A Comparison of U.S. and Hong Kong Cap-Floor Volatility Dynamics

Paul D. McNelis[*]and Salih Neftci[†]

November 2003

Abstract

In this paper we investigate the dynamics of Hong Kong Cap/Floor volatilities and compare their dynamics with the US Cap/Floor volatilities. We use linear and non-linear factor models and VAR's. The results show that the first principal components, both linear and non-linear, do a very good job in explaining the dynamics of the volatility curve and but there is not much to be gained by moving to non-linear models for the case of Hong Kong data. Secondly, we see that Hong Kong cap-floor volatilities cannot be obtained from the USD cap-floor volatilities by simply adding a volatility spread. The two sets of volatilities are non-trivially related to each other.

Key words: Cap-floor volatilities, linear and non-linear principal components

1   Introduction

Caps and floors provide a liquid platform for trading interest rate volatilities. In this paper we study the time series properties of cap-floor volatilities for Hong Kong and then compare them with those denominated in the US Dollar. Given the currency peg in use in Hong Kong, these volatilities should bear a significant resemblance to each other. The main difference would be due to

[*]Department of Economics, Georgetown University, Washington, D.C. 20057-1036. Email: McNelisP@georgetown.edu

[†]Department of Economics, Graduate Center, City University of New York, New York, N.Y. 10016-4309. Email: SNeftci@wwc.com

the risk of a brake-up in the currency peg in Hong Kong. We investigate if this and other idiosyncratic risks are important enough to make a difference in the time series behavior of quoted cap-floor volatilities. After all, the Hong Kong interest rate swap curve is priced off the USD swap curve, and as a first approximation Hong Kong cap-floor volatilities should closely resemble their USD counterparts.

U.S. caps and floors are very liquid instruments. In Hong Kong the liquidity is relatively less than in the US, but the markets still exist. In Hong Kong there is a relatively more liquid swaption market. Hong Kong authorities have also started to develop the local mortgage sector by establishing the Honk Kong Mortgage Corporation (HKMC). HKMC purchases mortgages from banks and finances these by issuing general debt securities. Any negative convexity that is due to pre-payment of mortgages has to be hedged. The cap-floor volatilities can potentially play an important role in this hedging activity, especially in the near future when the activities of the HKMC become more significant.

In order to study the dynamics of cap-floor volatilities in Hong Kong and compare them with the USD caps and floors, we estimate linear and non-linear principal components of the implied volatility series. We show that Hong Kong cap-floor volatilities can essentially be represented by a single factor. This single factor is made of a simple average of cap-floor volatilities of different maturities. It does not make much difference whether linear or non-linear principal components are used. These regularities associated with the Hong Kong cap-floor volatilities match to a large extent the observed properties of USD cap-floor volatilities.

The results presented here are relevant for several reasons. First of all we can see to what extent the volatility curve in Hong Kong is priced off the USD volatility curve. Second, the results can give hints about the effect of the currency peg on the implied volatilities. It is true that the risks associated with the currency peg introduce a spread between the swap rates in the USD and HKD. Yet, it is not clear to what extent the currency peg leads to a an extra risk in implied volatilities. The results may tell us if this risk can be accounted by a second factor. Third, it is clear from the above that, if the peg introduces a wedge in interest rates but not in volatilities of interest rates, then the pricing of Hong Kong mortgage instruments would be significantly simplified. In the opposite case when the currency peg introduces a wedge at the level of volatilities, the pricing of Hong Kong mortgage instruments would become significantly more complicated.

The paper is organized as follows. First, we give definitions of cap and floor volatilities to provide the necessary framework for the data. We then provide a description of our data sources. Next we introduce the non-linear principal component model and discuss the numerical issues that arise in this context. Following this, we give the results. The last section is the conclusion, where we assess the relevance of this approach.

2 Implied volatility of Caps and Floors

Let L_{t_i}, f_{t_o}, N be the relevant Libor rate, forward swap rate, and the forward swap notional respectively. Let δ be the reset interval.

