Correlation Smile Structures in Equity and FX Volatility Markets

Nasir Afaf

My current affiliation is with Commerzbank AG, where I am in charge of foreign exchange derivatives. This article was first written in late July 2004, when I took 3 months out of the market and was not affiliated with any institution.

Abstract

Reconciling and explaining observed volatility surfaces of equity indices from observed volatility surfaces of its constituents is an important issue for both relative value trading, and the pricing and hedging of equity options books. The same issues are to be found in the case of 'cross' rates in foreign exchange markets. This note develops the motivation behind moving to a correlation structure type approach and discusses its rationale in the case of both equity and foreign exchange markets. A method for estimating the correlation structure is then described in the local volatility framework. It is shown how the presence of liquidly quoted cross vanilla options in the foreign exchange market leads to unique correlation structures, and how this is not the case for equity options markets. It is further conjectured that a specific choice of fit, whereby pair wise correlations are dependent on the whole state of the system rather than just the two underlying stocks, may in fact have a basis in actual equity market dynamics.

Keywords

volatility, correlation, smiles, multi-factor derivatives, equity, foreign exchange.

Background

Volatility quotes in terms of full moneyness by term matrices are normally available for single market factors—these may be individual FX rates, say USD/JPY, or individual stocks, say ING Groep, as quoted on the EUROSTOXX 50 index. In equity markets as well, volatility matrices are available for stock indices—the S&P 500, EUROSTOXX50, and Nikkei are just some of the examples that leap to mind.

Equity and FX markets differ in a variety of ways. One difference is the typical definition of moneyness. FX markets prefer using a normalised version of moneyness of a vanilla option—often the delta of the underlying option. Equity markets in contrast have had a few years of debate about the merits of different version of moneyness—sticky strike and sticky delta are two of the more well known examples. It appears that sticky strike is still commonly used—especially on the short end of the curve, and many market participants have various arguments in its favour—not least the simplicity of definition. One simplicity is that pricing is consistent with Black-Scholes sensitivities—a trader's P/L at the end of day is consistent with the reported greeks (or at least ought to be!). This is not the case with sticky delta, for example, where under significant skews/smiles the Black-Scholes sensitivities of a single volatility may be significantly far removed from the actual sensitivities. Many a trader and institution has come to grief on this point—and it is fair to say not all have adapted, even at the time of writing. An option premium may be faked by inputting one volatility number into the Black-Scholes pricing formulas—but the sensitivities cannot be obtained in this manner.

All these differences notwithstanding, there are a number of similarities across FX and equity markets, albeit with some differences here as well. One is the presence of multi-factor derivatives instruments where correlation becomes an important input. Cross FX rates, while traded as a single factor, are really a multi-factor, albeit simple, instrument for traders who book their P/L in a currency different from either of the currencies that constitute the cross rate. The obvious example is the EUR/JPY cross for a trader who does his accounting in USD. Because of the liquidity of the EUR/JPY volatility market, the trader need not estimate the volatilities from USD/JPY and EUR/USD volatilities and an associated correlation (structure) between these two rates. Instead, the implied EUR/JPY volatility surface contains within it implicit information about this correlation (structure). Most often, and given the way both FX vanilla and exotics markets function at the time of writing, this is not a concern to most market participants.

While this is generally true, for anyone involved in pricing baskets on USD/JPY and EUR/USD, and perhaps with a precise or pedantic bent of mind, questions of consistency and value do arise—what correlation should one use? Relative value correlation traders and multi-factor exotics traders would be asking the same question in a different context—namely question of consistency and value across the three quoted implied volatility surfaces. To provide a concrete and still topical example, the concern might be the relative value of a EUR/JPY 10 delta butterfly trade against its 25 delta

The usual disclaimers apply. The views expressed are those of the author and not Commerzbank AG. The author is solely responsible for any inaccuracies, omissions or errors. Comments and opinions to nasir.afaf@commerzbank.com counterpart—given information about the whole EUR/USD and USD/JPY volatility surfaces.

