Does Risk Pay? An Analysis of Short Gamma Trading Strategies and Volatility Forecasting in the Swaptions Market

Tasha Stær Bollerslev and Michael Scott Kuritzky

Professor Emma Rasiel, Faculty Advisor

Honors Thesis submitted in partial fulfillment of the requirements
for Graduation with Distinction in Economics
in Trinity College of Duke University

Duke University
Durham, North Carolina

April 15, 2008

Acknowledgements

First and foremost we would like to thank Professor Emma Rasiel, our faculty advisor, for her thoughtful critiques and suggestions. Without her, this paper would not have been possible. In addition, we are especially grateful to Professor Michelle Connolly for her constant support and encouragement, and to Professor Tim Bollerslev for his guidance and expertise. Finally, we would like to thank Richie Prager and Bank of America for providing us with the data necessary to perform this analysis.

Abstract

We evaluate short gamma trading strategies in the interest rate swaptions market from January 4th, 1999 to January 19th, 2007, and test the effectiveness of swaption proprietary forecasted volatility at predicting future realized volatility. We find that swaptions market proprietary forecasted volatility is an effective estimator; there is no risk premium priced into swaption prices and hence short gamma strategies are not profitable. We find that the market on average underprices interest rate swaptions by underestimating forward realized volatility.

I. Introduction

A swaption is a financial instrument that gives its owner the right, but not the obligation, to enter into an interest rate swap.[1] The market for swaptions has become of increasing importance recently as interest rate swaps and swaptions have become ubiquitous hedging tools against exposure to interest rates for major corporations and institutions. The notional amount of derivatives (including caps, floors, interest rates, and swaptions) held by US commercial banks, according to the Option Clearing Corporation, is worth approximately $50 trillion, with interest rate derivative contracts accounting for 80% of this total value (Fan, Gupta and Ritchken, 2006).[2] The US swaptions market specifically has a notional amount of over $7 trillion for over-the-counter contracts, highlighting the fact that swaptions are amongst the most important interest rate derivatives (Fan, Gupta and Ritchken, 2006).[3] The size of this market, coupled with a relative lack of relevant academic research on the nature of the relationship between option-implied volatility and subsequent realized volatility, makes it an interesting area for analysis. While there has been significant research done on the relationship between forecasted and realized volatility in other asset markets there has been no research on the nature of this relationship in the interest rate swaptions market. In this paper, we test the efficiency of swaption PFVs provided by an international investment bank in two ways. First, we compare the PFVs to other econometric models for forecasting future realized volatility. Second, we test whether the swaptions market consistently over-estimates

[1] An interest rate swap is a contract between two counterparties where each party agrees to pay the other's interest payments on a previously specified principal amount of debt. Typically, interest rate swaps allow the parties to exchange fixed interest rates for floating interest rates, and vice-versa.
[2] A derivative is any financial instrument whose value is based on an underlying asset. For example, an option to buy Google stock is a derivative based on the price of the underlying asset—Google.
[3] The term over-the-counter refers to any financial instrument not traded on an established exchange, and can be specialized to accommodate the specific customer's needs.

future realized swap rate volatility, as is the case in many other volatility markets. If it does, and a risk premium exists, then short selling same strike calls and puts (straddles) on interest rate swaps should be a profitable trading strategy, on average.

The format of the paper is structured as follows. Section II discusses the literature that is relevant to our topic, placing our analysis in the context of prior research. Section III is a concise description of the data used in our research. Section IV presents the theoretical framework underpinning our research, and describes the empirical specifications of our model. Finally, Section V contains a discussion of our results and Section VI is the conclusion, placing our results in the context of the relevant academic literature.

II. Literature Review

Ever since volatility was observed to be stochastic in nature by Engel (1982) and Bollerslev (1986) economists have studied the efficacy of various volatility forecasting models. However, while many academic papers analyze the efficiency of implied volatility as a predictor of realized volatility, none of these papers focus on the interest rate swaptions market.

Christensen and Prabhala (1998) examined the forecasting ability of S&P 100 equity index options, and found results that diverged from most of the previous research. While prior researchers had found that implied volatility from options on the S&P 100 index was an inefficient forecaster of future realized volatility, Christensen and Prabhala conclude that implied volatility actually outperforms realized volatility in forecasting future volatility in the period after the October 1987 stock market crash.

One of the ways in which Christensen and Prabhala differentiate themselves from previous research is that they use volatility that is sampled over a longer time period in time than in previous studies. This increases the statistical power of their regressions. The authors also sampled the implied and realized volatility series at a lower frequency, monthly, than previous researchers permitting the construction of a volatility series with non-overlapping data. As such, there is exactly one implied and one realized volatility covering each time period in the sample—resulting in much more reliable regression estimates than in previous research.

In addition to their main finding, Christensen and Prabhala also find that realized volatility was much more variable in the period following the 1987 stock market crash, whereas implied volatility was much more variable than realized volatility in the precrash period. This is indeed an interesting result as it runs counter to what these authors claim is a common notion—that since implied volatility is a smoothed expectation of realized volatility, it should be less variable than realized volatility.

Although Christensen and Prabhala reach some interesting conclusions they also recognize that there are possible errors in their calculation. First the Black and Scholes (1973) option pricing model prices European style options which are known in advance to pay no dividends (as is the case in foreign exchange and interest rate swaption trading).[4] However, the S&P 100 index pays dividends, and since stock prices fall by the amount of the dividend after the disbursement date, dividends imply lower call premiums and higher put premiums. This in effect means that the implied volatility that results from using this model on call (put) options is biased upward (downward). Additionally, the Black-Scholes model assumes that the index levels follow a log-normal diffusion process with deterministic volatility. Unfortunately this may not be true empirically for two reasons. First, volatility is stochastic in nature as found by Engel (1983) and Bollerslev (1986) and second because index levels do not perfectly follow a log normal diffusion process, as Bates (2000) observed there exists jumps in index prices that often .

Heston and Nandi (2000) compare the closed form solution of their applied GARCH (1,1) model and an *ad hoc* Black-Scholes model "in which each option has its own implied volatility depending on the strike price and time to maturity."[5] They test both of these models on option price data on the S&P 500 index market to compare each

[4] A European style option differs from an American style option in that a European option can only be exercised on its expiration date.

[5] At this point it is necessary to point out that this *ad hoc* Black-Scholes model is not the same *ad hoc* Black-Scholes model given to us by the Investment Bank—but it is safe to assume that both of these models add explanatory variables to the traditional Black-Scholes standard five in order to increase the models forecasting power.

of the models' ability to predict in and out of sample option prices. The authors find results that run counter to previous empirical work—including the aforementioned paper by Christensen and Prabhala (1998)—that tested deterministic versus stochastic volatility models.[6] Heston and Nandi (2000) find that the GARCH model has smaller out of sample valuation errors; (i.e. is a more efficient estimator of option prices) than the *ad hoc* Black-Scholes model, even though the parameters on the Black-Scholes model are updated every period whereas the parameters of the GARCH model are held constant and "the variance is filtered from the history of asset prices." The authors believe that the increased accuracy of the GARCH forecasting model is due to its ability to capture the stochastic nature of volatility, as well as the fact that it incorporates the relationship between volatility and asset returns into its prediction. Heston and Nandi also hypothesize that GARCH is a better forecasting tool because it uses the history of asset prices to value an option on any given day, while the Black-Scholes model only uses a single implied volatility to make this forecast. We also test the forecasting ability of a GARCH(1,1) against a variation of the Black-Scholes option pricing model—but instead of testing these models' ability to predict option prices, as Heston and Nandi do, we test the model's ability to predict realized volatility.