The cash flows of a *caplet* with expiration date t_i and settlement date t_{i+1} will replicate the payoff of an option written on the Libor rate L_{t_i}. The client receives the difference at time t_{i+1} if the Libor observed at time t_1 exceeds a *cap level* κ which in this case is $\kappa = f_{t_o}$. This instrument is labelled a *caplet*, and is denoted by Cl^{t_i}. The floorlet is similar but plays the role of insurance against a drop of the Libor rate below the *floor level* κ. It is denoted by Fl^{t_i}.

By putting n caplets(floorlets) together we get the n period interest rate cap (floor), that starts at time t_o and ends at $T = t_n$:

$$Cap = \{Cl^{t_1}, ..., Cl^{t_n}\} \tag{1}$$

Similarly, the floorlets can be grouped as an n period interest rate floor:

$$Floor = \{Fl^{t_1}, ..., Fl^{t_n}\} \tag{2}$$

The market practice on pricing a typical caplet, Cl^{t_i}, written on L_{t_i}, proceeds as follows. First, let the $f(t, t_i)$, be the forward interest rate decided at time t on a loan that will begin at time t_i and end at time $t_i + \delta$. Market practice *assumes* that the forward Libor rate $f(t, t_i)$, has a Martingale representation with *constant* instantaneous percentage volatility σ. [1] Secondly, the market assumes that the arbitrage-free value of a t_{i+1} maturity default-free pure discount bond price denoted by $B(t_o, t_{i+1})$ can be calculated. Under these conditions the time t_o value of the caplet Cl^{t_i} is calculated using the expression:

---

[1]If one uses the correct forward measure this assumption will always be correct.

$$Cl_{t_o}^{t_i} = B(t_o, t_{i+1}) \left[F(t_o, t_i)N(h_1) - \kappa N(h_2) \right] \tag{3}$$

where

$$h_1 = \frac{\log(\frac{F(t_o,t_i)}{\kappa}) + \frac{1}{2}\sigma^2(t_i - t_o)}{\sigma\sqrt{t_i - t_o}} \tag{4}$$

$$h_2 = \frac{\log(\frac{F(t_o,t_i)}{\kappa}) + \frac{1}{2}\sigma^2(t_i - t_o)}{\sigma\sqrt{t_i - t_o}} - \sigma\sqrt{(t_i - t_o)} \tag{5}$$

To get the value of the Cap itself the market simply adds these caplet prices:

$$Cap_{t_o} = \sum_{i=1}^{n} B(t_o, t_{i+1}) \left[F(t_o, t_i)N(h_1) - \kappa N(h_2) \right] \tag{6}$$

where the h_i, g_i are as given above. We can now discuss the implied volatility.

Using these formulas we can derive the cap-floor volatility. The market quotes a single cap volatility σ such that:

$$\sum_{i=1}^{n} Cl(\sigma_i, F(t, t_i), t) = \sum_{i=1}^{n} Cl(\sigma, F(t, t_i), t) \tag{7}$$

This suggests that the cap-floor volatility will be a non-linear weighted average of the caplet volatilities. The highest weighted caplets being those that re-close to being in the money in terms of the single caprate κ. This implies that as the market fluctuates significantly, due to both the smile effect and the changing weights on the constituent caplets, the cap-floor volatilities may shift suddenly and in a non-linear fashion. This, in fact, is confirmed by observed stylized facts.

3   The data

We use data on Cap-Floor volatilities for the US and Hong Kong dollar for the period 1997-2003. We have 5 series for each economy. The data used for the study were obtained from historical cap-floor volatilities as reported in Bloomberg. We used bid, ask and the last prices for the volatilities. The results reported here are for the *last* prices. The 3,4,5,7 and 10 year cap-floor maturities were selected.

Historical Bloomberg data always contain some typos and had to be cleaned carefully. Also, some special holidays in US and Hong Kong introduce non-synchronous trading days. Rather than interpolating the data using simple averages for the previous and next trading days, we eliminated these periods for both data sets. In any case, the data being daily, the effect of such special holidays and days with missing data would be minimal. This is especially true since volatilities themselves are not very volatile during special holidays. Market practitioners usually take long weekends even when the underlying markets are technically open.

