				
	S&P	Oil	Aluminium	Yen	Eurodollars	S&P	Oil	Aluminium	Yen	Eurodollars
S&P	1	27%	-10%	19%	6%	1	15%	1%	5%	20%
Oil	27%	1	7%	6%	-2%	15%	1	2%	8%	8%
Aluminium	-10%	7%	1	-1%	-8%	1%	2%	1	2%	0%
Yen	19%	6%	-1%	1	-4%	5%	8%	2%	1	5%
Eurodollars	6%	-2%	-8%	-4%	1	20%	8%	0%	5%	1

Correlations of Price Volatilities: Daily Returns, Jan. 1990 - Jul. 2006										
	Volatility Values					Volatility Changes				
	S&P	Oil	Aluminium	Yen	Eurodollars	S&P	Oil	Aluminium	Yen	Eurodollars
S&P	1	17%	-10%	23%	3%	1	7%	3%	0%	12%
Oil	17%	1	31%	-11%	-11%	7%	1	1%	2%	6%
Aluminium	-10%	31%	1	-3%	-9%	3%	1%	1	-2%	5%
Yen	23%	-11%	-3%	1	4%	0%	2%	-2%	1	-1%
Eurodollars	3%	-11%	-9%	4%	1	12%	6%	5%	-1%	1

The only way out of this difficulty is to consider CTA returns from one specific market and compare them to that market's volatility. This information is very difficult to obtain for an individual CTA, impossible for a CTA index, and yet very easy to get from MFS testing.

The second interesting point in volatility testing is dependence on the trading system time scale. If we consider an option position, the time scale of volatility in question is obvious – the time to the option expiration. Both trading systems and volatility measurements have their own time scales. Moreover, different CTAs prefer different time scales, usually ranging from 5 to 200 days. Imagine that we want to test whether a CTA in question is long volatility. First, we have to choose a time scale. If the weekly volatility is falling, the monthly volatility is rising, and the annual volatility is flat, which one are we supposed to be long or short?

The difficulty compared to options is that we have no reason to say that an N-day system has to be influenced only by the N-day volatility. We need to test various volatility-system time scale combinations to understand their interaction and performance dependence.

Thus, in order to meaningfully test volatility exposure, we need to define the question as clearly as possible. First, we select a specific market. Second, we take each of the 20 MFS systems representing 20 time scales and calculate returns on the selected market for each of them separately. Third, we select several volatility time scales and calculate the volatility of the market for each time scale. Finally, we compare the volatility in each of the selected time scales to the performance of each of the 20 MFS systems, all of them representing the same market.

c) Time Scales of Volatility and Systems: Theoretical Considerations

Let us first understand the relationship between the system and volatility time scales conceptually.

Return volatility, as we have already discussed above and as can be shown in testing, is not directly relevant to trend or range price behavior and thus does not matter for CTA trading. Below we focus on price volatility, i.e. price ranges.

Consider an N-day breakout system, which is used in our testing and which is a benchmark trend-following system. It enters the market based on N days of historical data and on average holds its position for approximately 1.7N days. The similarity between the two numbers is not accidental – practically any trend-following system that stays with the trade for a period X makes its decision based on a time scale close to X. The only obvious exception is a system that is made to exit based purely on time due to some outside considerations like correlation to some index. Let us define our volatility on some time scale of the length X and require that our N-day system react to X-day volatility changes. On average, as discussed above, the price volatility expansion has some likelihood to be due to a trend and thus to become a profitable opportunity for our system.

We have determined above that the whole "trade cycle" for a system, from the first data point used to the last day in trade, lasts approximately 2.7N days. Our volatility measurement intervals are of length X. Thus, if we are to require that the full "trade cycle" of the system happen in time X, our system's N has to be lower than (1/2.7)X = 0.37X days. On the other hand, to consider a more balanced case, if the market starts a directional movement today, a breakout system does not need N days to react to it and enter a trade. In the worst case, that of a V-bottom, it will react in N days. In the best case, that of a range breakout, it will react immediately. Let us take N/4 as a crude average time that an N-day system needs to enter a trade. While already in a profitable trend, the system does not need the new market behavior to continue over the whole average trade length to be profitable. Let us, again as a crude average, require that our new profitable period last half the normal trade length, i.e. 0.8N. So, our average reasonable X for which a system responds to volatility changes on scale X should be of order 1.2N.