A quick and fast answer would be along the lines of supply and demand, and market efficiency etc. While not obviously incorrect, such an answer would fail to address the real underlying issue. A thoughtful relative value trader would (should) realise that this question concerns issues of unrealised correlation—what correlation does he expect (anticipate) given prevailing market and anticipated market conditions? At the time of writing the story of the US current account deficit is writ large everywhere—if and when the big depreciation in USD comes, what would it mean for volatility surfaces in a relative sense—across EUR/USD, USD/JPY and EUR/JPY respectively? Would it be sensible to expect the same correlation for large moves in the underlying markets that one would expect for smaller moves? A thoughtful relative value trader would, if he is lucky in the resources made available to him, be able to infer correlation information from the three volatility surfaces, and see if his expectations are priced in. If not, he could put on some trades in anticipation—expecting the markets to converge favourably to his positions as events unfold.

While in the above discussion I have highlighted some issues that are of concern to a relative value correlation trader, the same concerns would apply to a thoughtful exotics trader as well—one who is concerned with putting on appropriate vega hedges for the multi-factor exotics in his books. The issues remain the same, though the end objective is slightly different.

The foregoing discussion on FX markets leads us to into equity volatility (correlation) markets. Here, the differences lies primarily in lack of liquidity to hedge the correlation exposure (recent growth of explicit correlation exotics, like correlation swaps, not withstanding), and in the number of underlying factors at hand. In the FX example above, the trader was concerned only with EUR/JPY, which is a simple product of EUR/USD and USD/JPY, just two underlying factors. Secondly, while an exotic written on EUR/USD and USD/JPY has a correlation exposure, this can in principle be hedged away by trading the cross vega in the liquid implied EUR/JPY options market. By and large, the equity trader does not have such a luxury.

The equity markets trader, unlike his FX counterpart, does not have access to complete information—primarily resulting from the large number of possible combinations that can be readily created. EURSTOXX50 volatility surfaces for individual constituents may be fairly complete and readily available—as is the case for options on the index itself (the whole basket). Yet volatility surfaces for pair wise combinations are rarely fully available—which is where the missing information lies.

What an equity trader observes, on the other hand, are the full set of volatility matrices for the underlying stocks, and a volatility matrix for the index (basket). A thoughtful equity trader, like his FX counterpart, would be keen to reconcile the two. How does one move from individual volatility matrices to the index volatility surface? It would surely be nice if simple estimates of correlation would help create a volatility matrix for the index in line with the observed one.

This would perhaps be asking too much—and in fact, from anecdotal evidence, appears to be so. Consistency between volatility surfaces of individual stocks and the volatility surface of the overall index is not easily achieved by ascribing a single correlation number for each pair. One would expect risk premia to be an important factor in the determination of the index volatility surface. This in particular immediately leads us into the idea of a correlation structure.

Most market participants would generally agree that trending markets in equities generally exhibit higher correlation (at least on the downside)—and there are a number of easily conceivable reasons why this might be so. Non trending markets tend to get de-correlated—even though significant positive correlation is still the norm. Let's assume for the moment that it is so—what would one then expect for the relative observed skew between individual stocks and the basket?

A perceptive trader would say that the above scenario for correlation suggests that skew for larger moves (far out of the money options) on the index would be significantly above observed skews on the underlying constituents (for far out of the money options). For options not very far from the money, this effect would be less pronounced. The overall point is that introducing a reasonable correlation structure in terms of observed dynamics of equity markets, and market participant behaviour, would introduce a relative spread between skews(smiles) on the index and skews (smiles) in the underlyings.

The correlation structure that leads to such an index volatility surface, and conversely provides the explanation for observed implied data, would not just be an abstract mathematical construct and fit—with little or no financial justification. Quite the contrary in fact. In the hands of an astute correlation trader, such a structure would shed light on the underlying markets, and identify opportunities both for relative value trading, and for hedging correlation exposure.

With the above motivation in mind, I propose an *ansatz* for calculating implied or consistent correlation structures. I use the term *implied* when a *unique* correlation structure can be unveiled, in the sense of complete information and liquidity, as is the case in FX markets, and *consistent* for incomplete information, as is the case in equity markets. It should be clear that a *consistent* correlation structure so obtained may not be *unique*.