4 Non-linear Principal Components

The question of interest to us in this paper concerns the existence of linear and non-linear principal components in USD and HKD sectors. In particular, we would like to know what are the differences (if any) between the USD and HKD principal components. The question is of interest for several reasons, the most important relating to the Hong Kong currency peg. It is well known that the risks associated with the currency peg can add a spread over the USD sovereign yield curve. The main question is whether this risk results in a spread between the USD and HKD *volatilities* as well. We investigate this issue using the USD and HKD cap-floor volatilities.

Let $x_t^i, i = 1, .., 5$ denote the i'th cap volatility curve. Then we posit the following multivariate model:

$$\begin{pmatrix} x_t^1 \\ ... \\ x_t^5 \end{pmatrix} = \sum_{j=1}^5 \begin{pmatrix} \alpha_1^j \\ ... \\ \alpha_5^j \end{pmatrix} Z_t^j + \begin{pmatrix} \epsilon_t^1 \\ ... \\ \epsilon^5 \end{pmatrix} \tag{8}$$

Here the $Z_t^i, i = 1, ..., 5$ represent up to 5 linear or non-linear common principal components.[2] Their functional form will be specified below. The α_i^j represent the sensitivity of the i'th cap volatility to j'th common factor. Finally the ϵ_t^i are i'th volatility's innovation. These innovations have a non-zero contemporaneous variance covariance matrix, but are assumed to be uncorrelated over time:

$$E\left[\epsilon_t \epsilon_t'\right] = \Omega \tag{9}$$

---

[2]Some sort of normalization must be imposed on the coefficients, since as written the system will be un-identified.

$$E\left[\epsilon_t \epsilon_u'\right] = 0 \qquad \text{for } u \neq t \qquad (10)$$

The non-linear factors are modelled using a Logit type structure:

$$Z_t^j = \frac{1}{\left(1 + e^{-\sum_{i=1}^{5} \beta_i^j x_t^i}\right)} \qquad (11)$$

A typical plot of this function is shown in Figure 3. We see that the Z_t^j are highly non-linear functions of the elementary cap-floor volatilities and it is this property that may indeed help to reduce the number of common factors that was found in case of linear principal components.

The parameters of the model that are of interest to us are the sensitivities $\{\alpha_i^j\}$, the $\{\beta_i^j\}$ which determine the common factors and the estimate of the ω matrix. This latter is especially useful for performing statistical tests. In order to obtain estimates of these parameters the following objective function is minimized:

$$\min_{\{\alpha_i^j, \beta_i^j, \omega\}} \begin{pmatrix} \sum_{t=1}^{N}(\epsilon_t^1)^2 \\ \dots \\ \sum_{t=1}^{N}(\epsilon_t^5)^2 \end{pmatrix}' \begin{pmatrix} \dots & \dots & \dots \\ \dots & w_{ij} & \dots \\ \dots & \dots & \dots \end{pmatrix} \begin{pmatrix} \sum_{t=1}^{N}(\epsilon_t^1)^2 \\ \dots \\ \sum_{t=1}^{N}(\epsilon_t^5)^2 \end{pmatrix}' \qquad (12)$$

where the weights w_{ij} are obtained from the first stage of minimizing the same criterion function with W selected as the identity matrix. Once the estimates of α_i^j, β_i^j are obtained, we can form the estimates of the non-linear principal components using the equation (11). Statistical tests are based upon the assumption that the estimated weighted sum of squares shown in (12) is distributed as χ_{M-k}^2, where $M = 5N$ and k is the total number of estimated parameters.

Letting this estimated sum of squares in (12) be denoted by $S(M, k)$, we can test for the hypothesis that there is only one factor by calculating the $S(M, k)$ under the assumptions of 2 factors and 1 factor, and then looking at the difference of the respective sum of squares.