For N much larger than this number, the system reacting to volatility disturbances on a time scale X either won't have time to enter the trade or time to profit from it. Moreover, short volatility surges of a sufficient magnitude are likely to make a long-term system enter the market and exit it with a loss in the subsequent reversal. Short-term volatility should on average be detrimental to long-term systems, which is confirmed in our testing.

Let us summarize our conclusions. Consider market volatility on a time scale X. Generally, the shorter the time scale of a system, the more it profits from the volatility. Systems with N below 0.4X fall completely inside the volatility frame and should have similar high correlations to volatility levels. Systems with N between 0.4X and 1.2X respond to volatility strongly and positively, their correlation decreasing with N. For systems with N substantially longer that 1.2X the dependence on volatility falls to zero to negative levels.

d) Dependence on Volatility Levels – Results

According to the procedure discussed above, we selected 5 liquid representative markets from different futures/FX sectors: the yen, light crude oil, the S&P 500, eurodollars and aluminum.

For both price and return volatilities, we present results for (i) five volatility time scales for aluminum and (ii) the 3-month volatility time scale for the five markets.

The chart below shows correlations of returns to price volatility for aluminum and illustrates the effects discussed above. Each line on the chart corresponds to a specific volatility time scale. The time scales range from one week to 3 months. The charts show correlations versus the time scales N of MFS systems. As one can see, for each volatility time scale the shortest systems benefit the most. For example, for the 3-month volatility time scale, 0.4X and 1.2X equal 24 and 78 days, respectively. We can see that B24 does in fact mark the threshold of a short-term plateau, while at B82 the correlation falls by a little over a half.

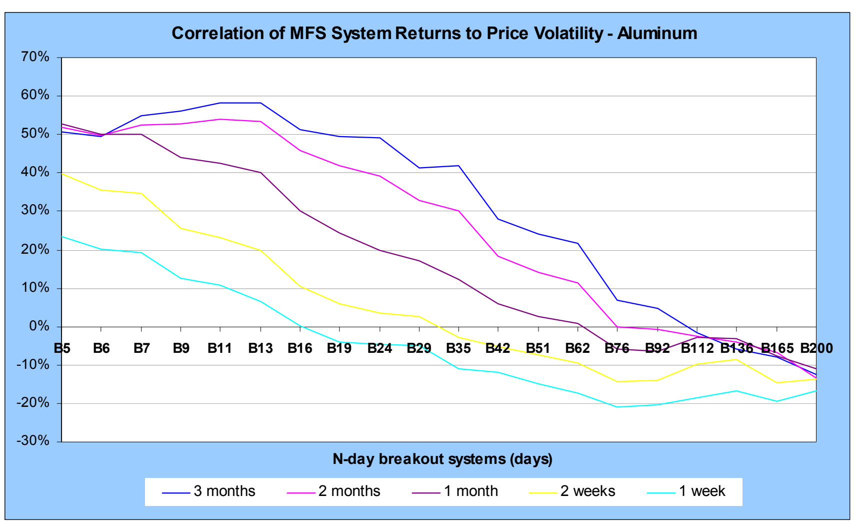

We ran the same tests for the other four markets. For some of them the results are less pronounced, but mostly they have similar falling curve patterns, presented below for the three-month volatility:

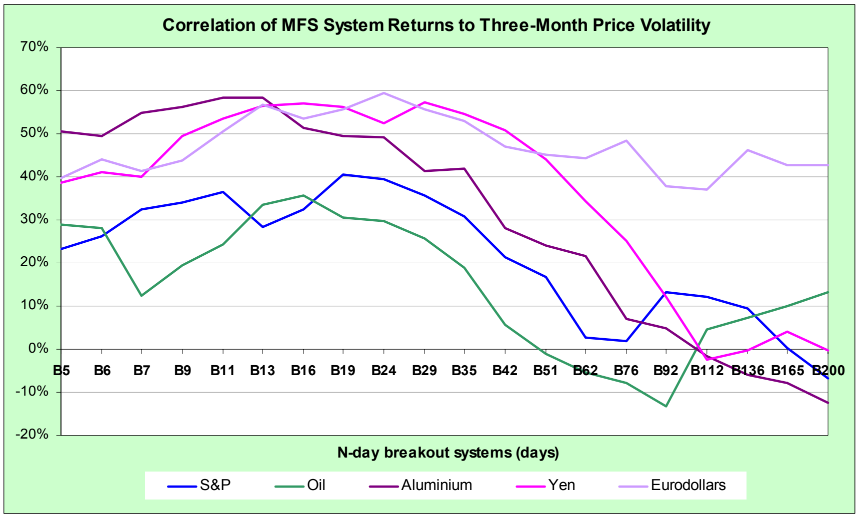

Those were the results for price volatility defined as the standard deviation of prices. Now we look at return volatility defined as the standard deviation of returns.

The charts below are similar to the charts for price volatility, but the correlations here are much lower, more negative and the effects are less pronounced and more dissimilar for different markets. As was discussed above, influence of return volatility on system performance is much more difficult to understand conceptually. The same concept holds true, although – the longer the time period N for the N-day breakout system, the lower the correlation. For long-term systems both short-term price and return volatilities create noise that is detrimental to their performance.

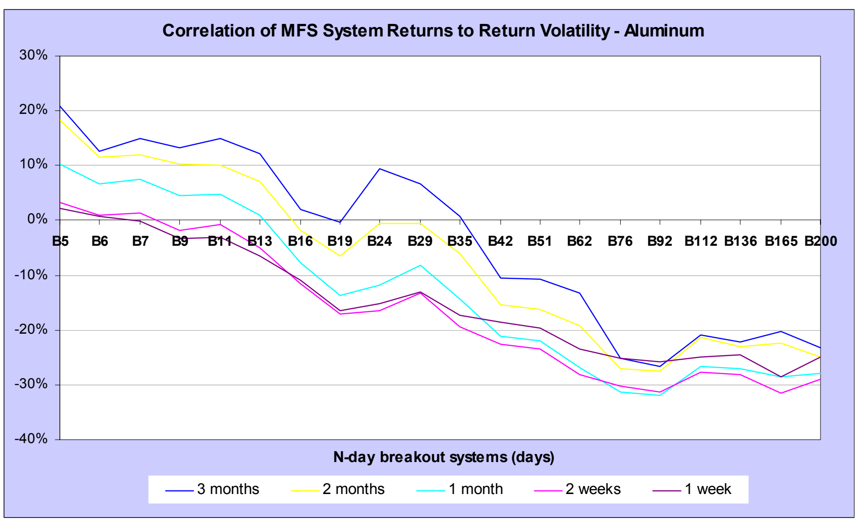

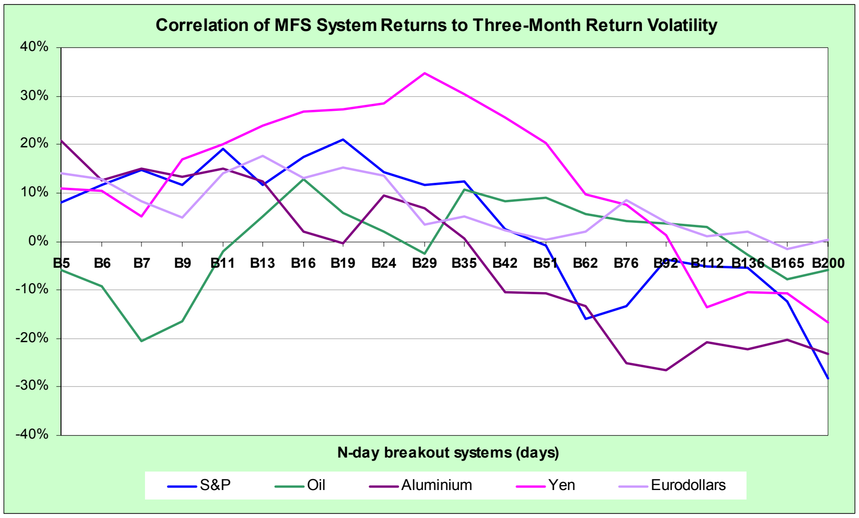

d) Volatility Exposure – Results

In this subsection we present the results for testing volatility exposure, i.e. correlations of MFS system returns to volatility changes. The results are presented in the same format as above. In all charts we see falling patterns similar to those presented and discussed in the previous subsection.