Correlation Structure in the Local Volatility Framework

I highlight the method in the local volatility[1] framework. It is useful to keep in mind the schematic chart depicted in Figure 1:

I shall take the above graph to mean that we can move from the volatility surface to implied distributions, or from the volatility surface to local volatility—in fact, in all possible directions. Namely, that knowledge of any one of the three circles above is enough for us to recreate the other two.[2]

Assume that the index Y is a function of N underlying factors:

$$Y \equiv Y(x_1, \ldots\ldots, x_N) \tag{1}$$

and that the SDE's for the single factor options (underlyings) are given as follows:

$$dx_i = a_i(x_i, t)dt + \sigma_i(x_i, t)dW_i; \quad 1 \leq i \leq N \tag{2}$$

Note that I have used a slightly different definition of local volatility above from the one conventionally used. I now assume the presence of an index, called Y, which is a function of the above x_i. The SDE for the index is then given by Ito's lemma:

$$dY = \sum_i \frac{\partial Y}{\partial x_i} a_i(x_i, t)dt + \sum_{i,j} \frac{\partial^2 Y}{\partial x_i \partial x_j} \sigma_i \sigma_j \rho_{ij} dt + \sum_i \frac{\partial Y}{\partial x_i} \sigma_i dW_i \qquad (3)$$

However, we also observe the volatility surface for the index Y directly. Assume that its SDE is given by:

$$dY = a(Y, t)dt + \sigma_y(Y, t)dW_y \qquad (4)$$

Equations 2 and 3 provide two different SDE's for the index Y. Taking the instantaneous variance of equations 2 and 3 and equating them gives:

$$\sigma_Y^2(x_i, ..., x_N; t) = \sum_{i,j} \frac{\partial Y}{\partial x_i} \frac{\partial Y}{\partial x_j} \sigma_i(x_i, t)\sigma_j(x_j, t)\rho_{i,j} \qquad (5)$$

Equation 5 shows that correlation $\rho_{i,j}$ contribute to the local volatility of Y, and hence may be regarded as an instantaneous and time and spot dependent correlation. Let us now term it the local correlation structure:

$$\rho_{ij} \equiv \rho_{ij}(x_1, ..., x_N; t) \qquad (6)$$

Case 1—"Cross FX"

In the preceding discussion we discussed the relation of EUR/JPY volatility surfaces to those of EUR/USD and USD/JPY. Regarding EUR/JPY as a function of just the two factors, EUR/USD and USD/JPY, it is clear that one can back out a unique correlation structure from equation 5.

In terms of a local volatility description, this correlation structure provides consistency across all three implied volatility surfaces. It may be used for relative value trading or for the pricing of other exotics.

From the above example, we can see that FX is a slightly easier case, as volatility surfaces on the major crosses are readily available, so implied local correlation structures may be inferred for all exchange-rate pairs.

Case 2—"Equity Indices"

Here equation 6 indicates that we have great freedom in choosing the local correlations, $\rho_{ij}(x_1, ..., x_N; t)$, to match the observed left hand side—the implied local volatility of the index as derived from the index volatility surface. For N factors, we have $N(N-1)/2$ correlation to play around with. While the flexibility is welcome, it is clear that some fits achieved may have little meaning from an economic point of view.

One way around this problem is to try a proportional fitting technique across all correlations. By this I mean that if fitting a flat correlation structure does not yield the local volatility of the index, all correlation should be perturbed in the same direction. In a world of just 3 underlying stocks, for example, if the chosen correlation (say the implied ATM correlations) do not yield the correct local volatility of the index (say are below), then all 3 correlations may be proportionately moved up till equality is achieved. This is equivalent to writing equation 5 in the following manner:

$$\sigma_Y^2(x_i, ..., x_N; t) = \sum_i \left(\frac{\partial Y}{\partial x_i}\right)^2 \sigma_i^2 + 2\sum_{i \neq j} \frac{\partial Y}{\partial x_i} \frac{\partial Y}{\partial x_j} \sigma_i(x_i, t)\sigma_j(x_j, t)\rho_{ij}(ATM)\alpha(x_1, .., x_N; t) \qquad (7)$$

Here the $\rho_{ij}(ATM)$ are defined to be implied correlations using ATM volatilities—which are known and the choice is reduced to finding the function $\alpha(x_1, .., x_N; t)$, which is the only unknown in equation 7.