In Busch, Christensen, and Nielsen (2006), the authors look at the forecasting of future realized volatility in the stock, bond and foreign exchange markets using high frequency data over the time period from 1990 to 2002 for the S&P 500 futures options and 30 year Treasury Bond futures options, and over the time period from 1987 to 1999 in the foreign exchange market. Busch, Christensen and Nielsen find that when implied volatility is added into regressions of past realized volatility, the implied volatility still

[6]Stochastic volatility models include Engel (1983) and Bollerslev (1986).

has explanatory power above and beyond that of realized volatility in certain markets. This paper is of particular interest to us as we are dealing with another fixed income financial instrument—interest rate swaptions.

Bausch, Christensen and Nielsen include implied volatility as an explanatory variable when forecasting future volatility. They find that implied volatility has significant information about future realized volatility above and beyond that contained in both aspects of past realized volatility in all three markets. Additionally, they find that implied volatility is an unbiased forecaster of future realized volatility, in the stock and foreign exchange markets. They find that the best forecasts combining return and option prices information vary between markets. For the stock market, implied volatility together with the most recent one-day realized volatility measure is the best. In the bond market implied volatility together with monthly realized volatility measures is optimal, and finally in the foreign exchange market implied volatility completely accounts for the informational content of (daily, weekly, and monthly) realized volatility.

The author's results show that even the jump component of realized volatility is somewhat predictable and incorporated into option pricing in all three markets. Therefore traders in the option market must speculate and base their trading strategies on information about future jumps that occur in the stock, foreign exchange, and bond markets. This clearly has a profound impact upon the dynamic dependencies in the given markets, and again serves as a solid foundation for our analysis of the efficiency of the swaptions market.

No definitive conclusions have been reached on whether GARCH or Black-Scholes is a more efficient predictor of realized volatility, and more fundamentally as to whether implied volatility or past volatility is a better predictor of future realized volatility. For example, the aforementioned findings of Christensen and Prabhala (1998) and Busch, Christensen and Nielson (2006), differ from those of Day and Lewis (1992) and Lamoreux and Lastrapes (1993), who find that past volatility outperforms dealer forecasted volatility in the market for options on equities and equity indices.

It is in this space that we see our research adding to the literature as we compare the efficacy of proprietary forecasted volatility versus past volatility at predicting future realized volatility in the market for interest rate swaptions.

III. Data Description

The data provided by the Investment Bank consists of the daily closing levels of LIBOR, forward swap rates, discount factors, and proprietary forecasted volatilities for various maturity options on interest rate swaps. The data includes values for 1m, 3m, 6m, 9m and 1y LIBOR and 2y, 3y 5y and 10y swap rates[7] as well as forward rates for 1m, 3m and 6m options on 2y interest rate swaps and monthly discount factors.[8] Finally the data also includes proprietary forecasted volatility estimates for 1m, 3m and 6m options on 2y swaps, reported in basis points of volatility (BPVol).

The data covers the time period from January 4th, 1999 to January 19th, 2007. This time period is especially interesting since it includes data points from before and after the "Dot Com" crash, as well as a period in which the Federal Funds rate varied from a high of 6.50% to a low of 1.25% in 2002 and then back up to 5.25% in 2006.

Below are descriptive statistics and graphs of our data set. **Table 1** includes the mean and standard deviation for the following variables from our data over the entire time horizon: 1m, 3m and 6m LIBOR, the forward swap rate for a 1m, 3m and 6m options on a 2 year swap, both expressed in basis points and the proprietary forecasted volatility for the 1m, 3m and 6m options on a 2 year swap expressed in basis points and Black-Scholes volatility. **Figure 1** is a graph of 1m LIBOR and the 1m forward swap rate for a 2 year swap. The graphs of 3m and 6m LIBOR and the forward swap rates for 3m

[7] The LIBOR-Swap curve switches at one year from using LIBOR to using swap rates because at this time swaps become the more liquid assets.

[8] These discount rates are LIBOR zero coupon interest rates. "When swaps and other over-the-counter derivatives are valued, the cash flows are usually discounted using LIBOR zero coupon interest rates . . . because LIBOR is the cost of funds for a financial institution. The implicit assumption is that the risk associated with the cash flows of the [swap] is the same associated with a loan in the inter-bank market (Hull, 2006)."

and 6m options on a 2 year swaps, look almost identical to those of the 1m, and as such, are omitted for brevity.

Table 1:

Investment Bank Data *Time Period: 01/04/1999 - 01/19/2007*												
	LIBOR			**Forward Swap Rates**			**Proprietary Forecasted Volatility (Basis Points)**			**Proprietary Forecasted Volatility (Black-Scholes)**		
	1m	3m	6m	1m2y	3m2y	6m2y	1m2y	3m2y	6m2y	1m2y	3m2y	6m2y
Mean	3.58	3.65	3.73	4.36	4.47	4.63	98.26	102.04	106.24	0.279	0.276	0.269
StdDev	1.91	1.92	1.93	1.64	1.58	1.49	23.70	22.88	22.11	0.162	0.152	0.137

Figure 1:

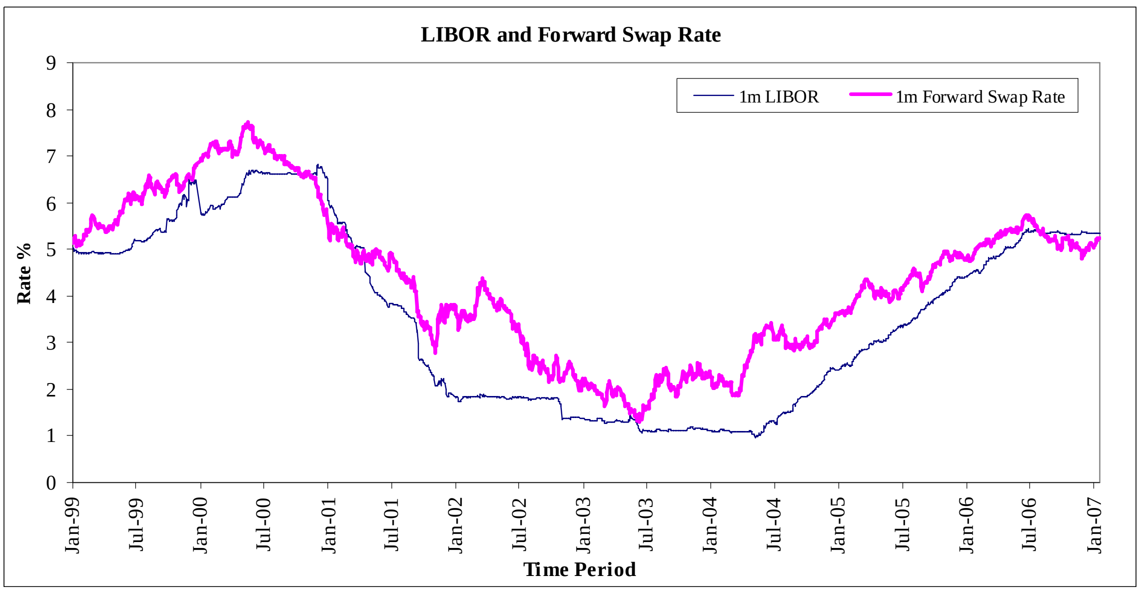

IV. Theoretical Framework and Empirical Specification

1. Trading Strategies Analysis

The first step of our analysis was to use a Black-Scholes model to price options based on the proprietary data. The model prices an interest rate swaption that gives its owner the right, but not the obligation, to pay a fixed interest rate ***X***, and receive a variable rate, LIBOR, on a swap that will last 2 years, starting ***T*** years from the option origination date. The model assumes that there are two payments per year associated with the swap, and that the notional amount is $100M. Let F represent the value of the forward swap rate at time 0, BPVol represent the basis point volatility of the forward rate, and PFV represent the Black-Scholes volatility of the forward rate. In order to compute the Black-Scholes volatility on each date, we divide the basis point volatility by the forward rate on that date.