As described above, let us consider correlations of MFS system **daily returns** to **daily changes** in X-day volatility. Here the time scale X takes the same five values as above from a week to three months. The first two charts show results for price volatility. The correlations are very low and do not imply any significant dependence:

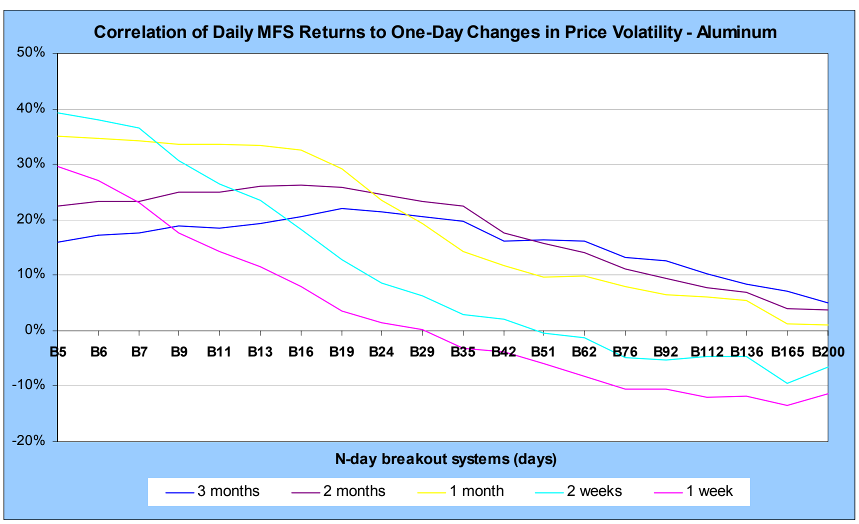

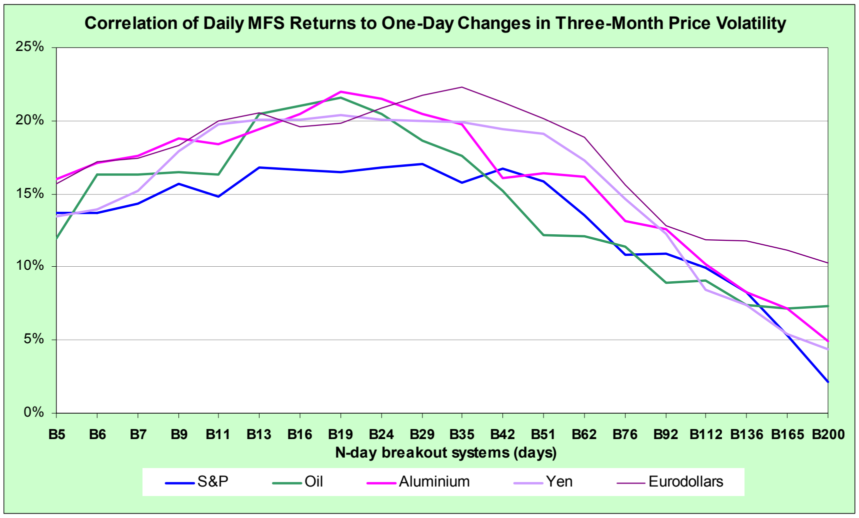

The following two charts present correlations to daily changes in return volatility, also very low:

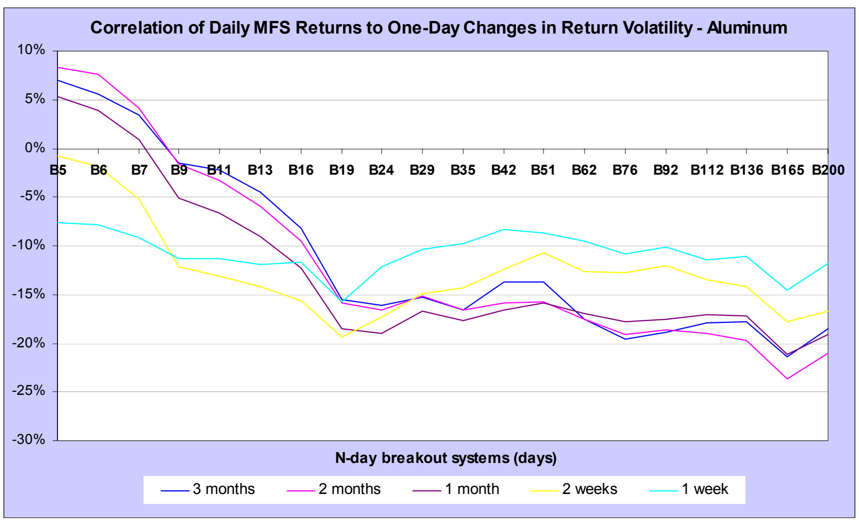

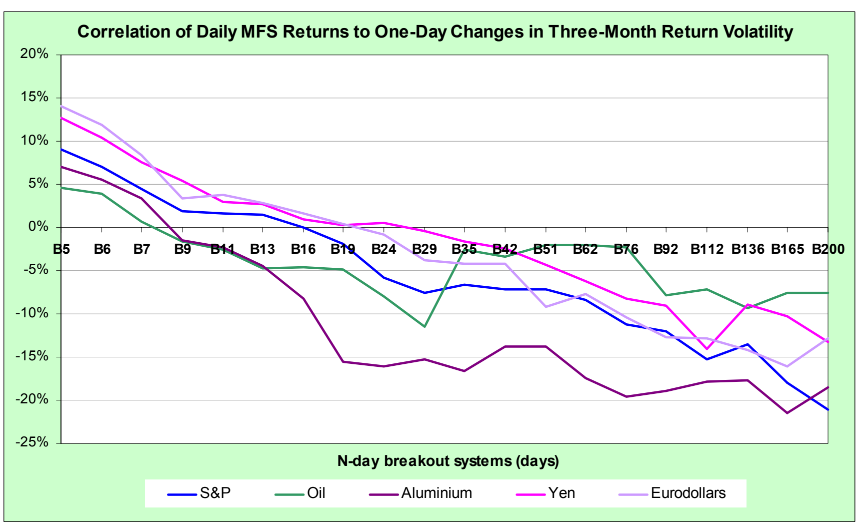

As discussed above, we can also define volatility exposure as dependence of **monthly returns** of MFS systems on **one-month changes** of **monthly volatility**. We choose the same time scale for return and volatility calculations. The time interval for calculating volatility changes is also the same, since we want the two consecutive volatility measurements to be as close to each other as possible without any overlap in the data. Instead of a month, we can substitute any of our five X-day time scales.

The highest and the only really significant volatility exposure is that of short-term system returns to long-term price volatility changes. The reason for it is simple, as was shown in Section 8. As a trend is developing from a trading range, each day the price volatility changes by the same absolute amount as the price itself. The price changes, on the other hand, coincide with returns of a trend-following system that takes a position due to the trend. And, contrary to price moves themselves, in such a trend both price volatility changes and system returns are positive and thus are highly correlated. So, on the one hand, short-term systems do have a high exposure to long-term price volatility. On the other hand, this exposure is a direct consequence of price volatility definition, which is closely related to price moves. Also, as one can see from the charts below, the correlation falls off quite fast and when one averages such dependencies for all systems in a CTA portfolio, the effects are averaged away to insignificant levels.