Discussion

One important point to note about equations 6 and equations 7 is that the pair wise correlations, $\rho_{ij}(x_1, ..., x_N; t)$, may be state-dependent. In other words, nothing precludes ρ_{ij} from depending on all the $(x_1, ..., x_N; t)$ respectively rather than just $(x_i, x_j; t)$. While there is no reason why this should be the case (indeed equation 5 could conceivably be fitted with the constraint that the individual ρ_{ij} are functions of only $(x_i, x_j; t)$), it is clear that we can make the correlations dependent on the information for the entire state at a given time—in other words, the $(x_1, ..., x_N; t)$.

The method suggested in equation 7 to achieve a fit in fact achieves this explicitly. All correlations move up or down by the same multiplicative factor $\alpha(x_1, .., x_N; t)$. In fact we have the equation:

$$\rho_{ij}(x_1, ..., x_N; t) = \rho_{ij}(ATM)\alpha(x_1, .., x_N; t) \qquad (8)$$

This simply achieves the following: as the market moves from state to state, all instantaneous correlations move up and down proportionately as determined by $\alpha(x_1, .., x_N; t)$. Instantaneous pair wise correlations become dependent on the entire state of the market, $(x_1, ..., x_N; t)$, not just on the sub-state $(x_i, x_j; t)$.

This is perhaps not an unwelcome effect of the fitting method chosen above. Anecdotal evidence suggests that pair wise correlations tend to move together. In other words, as the market moves from one state to another, pair wise correlations tend to move in unison—at least in some average manner.

The following two graphs from the Global Titans Index are illustrative.

The time series of correlation above is on a data set of 3500 days—so roughly 10 years. Each data point was constructed from 90 day periods—chosen to be non overlapping here. A quick glance suggests that pair-wise correlations do tend to move together. In fact, it is interesting to look at the above graph in light of correlation of the pair wise correlation time series.

Table 1 is interesting in that the correlation of correlation is roughly around 50% on average. Note that since IBM is the base stock above, one would really expect zero correlation across the grey row and column—that is indeed the case, but rounding errors in Excel give non-zero numbers. Note as well that the average correlation of correlation is pretty

TABLE 1: CORRELATION OF CORRELATION TIME SERIES WITH IBM AS BASE STOCK—RELATING TO FIGURE 2 ABOVE

CORRELATION of CORRELATION - 90 day non overlapping periods, 10 year history back from July 22, 2004

	C UN	HSBA LN	**IBM UN**	NESN VX	PEP UN	PFE UN	PG UN	ROG VX	RDA NA	UNA NA
C UN	100%	45%	7%	88%	31%	50%	47%	50%	48%	44%
HSBA LN	45%	100%	-17%	44%	46%	41%	46%	46%	58%	44%
IBM UN	7%	-17%	100%	9%	-11%	-9%	-15%	-11%	-10%	0%
NESN VX	88%	44%	9%	100%	37%	56%	50%	66%	53%	50%
PEP UN	31%	46%	-11%	37%	100%	66%	62%	40%	34%	61%
PFE UN	50%	41%	-9%	56%	66%	100%	71%	60%	45%	60%
PG UN	47%	46%	-15%	50%	62%	71%	100%	41%	61%	51%
ROG VX	50%	46%	-11%	66%	40%	60%	41%	100%	51%	54%
RDA NA	48%	58%	-10%	53%	34%	45%	61%	51%	100%	48%
UNA NA	44%	44%	0%	50%	61%	60%	51%	54%	48%	100%
Average	**50.39%**	**46.16%**		**55.46%**	**47.01%**	**56.18%**		**51.03%**	**49.58%**	**51.52%**

(Data Source: Bloomberg)

TABLE 2: CORRELATION OF CORRELATION TIME SERIES WITH PG AS BASE STOCK—RELATING TO FIGURE 3 ABOVE

CORRELATION of CORRELATION - 90 day non overlapping periods, 10 year history back from July 22, 2004