$$PFV = \frac{BPVol}{F} \tag{1}$$

Since our analysis focuses on at-the-money forward options, we set the strike of the option, ***X***, equal to the forward rate.

$$X = F \tag{2}$$

Next, using the formulas stated below we compute d_1 and d_2 and hence N(d_1) and N(d_2), respectively, where N(*d*) is the standard normal distribution function.

$$d_1 = \frac{\ln(F / X) + PFV^2 (T/2)}{PFV \sqrt{T}} \tag{3}$$

$$d_2 = \frac{\ln(F / X) - PFV^2 (T/2)}{PFV \sqrt{T}} = d_1 - PFV\sqrt{T} \tag{4}$$

Finally, we calculate the option premium, ***P***, using the Black-Scholes model for pricing a call. The standard Black-Scholes model using these inputs is defined as $P = \left[F * N(d_1) - X * N(d_2)\right]$.[9] However, in the swaps market, F and ***X*** are rates rather than prices. Thus, in the swaps world, we adopt the following option price formula, suggested by Hull (2006).

$$P = \$100M \times A \times \left[F \ * N(d_1) - X * N(d_2)\right] \tag{5}$$

where A is given by

$$A = \frac{1}{2}\sum_{j=1}^{4} df_j \tag{6}$$

and where df_j is the discount factor corresponding to period ***j***.

As defined, *A* is the sum of all the discount factors divided by the number of times the swap's cash flows are compounded per year. Since cash flows are received semi-annually over the life of the two year swap, we sum the four discount factors corresponding to the payment dates that occur, in the case of a three month option, 9m, 15m, 21m, and 27m from swaption initiation. Intuitively, *A* is the parameter that converts the interest rate information into a price.

After creating the swaptions pricing model, the value of the swap at expiration must also be calculated. Since our model assumes that the option holder is short a straddle position, the call (put) is in-the-money if the swap rate at expiration is greater (less than) then the strike rate on the option. Since the market swap rate at expiration

---

[9] While a standard Black-Scholes model only uses five inputs to price a swaption—the current market swap rate (spot), the strike, time to maturity, the risk-free interest rate and volatility—it is clear that the bank's proprietary *ad hoc* version of Black-Scholes incorporates other variables into its calculation. For example, while we were not privy to the specifics of the model (only its outputs), we observe that the bank is tweaking the traditional Black-Scholes model by taking into account the term structure of volatility—the fact that implied volatility differs for related options with different maturities.

must be either above or below the strike it is apparent that only one of the options—either the call or the put—can be in-the-money (ITM) at expiration. As such, we only need to price the swap underlying the in-the-money option.[10]

Since the value of any swap is proportional to the spread between the forward swap rate and the option's strike we define the value of the ITM swap, ***ITMSwap***, explicitly as[11]

$$ITMSwap = \left| F - X \right| \times A \tag{7}$$

Finally we calculate the profit & loss at expiration, PnL_E, from the perspective of the investment bank, the straddle seller. We first calculate the future value of the premium, FV_E as

$$FV_E = Pe^{(r_1 t)} \tag{8}$$

Where r_1 is one month LIBOR on the swaption initiation date.

Combining (7) and (8), PnL_E is given by

$$PnL_E = FV_E - ITMSwap \tag{9}$$

In order to analyze an alternate trading strategy, we re-price all of the options—using the formulas outlined above—with one month left until expiration, assuming that the investment bank might *retire* (buy back) the options before expiration. Additionally, now both the put and call premiums, ***PutP & CallP***, must be calculated since both may be non-zero. With this in mind we recalculate the profit & loss with one month left to expiration, PnL_{1MO}, as

[10] The swap underlying the out-of-the-money option has a negative NPV and would never be exercised.
[11] The absolute value function here allows us to price whichever option—the call or the put—that will be exercised, without determining *ex ante* which option will be in-the-money.

$$PnL_{1MO} = FV_R - PutP - CallP \tag{10}$$

2. Forecasting Analysis

(i) Relationship between Proprietary Forecasted and Realized Volatility

Our second avenue of research explores the relationship between proprietary forecasted and realized volatility. We first compute the actual realized volatilities of the forward rates over the life of the option. We then compare these realized volatilities to the proprietary forecasted volatilities from our model in order to determine how effectively the bank's proprietary pricing model forecasts realized volatility. For example, we calculate realized 3 month variance of the forward rate from date t to t+60, V_{3t} as:

$$V_{3t} = \frac{1}{n-1} \times \sum_{i=t}^{t+60} U_i^{\,2} - \frac{1}{n(n-1)} \times \left(\sum_{i=t}^{t+60} U_i \right)^2 \tag{11}$$

where n is the number of days during the life of the option, in this case 60 days and

$$U_i = \ln\left(\frac{F_i}{F_{i-1}} \right) \tag{12}$$

where F_i is the forward rate on day i.

Next, we take the square root of the 3 month variance to obtain 3 month volatility, and multiply that result by the square root of the number of trading days in one year to obtain 3 month volatility on date t, thus:[12]

$$RV_{3t} = \sqrt{V_{3t}} \times \sqrt{250} \tag{13}$$

We calculate 1 month and 6 month realized volatility in a similar manner.

[12] We use 250 here because there are approximately two hundred and fifty trading days in the year, and thus the daily volatility observations we annualize can only be made 250 times. We use a different number of days in our calculation of the future value of the premium because interest is compounded continually (and not just on trading days).

If the swaptions market is efficient, then proprietary forecasted volatility should be an effective estimator of future realized volatility. In addition, we expect proprietary forecasted volatility to be less volatile than realized volatility.[13]

(ii) Auto-Regression (AR) Forecasting Model

Since volatility is known to be path-dependent one econometric model we employ is a standard auto-regression with one lag term (Engle, 1982). This model regresses future realized volatility on past realized volatility for each option maturity to test whether past volatility contains significant information about future realized volatility. To make this more concrete: for the one month option, we test whether realized volatility from the previous thirty days is an effective predictor of realized volatility over the next thirty days. The functional form of the standard AR(1) model is as follows, where ε_t is an error term (Engel, 1982):

$$RV_{kt} = a + b(RV_{k,(t-n)}) + \varepsilon_t \qquad (14)$$

Where *k* equals 1m, 3m or 6m and *n* represents the number of business days associated with the respective time periods. It should be noted that this model uses overlapping data in that the forecast horizon is longer than the observation interval. However, we use Newey-West standard errors to control for the serial correlation that results from using overlapping data.

[13]Christensen and Prabhala (1998) reiterate the common notion that implied volatility is a smoothed expectation of realized volatility.