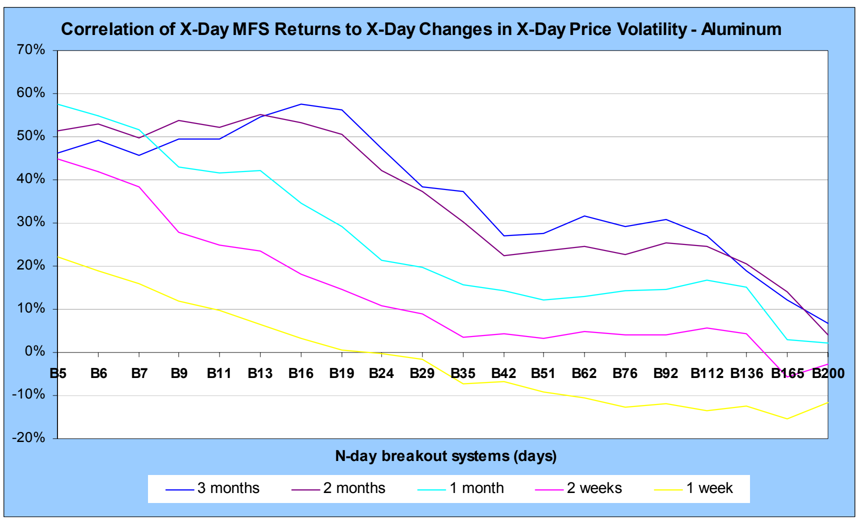

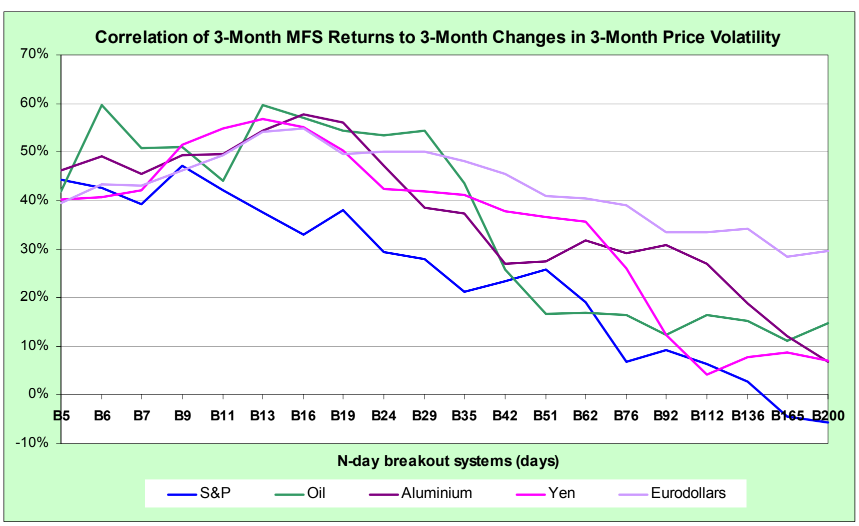

The following two charts show MFS system exposures to return volatility. Return volatility is exactly the kind of volatility that is used in option trading. If CTAs were long volatility in the same way as options are, these charts would show high correlations. Instead, their correlations never rise above 30-40%, and even those only for special time scale combinations:

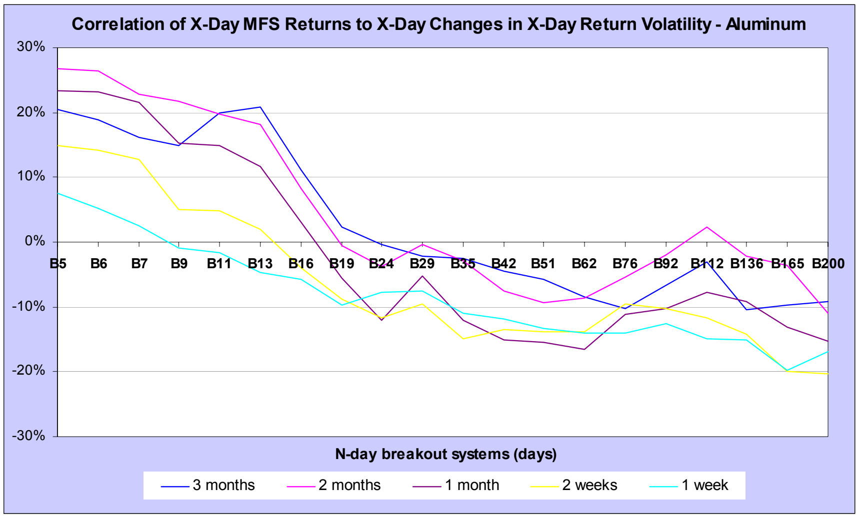

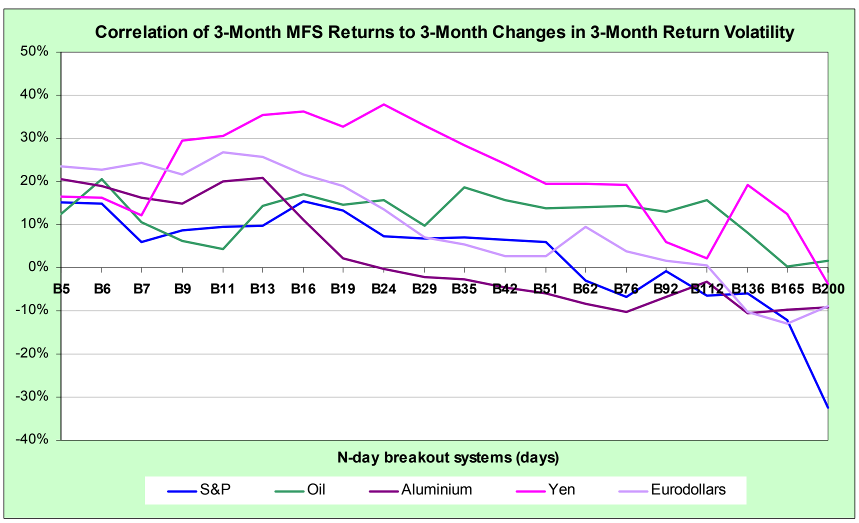

10. Conclusion

In this article, we have discussed volatility and its relevance to CTA performance. Let us reiterate the most important conclusions that were reached in our discussion:

- The notion of CTAs being long volatility due to the very nature of their trend-following strategy came from option trading similarities. It is incorrect and confusing. Using option terminology, CTAs are long gamma.
- Other hedge fund strategies that are considered to be long and short volatility are actually positively and negatively exposed to directional movements away from market equilibrium. If one has to use analogies, their

behavior is closest to buying and selling insurance on the stability of some market equilibriums. If the analogy absolutely has to be from the field of option trading, the funds are long and short gamma.

- The usefulness of exposure of hedge fund strategies to prices, volatilities or any other parameters of their underlying markets is severely undermined by the fact that hedge fund strategies usually are applied to a wide variety of uncorrelated, often unknown, markets. The practical application of knowing about any such exposure is thus very limited. To hedge strategies against each other in a portfolio, one has to check their return correlations directly.

- Volatility exposure, i.e. dependence on volatility changes, should not be confused with dependence on volatility levels. Yet, the concepts are strongly related and should be studied together.

- If a trend-following system is applied to a specific market, the only kind of volatility that it is significantly correlated to is long-term price volatility on a scale approximately equal to at least twice the scale of the system. The correlation decreases when the system time scale is increasing or the volatility time scale is decreasing. The correlations are significant both for volatility levels and volatility exposures.

- Return volatility correlations to performance are too low to imply any significant dependence even for the same markets for any time scales. Return volatility in general is unrelated to trending behavior. Price volatility, on the other hand, is so closely related to price changes as to make most of its properties obvious enough to question its usefulness as something separate from directional price action.

- If one considers long-term systems versus short-term volatility, then both return and price correlations for both volatility levels and exposures are negative. Short-term volatility for a long-term system is noise, causing false trade entries and deteriorating performance.

- And, most importantly, the overall performance of a CTA is not correlated to market volatility or its changes. Individual market volatilities are mostly uncorrelated, and therefore any volatility dependencies and exposures for individual markets and time scales are averaged to zero on the level of a typical diversified CTA portfolio.

Thus, the very notion of volatility exposure when taken out of option trading and applied to directional zero-vega strategies like CTAs is at best confusing and problematic. CTAs are "long" directional market trends and are "short" directionless choppy conditions and sharp price reversals. Such conditions are not related to market volatility in any unique predictable way. The most one can say in this respect is that CTAs are long gamma, as trend-following position entries are very similar to changing deltas of long gamma option strategies.