	C UN	HSBA LN	**IBM UN**	NESN VX	PEP UN	PFE UN	PG UN	ROG VX	RDA NA	UNA NA
C UN	100%	15%	26%	100%	27%	45%	-24%	55%	34%	34%
HSBA LN	15%	100%	36%	16%	22%	49%	-2%	18%	36%	-3%
IBM UN	26%	36%	100%	27%	38%	55%	-7%	9%	6%	-9%
NESN VX	100%	16%	27%	100%	26%	46%	-24%	56%	35%	34%
PEP UN	27%	22%	38%	26%	100%	30%	-2%	17%	26%	5%
PFE UN	45%	49%	55%	46%	30%	100%	-12%	21%	28%	-12%
PG UN	-24%	-2%	-7%	-24%	-2%	-12%	100%	-19%	-11%	-6%
ROG VX	55%	18%	9%	56%	17%	21%	-19%	100%	38%	29%
RDA NA	34%	36%	6%	35%	26%	28%	-11%	38%	100%	7%
UNA NA	34%	-3%	-9%	34%	5%	-12%	-6%	29%	7%	100%
Average	**41.87%**	**23.64%**	**23.44%**	**42.39%**	**23.79%**	**32.58%**		**30.31%**	**26.10%**	**10.69%**

(Data Source: Bloomberg)

close to 50%. It is in fact 51.24%, excluding self-correlations and the greyed cells.

IBM is a random choice and perhaps the numbers say more about the base stock chosen than highlighting the size of the effect. It is clear that credit ratings, sectors and other essential factors and information would be expected to have an effect on the numbers obtained.

Over the same period, it would be instructive to use some other stock as a base—just for the sake of comparison. I have randomly chosen

>

TABLE 3: CORRELATION OF CORRELATION TIME SERIES WITH HSBC AS BASE STOCK—RELATING TO FIGURE 4 ABOVE

CORRELATION of CORRELATION - 90 day non overlapping periods, 10 year history back from July 22, 2004

	C UN	*HSBA LN*	***IBM UN***	*NESN VX*	*PEP UN*	*PFE UN*	*PG UN*	*ROG VX*	*RDA NA*	*UNA NA*
C UN	100%	-9%	20%	100%	56%	38%	59%	72%	54%	71%
HSBA LN	-9%	100%	-11%	-11%	4%	-18%	17%	5%	-2%	-15%
IBM UN	20%	-11%	100%	21%	25%	50%	25%	27%	26%	39%
NESN VX	100%	-11%	21%	100%	55%	39%	56%	72%	55%	71%
PEP UN	56%	4%	25%	55%	100%	35%	64%	63%	50%	52%
PFE UN	38%	-18%	50%	39%	35%	100%	35%	53%	35%	39%
PG UN	59%	17%	25%	56%	64%	35%	100%	53%	32%	58%
ROG VX	72%	5%	27%	72%	63%	53%	53%	100%	61%	69%
RDA NA	54%	-2%	26%	55%	50%	35%	32%	61%	100%	53%
UNA NA	71%	-15%	39%	71%	52%	39%	58%	69%	53%	100%
Average	**58.70%**		**29.19%**	**58.46%**	**49.86%**	**40.40%**	**47.72%**	**58.77%**	**45.75%**	**56.24%**

(Data Source: Bloomberg)

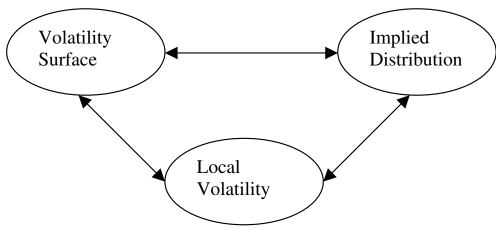

Figure 1: The relation between volatility surface, implied distributions and local volatility.

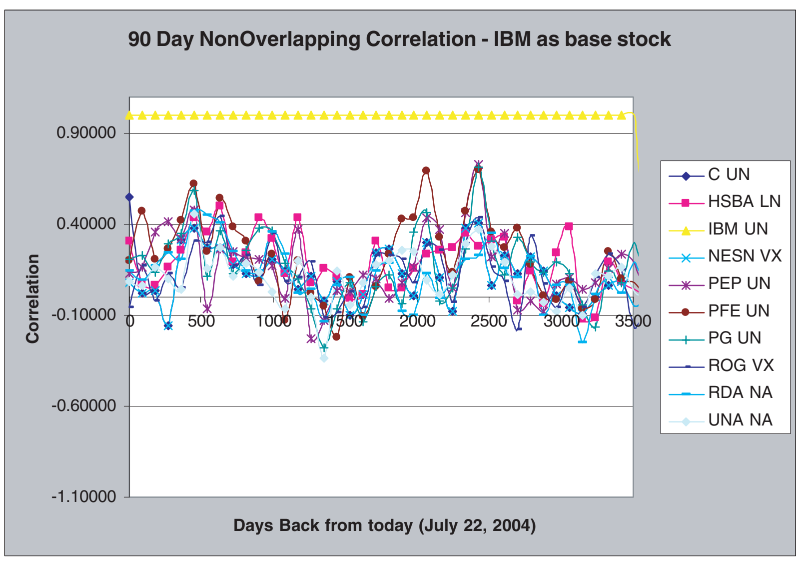

(Data Source: Bloomberg)