(iii) Generalized Auto-Regressive Conditional Heteroskedasticity Model

The second model we test is the GARCH(1,1) model which expands the AR(1) model so that future volatility depends on both past squared volatility and past squared returns. The GARCH(1,1) model also controls for the time-varying heteroskedasticity of the error term in the above regression. The model regresses squared future realized volatility, $GV_t^{\,2}$, on squared past volatility, $(GV_{t-1})^2$, and squared past returns, $(u_{t-1})^2$. GARCH models are estimated by the method of maximum likelihood. The functional form of the GARCH(1,1) model, first postulated by Bollerslev (1986) is:

$$GV_t^{\,2} = \omega + \alpha(u_{t-1})^2 + \beta(GV_{t-1})^2 \tag{15}$$

The GARCH model requires that the returns on the underlying asset (in this case the swap), be serially uncorrelated. While this is a reasonable assumption in equity markets, it is violated in the swaps market, where returns are significantly auto-correlated. Consequently, when this model is applied to the swaps market the estimated coefficients on both explanatory variables will be biased. As a result, in order to avoid intentionally introducing a bias in our regression, we do not estimate the GARCH coefficients using the standard maximum likelihood approach. Instead, we use empirically specified coefficients found in relevant literature, such as Anderson, Bollerslev, Christoffersen and Diebold (2006). These coefficients are $\omega = 0, \alpha = .06,$ and $\beta = .94$

V. Results

I. Trading Strategies Analysis

We analyze the trading strategies discussed above on three swaptions—one, three and six month options on two year swaps. For completeness we also tested options on five and ten year swaps, but since the results were similar to our findings on the two year swap we have omitted the results for brevity. **Table 2** includes profit and loss statistics for these trading strategies.

Table 2:

Profit and Loss Statistics on Various Swaptions
Time Period: 01/04/1999 - 01/19/2007

	1m2y	3m2y		6m2y	
	Expiration	1mo to Expiration	Expiration	1mo to Expiration	Expiration
Number Profitable	573	409	731	316	632
Number Not Profitable	507	554	706	611	759
% ITM of Total	53%	42%	51%	34%	45%
Mean	-0.03	-0.20	-0.13	-0.48	-0.41
Std	0.35	0.51	0.72	0.95	1.16
Min	-1.86	-2.31	-2.90	-3.64	-3.90
Q1	-0.21	-0.45	-0.39	-0.95	-0.99
Median	0.02	-0.16	0.01	-0.34	-0.10
Q3	0.23	0.22	0.35	0.23	0.43
Max	0.68	0.63	1.22	1.03	1.39

Table 2 clearly shows that none of the trading strategies are profitable on average. The means of the trading strategies' payouts are skewed by large negative observations—in every case, the minimum payoff is at least twice as large in magnitude as the maximum payoff over the sample period. This is consistent with the fact that selling options offers bounded upside potential with unlimited downside risk.[14] As a result, the median payout in each trading strategy is more profitable than the mean, and in the case

[14] An option sellers' maximum profit is the premium they originally made from selling the options.

of letting one and three month options expire the median is actually marginally positive. Thus, holding the short positions to expiration may result in a higher mean profitability, in absolute terms, but this increase in the mean comes at the cost of a higher standard deviation.

From this initial analysis it appears that selling volatility in the interest rate swaptions market is not a consistently profitable trading strategy in any of these cases. Our results are consistent with Christensen and Prabhala's (1998) findings for the foreign exchange market—that there is no risk premium associated with interest rate swaptions. If there was a risk premium—and proprietary forecasted volatility consistently over estimated realized volatility—then selling volatility would be profitable over the long run.

Table 3:

Volatility Statistics on Various Swaptions
Time Period: 01/04/1999 - 01/19/2007

	1m2y		**3m2y**		**6m2y**	
	Proprietary Forecasted Volatility	Realized Volatility	Proprietary Forecasted Volatility	Realized Volatility	Proprietary Forecasted Volatility	Realized Volatility
Mean	28.35%	28.15%	28.01%	28.66%	27.21%	28.88%
Std	16.27%	17.90%	15.22%	17.05%	13.82%	16.65%
Min	10.25%	5.96%	11.09%	8.06%	11.79%	9.38%
Q1	14.05%	13.81%	14.76%	14.30%	15.20%	14.66%
Median	21.19%	20.35%	21.35%	21.44%	21.44%	21.14%
Q3	41.69%	41.53%	39.60%	44.55%	36.46%	42.21%
Max	69.54%	91.07%	63.04%	74.91%	63.77%	68.34%

From **Table 3** we find that proprietary forecasted volatility and realized volatility are similar in magnitude on average across all three maturities. However, we also find that the volatility of the shorter dated options is itself more volatile than the volatility of longer dated options—manifested through a higher standard deviation and a wider spread between the minimum and maximum values. In addition, by looking at the same two metrics we see that realized volatility tends to be more volatile than proprietary forecasted volatility over the sample period. This fact further reinforces our belief that no risk premium exists in this market, because if one did exist we would not expect proprietary forecasted volatility to be a smoothed expectation of realized volatility (Christensen and Prabhala, 1998). These findings are illustrated in the graphs below.

Figure 2:

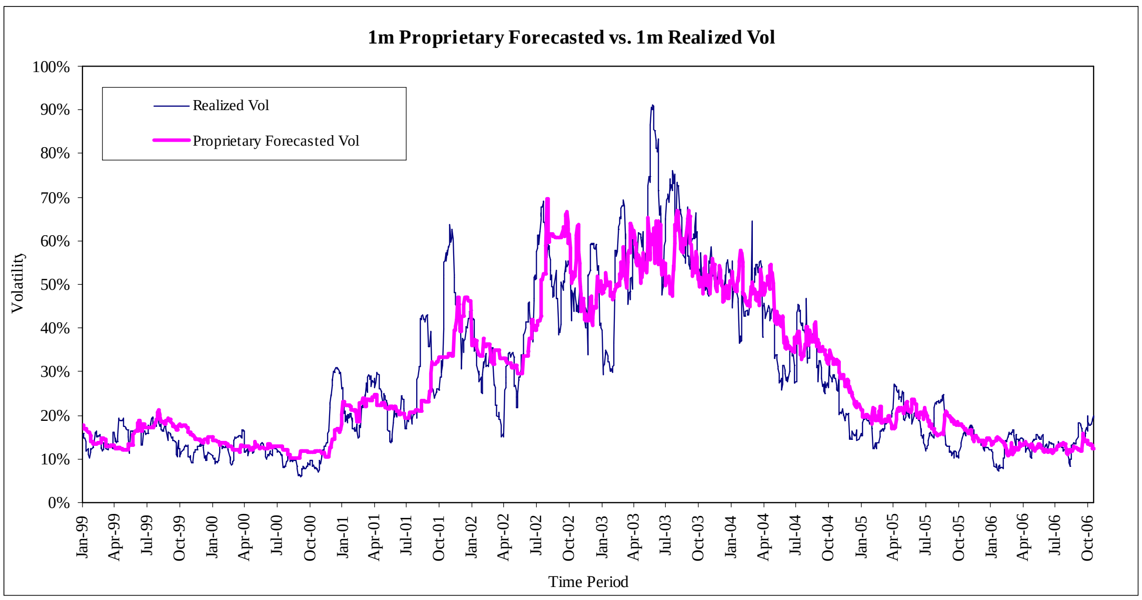

Figure 3:

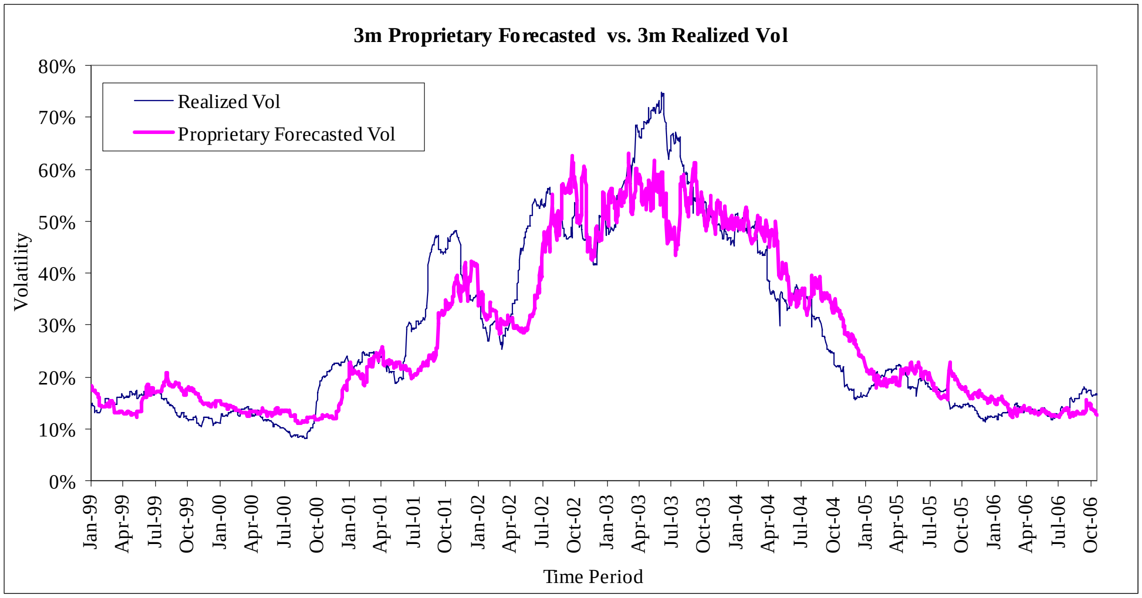

Figure 4:

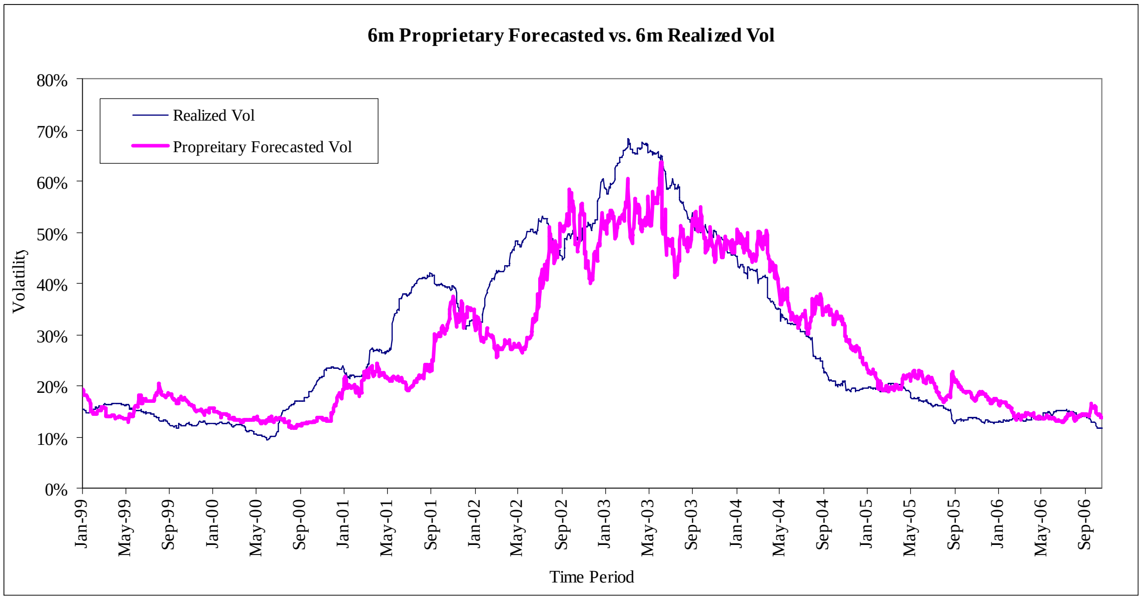

However, looking at average payouts and volatilities does not tell the entire story. Below we report graphs of the profitability of the different trading strategies across the entire time period.

Figure 5:

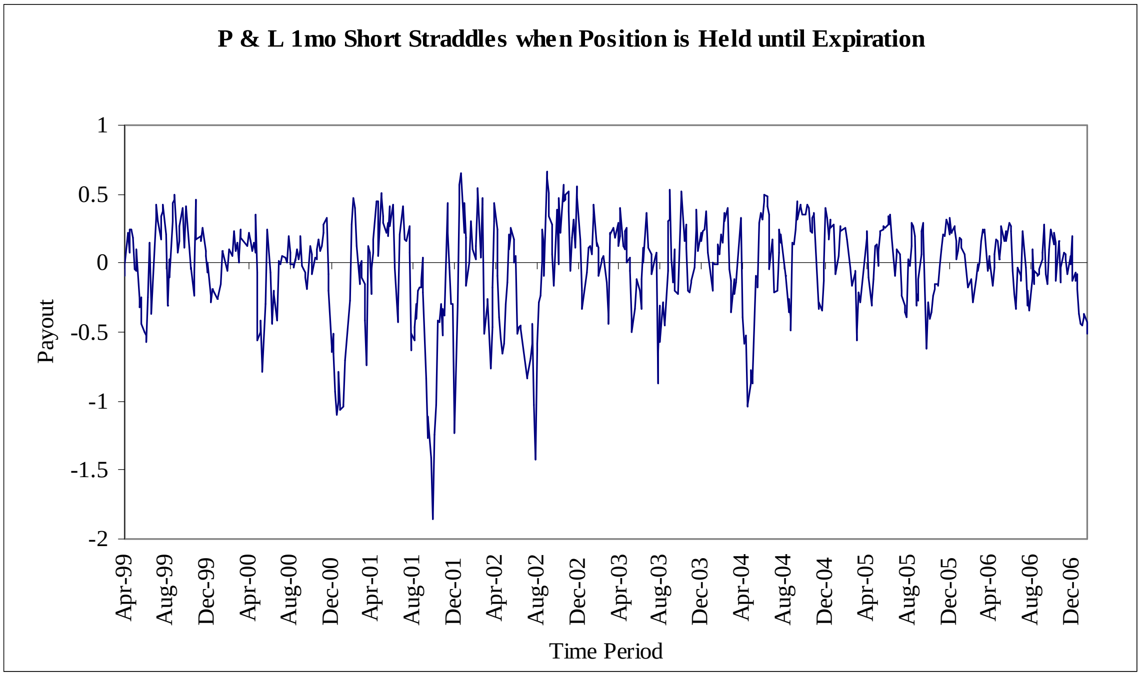

Figure 6:

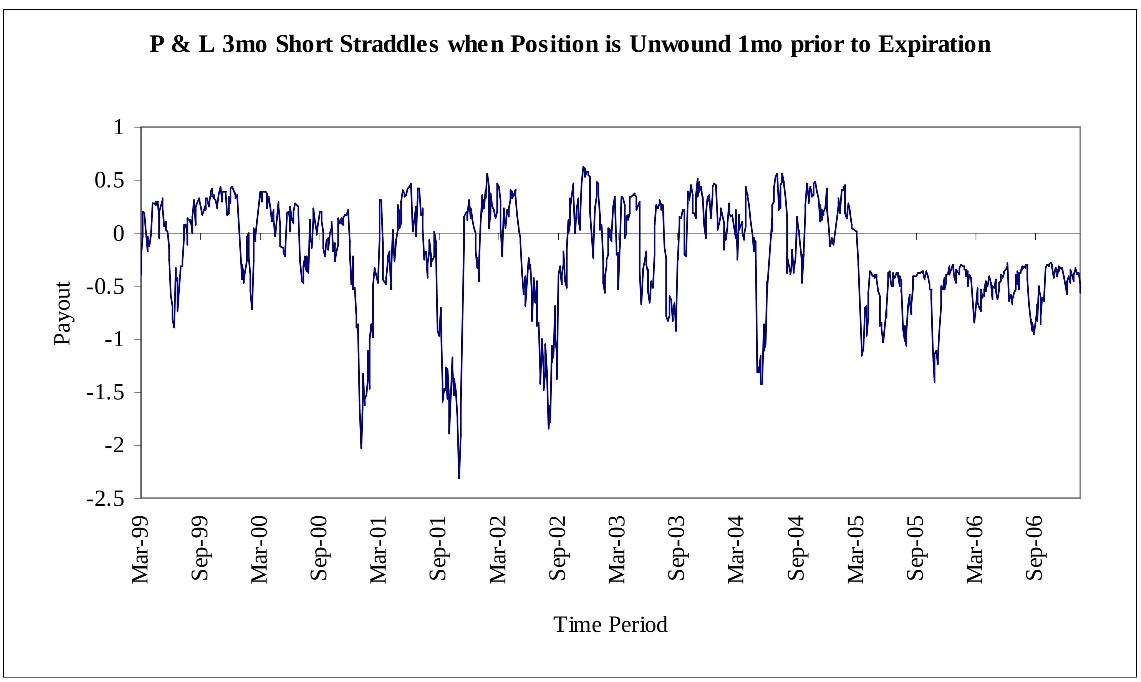

Figure 7:

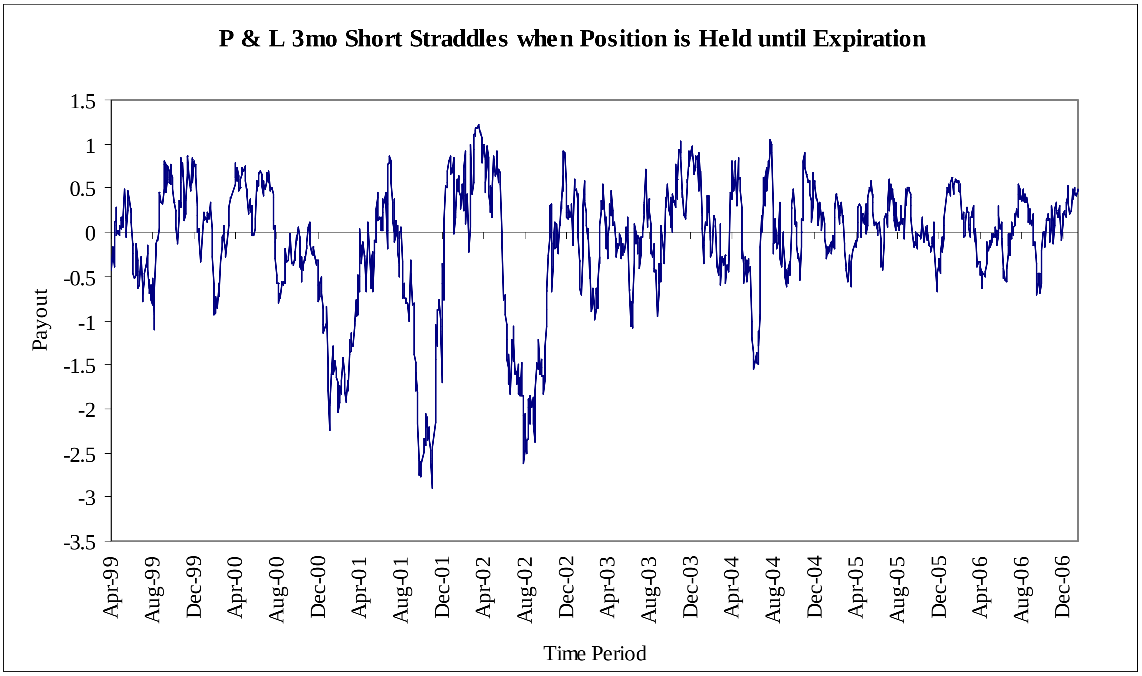

Figure 8:

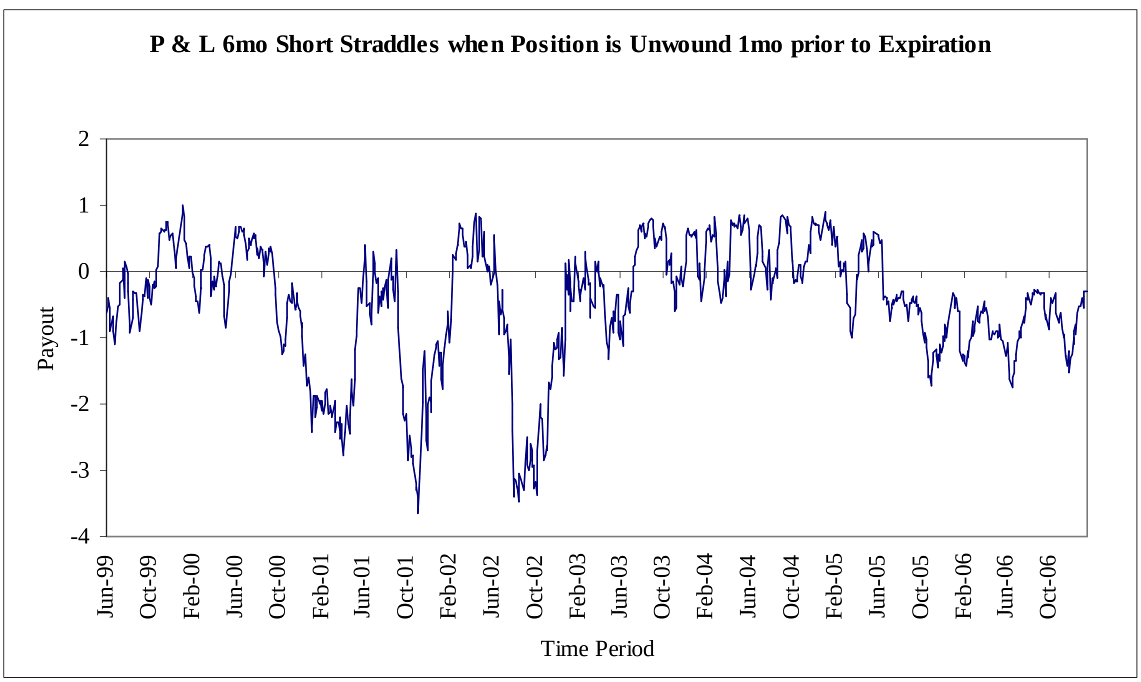

Figure 9:

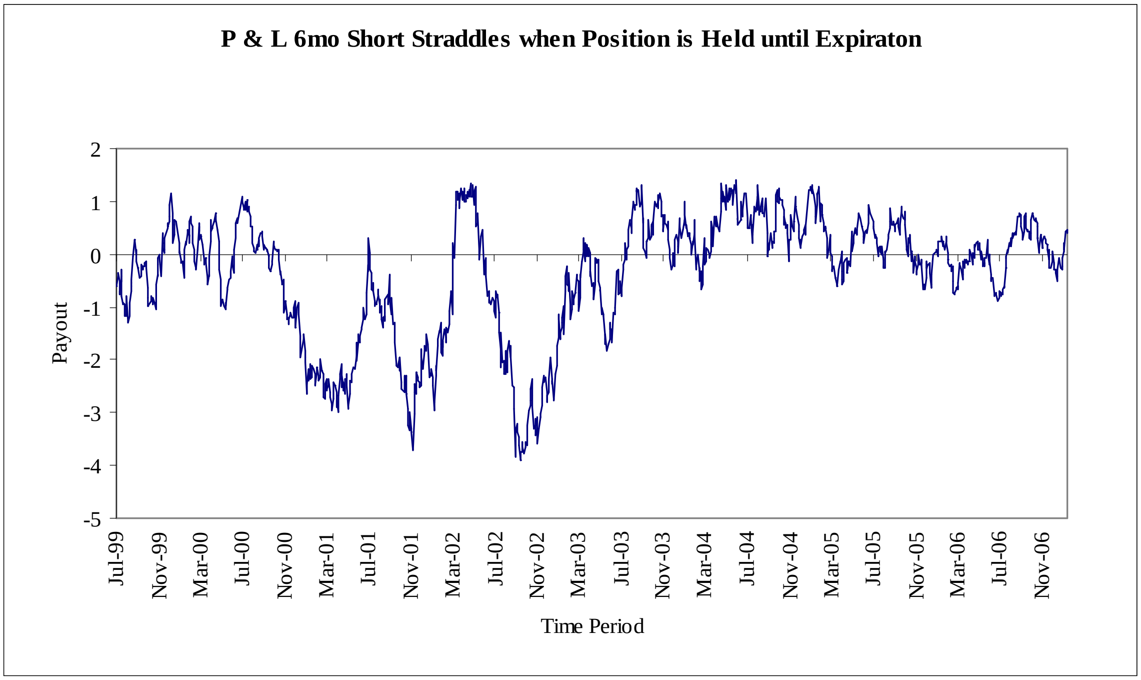

In **Figures 8 & 9** we see that, from July 2000 to July 2002, both 6mo short straddle strategies result in consistently large negative payouts. Since this period corresponds to the dot com bubble we decided to examine this interval in more detail. Our results show that during this period proprietary forecasted volatility underestimates realized volatility. This discovery leads us to believe that there may be a fundamental breakdown in the market's ability to accurately forecast volatility during periods where the economy and the stock markets undergo a substantial correction. We examine this hypothesis in more detail in the next section.

II. Relationship between Proprietary Forecasted Volatility and Realized Volatility

In the previous section, we hypothesized that the market is effective at forecasting volatility except during periods of structural uncertainty and high volatility. We now test this hypothesis by regressing realized volatility (RV) on proprietary forecasted volatility

(PFV) over the entire period, and over the high volatility sub-period, in order to test how well proprietary forecasted volatility's predicts future realized volatility. For example, using 3mo volatility we estimate

$$RV_{3t} = a + b(PFV_{3t}) + \varepsilon_t . \qquad (16)$$

The results of these regressions, taking into account heteroskedastic error terms, are reported below.

Table 4:

$RV_t = a + b * PFV_t$

(1/4/1999-1/19/2007)

# of Obs.	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
1934	1 Month	(a)	-0.011	0.003	-3.88	0.000	0.8217
		(b)	1.006	0.013	77.75	0.000	
1934	3 Months	(a)	-0.007	0.003	-2.74	0.006	0.8507
		(b)	1.043	0.011	97.03	0.000	
1934	6 Months	(a)	-0.001	0.003	-0.32	0.746	0.8102
		(b)	1.057	0.011	98.55	0.000	

As we can see from the regressions in **Table 4**, proprietary forecasted volatility is extremely highly correlated with realized volatility—thus, proprietary forecasted volatility is a good predictor of realized volatility. In each of the regressions the T-Statistic on the proprietary forecasted volatility coefficient is extremely large, which indicates that proprietary forecasted volatility is significantly different from zero at the one percent confidence level. In other words, this confirms our belief that proprietary forecasted volatility contains significant information about realized volatility. Furthermore, we see from the R-Squared statistic that across all three regressions proprietary forecasted volatility correctly predicts between 80-86% of the variation in realized volatility.

It is also interesting to note that the coefficients on the forecasted volatility regressors in all the cases are extremely close to one, because the magnitude of proprietary forecasted volatility is very close to the magnitude of realized volatility. The fact that the coefficients on the 3 month and 6 month proprietary forecasted volatility are greater than one suggests that realized volatility is slightly higher than forecasted volatility in this case because the constant term is indistinguishable from zero. However, even though the coefficients are greater than one for the longer dated options, there is unlikely to be a risk premium large enough to generate arbitrage profits—because the coefficient is so close to one transaction costs would likely wipe out any possible profits. Finally, it is important to note that realized volatility may have a causal effect on proprietary forecasted volatility as market makers continue to use past realized volatility as one of the inputs in their volatility forecasting models

We then re-run these regressions using only the data during the time period from July 3, 2000 to July 31, 2002 in order to test our hypothesis that the market's pricing mechanism is breaking down during this time.

Table 5:

$RV_t = a + b * PFV_t$

(7/3/2000-7/31/2002)

# of Obs.	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
502	1 Month	(a)	-0.020	0.008	-2.49	0.013	0.6518
		(b)	1.139	0.037	30.97	0.000	
502	3 Months	(a)	0.017	0.008	2.23	0.027	0.6832
		(b)	1.112	0.031	35.99	0.000	
502	6 Months	(a)	0.077	0.007	10.99	0.000	0.6090
		(b)	1.017	0.032	31.73	0.000	

The regressions in **Table 5** confirm that the market's ability to predict actual realized volatility appears to break down somewhat during periods of above average volatility. Although proprietary forecasted volatility remains statistically significant across all the different maturity options, the predictive power of all three regressions drops by approximately 18% during this time period—that is, during this period of higher volatility, proprietary forecasted volatility explains about 18% less of the variation in realized volatility.

III. Relationship between Econometric Forecasting Models and Realized Volatility

The final area of our research is concerned with whether market makers are efficiently forecasting future realized volatility. In order to test this hypothesis we look at other volatility forecasting models, specifically AR(1) and GARCH(1,1), to see whether they are more effective at forecasting future realized volatility in periods of both high and low volatility.

The results of the AR(1) model, as defined in equation (14), which uses lagged realized volatility to predict future volatility, over both the entire sample period and the high volatility sub-period, are reported below in **Table 6 & Table 7**.

Table 6:

$RV_t = a + b * RV_{t-n}$

(1/4/1999-1/19/2007)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
1080	1 Month	(a)	0.042	0.004	9.78	0.000	0.7119
		(b)	0.844	0.018	45.82	0.000	
1437	3 Months	(a)	0.033	0.003	9.56	0.000	0.7604
		(b)	0.879	0.013	69.44	0.000	
1019	6 Months	(a)	0.042	0.005	8.88	0.000	0.6868
		(b)	0.847	0.018	46.62	0.000	

Table 7:

$RV_t = a + b * RV_{t-n}$

(7/3/2000-7/31/2002)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
281	1 Month	(a)	0.104	0.013	8.17	0.000	0.3699
		(b)	0.680	0.055	12.45	0.000	
385	3 Months	(a)	0.125	0.011	11.21	0.000	0.3532
		(b)	0.684	0.046	14.74	0.000	
284	6 Months	(a)	0.115	0.009	13.17	0.000	0.5957
		(b)	0.847	0.041	20.55	0.000	

Looking first at **Table 6**, all the coefficients on the lag terms are statistically different from zero at the one percent confidence level. We also find that one lag of past realized volatility explains between 68% and 77% of the variance in future realized volatility. While this R^2 is extremely high by conventional standards, it is still substantially lower then the R^2 reported in the regressions in **Table 4** across all maturity options.

In addition, when looking at **Table 7** we see that the coefficients on the lag terms are still statistically significant from zero during the high volatility sub-period. However, the R^2 during this sub-period are considerably lower, ranging from 35% to 60%. When the AR(1) model is compared to the regressions in **Table 5** we see that proprietary forecasted volatility even further outperforms lagged volatility during time periods of high volatility—illustrated by a much larger R^2.

It appears from these results that proprietary forecasted volatility contains information above and beyond the information contained in past realized volatility in all periods. In order to test this hypothesis we regress future realized volatility on both a lag of realized volatility and proprietary forecasted volatility. This regression is reported below.