The parameters of the non-linear principal component representation are obtained by hybridization. As is well-known in any non-linear estimation problem, there is no "silver bullet" to prevent convergence to local, rather than global, optima. To minimize the possibility of convergence to local optima, we make use of a "global" stochastic evolutionary search process,

the genetic algorithm, to find a net of parameters as the initial conditions for the more conventional "local" gradient-descent minimization. The version of the genetic algorithm used for our estimation was adapted by Michalewicz (1996) for encoding problems with real, as opposed to binary, solutions. Quagliarella and Vicini (1998) call this coupling of the results of global stochastic search point with "local" gradient methods "hybridization" They point out that hybridization may lead to better solutions than those obtainable using the two methods individually.

5   Results

The major statistical properties of the cap-floor volatilities for the United States and Hong Kong appear in Tables I and II. These tables give the average volatility curves for each economy. We see that the volatility curves are essentially downward sloping. The volatilities of the volatilities are estimated by the standard deviations. In each case the long end volatilities are lower than the short end volatilities.

Table I: US Data Properties		
Maturity:	Mean	Std. Deviation
2	24.7456	14.6211
3	23.8639	11.9251
4	22.7989	9.7582
5	21.8659	8.1375
7	20.3504	6.1062
10	18.8911	4.5064

Table II: Hong Kong Data Properties		
Maturity	Mean	Std Deviation
2	28.5810	12.9065
3	26.1922	10.1833
4	24.2856	8.1229
5	22.9511	6.7189
7	21.2948	5.2376
10	19.9358	4.3025

What about the principal components of the US Dollar and Hong Kong Dollar cap/floor volatilities? This information is provided in Figure 1 and Figure 2. Each has three parts. The top portion of Figure 1 shows the time series behavior of the USD volatilities during the sample period. The middle part is the time series behavior of the first *linear* principal component. The third part is the time series behavior of the non-linear principal component. Figure 2 provides the same results for the Hong Dollar cap-floor volatilities.
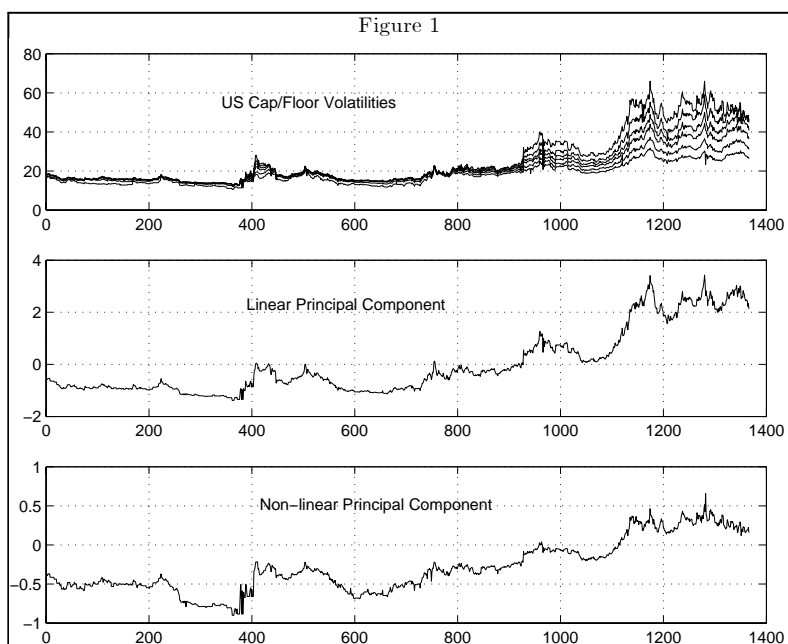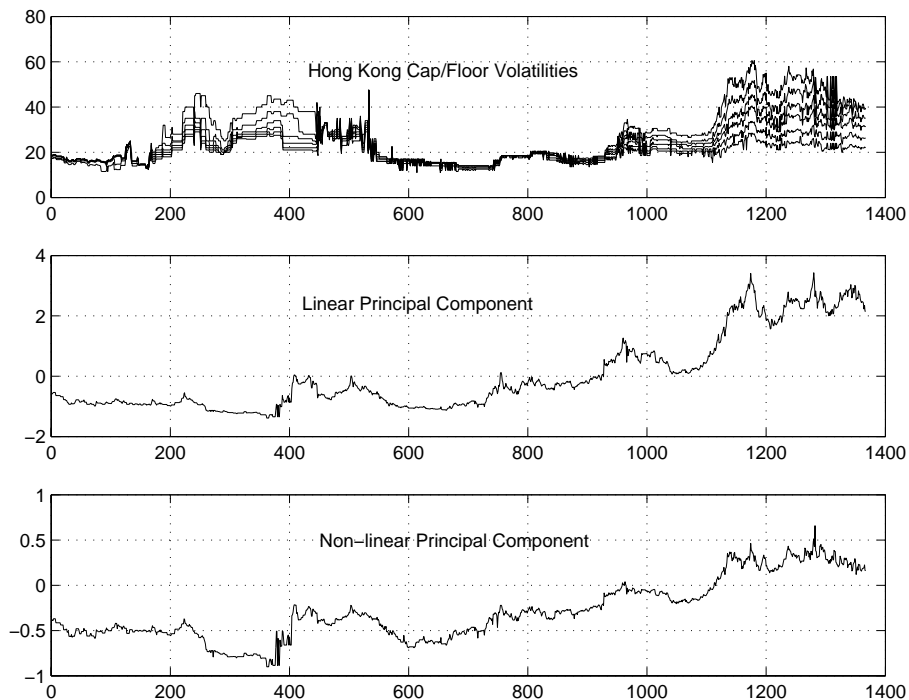