Figure 2: 90 Day correlation on individual Global Titans against IBM over last 10 years.

PG here. This time we see greater dispersion in the correlation time series chosen in Table 2:

Comparing Figure 2 with Figure 3, we can see that the envelope of time series is a little broader. The correlation of correlation time series with PG as a base stock is shown below:

Here it is clear that the correlation of pair-wise correlation with PG as base stock is lower than with IBM as base stock. In fact the average is now 28.31%, so dropping by around 24% from the results for the corresponding period with IBM.

For one final example I now choose HSBC as a base stock keeping the same period. Table 3 depicts the numbers and Figure 4 the graphs associated with this choice.

The correlation of correlation table is given below. In this case the average correlation is now back up at 50% (in fact 49.45%)

Discussion of Method and Suggestions for Further Research

The examples chosen above were for illustrative purposes only—and were randomly chosen. A full statistical study would be in order before any level of confidence can be achieved. Nevertheless the above numbers are encouraging in that they seem to point to a pattern that suggests that pair wise correlations tend to move together in unison—depending on the full state of the system. In particular it does seem that pair wise ρ_{ij} depend on $(x_1, \ldots, x_N; t)$ rather than just $(x_i, x_j; t)$. In other words the full dependence may be written as $\rho_{ij}(x_1, \ldots, x_N; t)$. It further appears that, in the local volatility framework described elsewhere in this paper, the choice of fit suggested in equation 8 may have some underlying meaning in terms of dynamics of actual markets.[3]

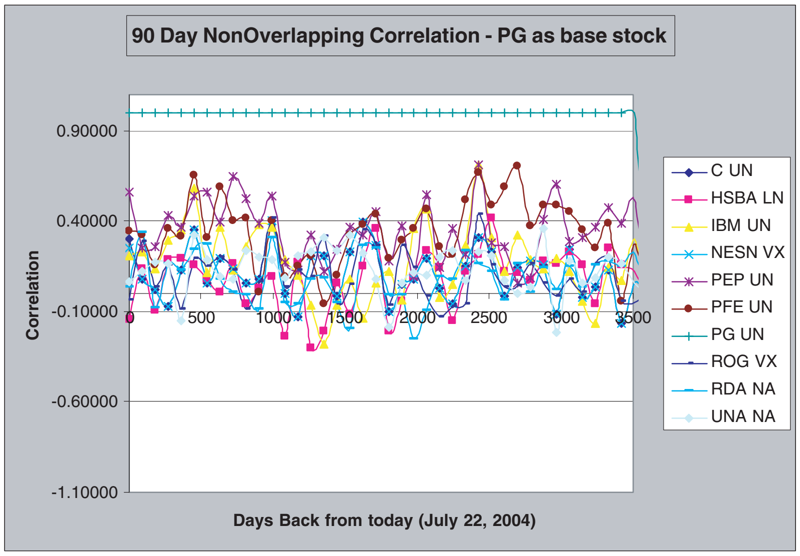

(Data Source: Bloomberg)

Figure 3: 90 Day correlation on individual Global Titans against PG over last 10 years.

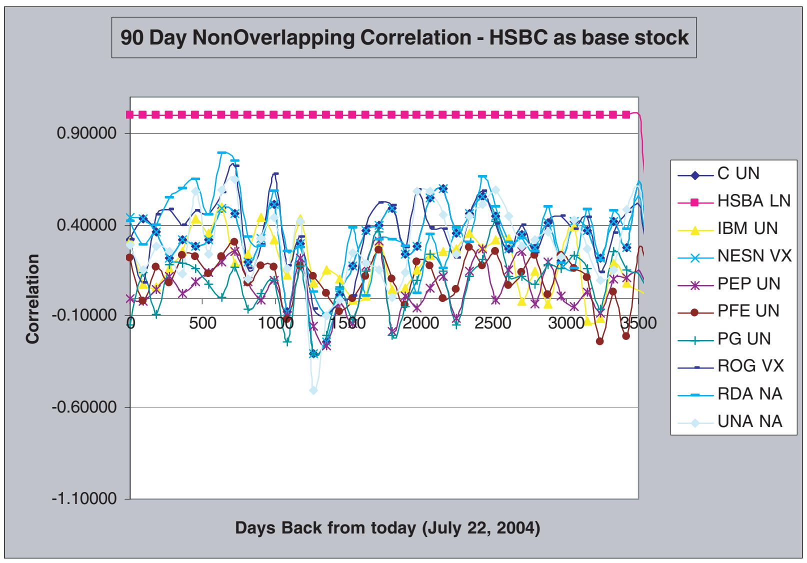

(Data Source: Bloomberg)

Figure 4: 90 Day correlation on individual Global Titans against HSBC over last 10 years.

Appendix

Equations 5 and 7 can be solved explicitly in simple cases. Say we first set

$$\rho_{ij}(x_1,\ldots,x_N;t) = \rho(x_1,..,x_N;t) \quad \text{for all } i,j \qquad (9)$$

This means that all pair wise correlations are the same, though state dependent—clearly not a very precise assumption, but one which sheds some light on the correlation structure so obtained. From equation 5, we then get:

$$\rho(x_1,..,x_N;t) = \frac{\sigma_Y^2(Y(x_1,\ldots,x_N);t) - \sum_i \left(\frac{\partial Y}{\partial x_i}\right)^2 \sigma_i^2(x_i;t)}{2\sum_{i\neq j} \frac{\partial Y}{\partial x_i}\frac{\partial Y}{\partial x_j}\sigma_i(x_i;t)\sigma_j(x_j;t)} \qquad (10)$$

Equation 10 is useful in that we see the correlation structure obtained is not entirely unintuitive. It is simply the difference between the instantaneous variance on the index and the sum of the instantaneous variances of its constituents (appropriately weighted).

Keeping our goal of simplest possible fits we can go a touch better by rewriting the pair-wise state dependent correlations as:

$$\rho_{ij}(x_1,\ldots,x_N;t) = \langle\rho_{ij}\rangle\,\alpha(x_1,..,x_N;t)$$

where $\langle\rho_{ij}\rangle$ are some mean level of correlation per chosen pair. In this case we obtain:

$$\alpha(x_1,..,x_N;t) = \frac{\sigma_Y^2(Y(x_1,\ldots,x_N);t) - \sum_i \left(\frac{\partial Y}{\partial x_i}\right)^2 \sigma_i^2(x_i;t)}{2\sum_{i\neq j} \langle\rho_{ij}\rangle \frac{\partial Y}{\partial x_i}\frac{\partial Y}{\partial x_j}\sigma_i(x_i;t)\sigma_j(x_j;t)}$$

and we can then use equation 11 to get the resulting correlation structure. Not that in both cases above, the pair-wise state dependent correlations so obtained will be 100% correlated with each other—which is clearly not the case in reality. However, it does capture some essence of underlying markets, while achieving a consistency of fit between the index volatility surface and its constituent surfaces.

Any information in this report is based on data obtained from sources considered to be reliable, but no representations or guarantees are made by the author or Commerzbank AG with regard to the accuracy or completeness of the data. The opinions, statements and calculations contained herein constitute the author's opinion and work at this date and time, and are subject to change without notice. This report is for information purposes, it is not intended to be and should not be construed as a recommendation, offer or solicitation to acquire, or dispose of, any particular securities or a recommendation to adopt a particular trading strategy.

FOOTNOTES

1. Dupire, Bruno (1994). Pricing with a smile, *Risk*, 7 (1), 18-20.
2. Technically, constraints have to be imposed in terms of integrability, choice of stochastic differential equation etc—but I will assume that they have been appropriately imposed.
3. Where we had set $\rho_{ij}(x_1,\ldots,x_N;t) = \rho_{ij}(ATM)\alpha(x_1,\ldots,x_N;t)$