Figure 1

Figure 2

- First, we see the effect of the 1998 Asian crisis very clearly in the Hong Kong data. During the summer of 1998 the cap-floor volatilities increase significantly. This effect is much less visible in the USD cap-floor volatilities.

- The second important observation is the similarity of the linear and non-linear principal components. Especially for the case of Honk Kong, the first principal component does not change dramatically when we switch from linear to non-linear models. In case of US this is not the case. The non-linear principal component seems to capture sudden changes in cap-floor volatilities better than the linear principal component.

- More interestingly the USD and Honk Kong Dollar principal components seem to be quite similar. This suggests that the Hong Kong dollar cap-floor volatilities are priced off the USD cap-floor volatilities. Yet, this is not necessarily done by *adding* a spread to the USD cap-floor volatilities.

9

The last point is further supported by Figure 3. Here we plot the USD and Hong Kong dollar principal components on the same graph.

Figure 3
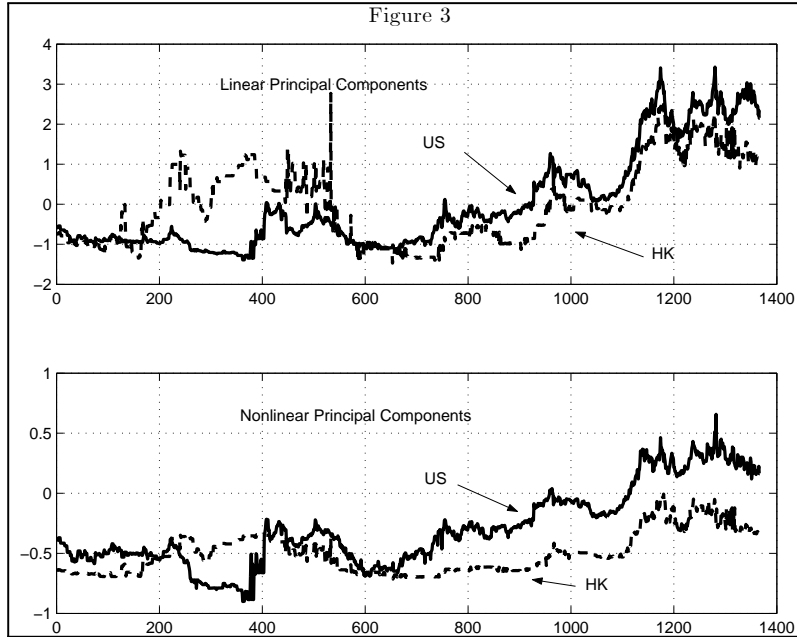
It is interesting to see that the two principal components are not "parallel" to each other. It turns out that, if Hong Kong cap-floor volatilities were priced off the USD volatilities by adding a "volatility spread" then the two indicators would display a more or less parallel time series paths. In fact, the principal components intersect at several points. This suggests that there is independent information in these two indicators. A Granger (1969, 1980) causality test indicates that there is bi-directional feedback for the non-linear principal components, whereas the US principal component is exogenous or a uni-directional cause of the Hong Kong principal component. Table III and IV, below, give the F-statistics of joint significance, as well as marginal significance levels for the causal relationships between the linear and non-linear principal component series, for five lags, representing a week of trading days.

Table III: Granger Causality, Linear Principal Components	US	Hong Kong
Regressor		
US	32263 (0)	8.3 (9e-8)
Hong Kong	.09 (.35)	4226 (0)

Table IV: Granger Causality, Non-linear Principal Components	US	Hong Kong
Regressor		
US	18249 (0.000)	11.5 (0)
Hong Kong	3.8 (.001)	24102 (0)

Tables V and VI display the explanatory power of the first linear and non-linear principal components for US and Hong Kong.

Table V: US Data, Explanatory Power of Principal Components		
US Data	Linear	Non-linear
2	0.9831	0.9892
3	0.9949	0.9809
4	0.9970	0.9750
5	0.9978	0.9737
7	0.9944	0.9714
10	0.9779	0.9676

Table VI: Hong Kong Data, Explanatory Power of Principal Components		
HK Data Set	Linear	Non-linear
2	0.9650	0.9932
3	0.9863	0.9796
4	0.9897	0.9438
5	0.9811	0.9098
7	0.9231	0.8336
10	0.7512	0.7297

We see that the R-squared statistics are all very high and similar across linear and non-linear principal components. As a matter of fact, it the principal components do surprisingly well in capturing the volatility dynamics. In the case of the US the non-linear principal component seems to have more explanatory power at the short and long end of the volatility curve.

Table VII displays the correlations of the principal components for the two economies during and after the Asian Crisis. It is interesting to see that the USD and Hong Kong principal components are very highly correlated after the Asian crisis. Yet, the correlations were significantly lower significantly during the periods of the Asian crisis.

Table VII: Correlations of Linear and Non-linear Principal Components		
	Linear	Non-linear
Total Sample	0.6903	0.6481
Post-Asian Crisis	0.9664	0.9345

6   Conclusions

There are some noticeable conclusions from processing the cap-floor volatility data. First we see that the first principal components, both linear and non-linear, do a very good job in explaining the dynamics of the volatility curve and that there is not much to be gained by moving to non-linear models in case of Hong Kong data. Secondly, we see that Hong Kong cap-floor volatilities cannot be obtained from the USD cap-floor volatilities by simply adding a volatility spread. The information contents of the two volatility curves appear to be non-trivially related to each other. The effect of the Hong Kong "peg" appears to be much less homogenous in case of volatilities.

References

[1] Fotheringame, David and Roland Baddeley (1997), "Nonlinear Principal Components Analysis of Neuronal Spike Tran Data". Working Paper, Department of Physiology, Univeristy of Oxford.

[2] Granger, C.W.J. (1969), "Investigating Causal Relations by Econometric Models and Cross Spectral Methods", Econometrica 37, 424-439.

[3] _____ (1980), "Testing for Causality: A Personal Viewpoint", Journal of Economic Dynamcis and Control 2, 329-352.

[4] Michalewicz, Zbigniew (1996), *Genetic Algorithms + Data Structures = Evolution Programs*. Third Edition. New York: Springer-Verlag.

[5] Quagliarella, Domenico and Alessandro Vicini (1998), "Coupling Genetic Algorithms and Gradient Based Optimization Techniques" in Quagliarella, D., J. Periaux, C. Poloni, and G. Winter, editors, *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science: Recent Advances and Industrial Applications*. West Sussex: England: Johns Wiles and Sons, Ltd.