Table 8:

$RV_t = a + b * RV_{t-n} + c * PFV_t$

(1/4/1999-1/19/2007)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
1080	1 Month	(a)	-0.010	0.004	-2.53	0.012	0.8192
		(b)	-0.081	0.055	-1.47	0.141	
		(c)	1.081	0.059	18.29	0.000	
1437	3 Months	(a)	-0.008	0.003	-2.59	0.010	0.8469
		(b)	-0.020	0.048	-0.43	0.669	
		(c)	1.067	0.052	20.6	0.000	
1019	6 Months	(a)	-0.002	0.004	-0.42	0.673	0.8066
		(b)	-0.037	0.038	-0.98	0.327	
		(c)	1.098	0.045	24.47	0.000	

These regression results confirm our previous hypothesis. In this regression we find that when we add proprietary forecasted volatility to the AR(1) model the R^2 increases by only ~10% to 85%. However, this is the same R^2 that results if only proprietary forecasted volatility is included in the model, as it is in the regressions in **Table 4,** suggesting that the lag term is not improving the model's forecasting ability. In addition, when we include both of these explanatory variables we find that only the coefficient on proprietary forecasted volatility is statistically significant. While we recognize that the regressors are themselves highly correlated, and as a result our regressions could potentially suffer from multicollinearity issues, the results of the regression are statistically significant enough to conclude that proprietary forecasted volatility subsumes all the information contained in a lag of realized volatility.

We continue with our analysis by also testing the GARCH(1,1) model's ability to forecast future realized volatility. The results of the GARCH(1,1) model, as defined in equation (15), with empirically determined coefficients, as specified in **Section IV Part 3,** are shown below.

Table 9:

$RV_t = a + b * GV_t$

(1/4/1999-1/19/2007)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
1856	1 Month	(a)	0.070	0.006	11.79	0.000	0.5298
		(b)	0.747	0.025	29.37	0.000	
1816	3 Months	(a)	0.126	0.005	23.29	0.000	0.3782
		(b)	0.582	0.019	30.85	0.000	
1756	6 Months	(a)	0.111	0.005	21.56	0.000	0.3491
		(b)	0.620	0.015	40.13	0.000	

First, we observe that the coefficients on all of the GARCH terms are statistically different from zero at the one percent confidence level. However, we see that although the coefficients on the explanatory variables are significant, they are not close in magnitude to one, as they are in the other forecasting models. Furthermore, from these regressions we find that the GARCH(1,1) model underperforms both other forecasting models in terms of R^2—using the GARCH forecast as a regressor explains only 34% to 53% of the variation in future realized volatility.

Finally, we add the GARCH forecasted volatility into the regression from **Table 8** to determine which of the three forecasting models most effectively predicts future realized volatility. **Table 10** below shows the results of these regressions:

Table 10:

$RV_t = a + b * RV_{t-n} + c * PFV_t + d * GV_t$

(1/4/1999-1/19/2007)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
1044	1 Month	(a)	-0.011	0.004	-2.72	0.007	0.8185
		(b)	-0.113	0.056	-2.01	0.045	
		(c)	1.052	0.060	17.55	0.000	
		(d)	0.070	0.031	2.28	0.023	
1391	3 Months	(a)	-0.003	0.003	-0.86	0.391	0.8460
		(b)	-0.004	0.049	-0.07	0.942	
		(c)	1.048	0.051	20.44	0.000	
		(d)	-0.011	0.013	-0.85	0.393	
985	6 Months	(a)	0.010	0.004	2.34	0.020	0.8070
		(b)	-0.050	0.037	-1.34	0.179	
		(c)	1.122	0.051	22.2	0.000	
		(d)	-0.038	0.020	-1.89	0.059	

By comparing the R^2 of these regressions with those in **Table 4** we see that adding the realized volatility lag term and the GARCH forecast term did not improve the R^2 of the regression. Furthermore we see that, in the regressions for the three and six month options, the coefficients on the lag term and the GARCH forecast term are statistically insignificant from zero at the one percent confidence level—only for the one month option are all three regressors significant, and only at the five percent confidence level. Again, while we recognize the potential multicollinearity issues arising from using regressors that are themselves highly correlated it is beyond the scope of this paper to employ the econometric techniques required to rectify this issue.

Additionally, we re-run the regression in **Table 10** using only the high volatility sub-period. The results of this regression are shown below.

Table 11:

$RV_t = a + b * RV_{t-n} + c * PFV_t + d * GV_t$

(7/3/2000-7/31/2002)

# of Obs	Time Period	Coef.	Coef. Value	Std. Err.	T-Statistic	Pr(P>|t|)	R-Squared
281	1 Month	(a)	-0.025	0.010	-2.44	0.015	0.6749
		(b)	-0.490	0.077	-6.36	0.000	
		(c)	1.605	0.085	18.89	0.000	
		(d)	0.038	0.076	0.5	0.616	
385	3 Months	(a)	0.012	0.009	1.37	0.170	0.7712
		(b)	-0.691	0.058	-11.87	0.000	
		(c)	1.953	0.096	20.4	0.000	
		(d)	-0.131	0.032	-4.03	0.000	
284	6 Months	(a)	0.098	0.011	9.13	0.000	0.7259
		(b)	0.408	0.046	8.96	0.000	
		(c)	0.795	0.052	15.17	0.000	
		(d)	-0.239	0.024	-9.91	0.000	

In this high volatility sub-sample we find that the coefficients on the variables in this regression, with the exception of one month GARCH volatility, are all statistically significant at the one percent confidence level. However, the magnitudes of the coefficients are not meaningful in this regression because of aforementioned multicollinearity issues. Finally, we note that, regardless of this multicolleanearity issue, proprietary forecasted volatility is still the most powerful explanatory variable in these three regressions. The coefficients on this variable are the most strongly statistically significant in all three regressions, and moreover, single variable regressions using only proprietary forecasted volatility as an explanatory variable have much higher R^2 values then the regressions that use only past or GARCH volatility.

The results of these regressions lead us to conclude that the Investment Bank's proprietary forecasted volatility is a better predictor of realized volatility than either an AR(1) or simplified GARCH (1,1) model during our sample period across all maturity options—it contains information above and beyond that contained in past volatility or GARCH forecasted volatility. We also find that during the high volatility sub-period the bank's proprietary forecasted volatility does not subsume all the information contained in past realized volatility or GARCH forecasted volatility, but is still the most effective estimator of future realized volatility. Intuitively, this suggests that the market is efficiently pricing volatility in the swaptions market.

VI. Conclusion

Our finding, that no risk premium exists in the interest rate swaptions market, implies that short gamma trading strategies are not profitable in this market. This finding runs counter to what other economists have found for other options markets. For example, Heston and Nandi (2000) find that a risk premium exists in the market for S&P 500 index options.

In addition, our conclusion that the swaptions market proprietary forecasted volatility is an effective estimator of future realized volatility, except during periods of structural uncertainty and high volatility is consistent with what Christensen and Prabhala (1998) find for options on the S&P 100 Index as well as what Busch, Christensen and Nielsen (2007) find for options on foreign exchange markets.

Our results, that proprietary forecasted volatility outperforms both past realized volatility and GARCH forecasted volatility in high and low volatility periods, are significantly different from other literature. For example, our findings run counter to Day and Lewis (1992) and Lamoureux and Lastrapes (1993) who find that past volatility actually outperforms proprietary forecasted volatility in forecasting future realized volatility. Similarly, our findings diverge from the work of Heston and Nandi (2000), who find that the GARCH forecast estimates outperform both past and implied volatility forecasts. Regardless, our research is the first paper to examine volatility forecasting specifically in the swaptions market and, as such, should serve as a jumping off point for further research in this arena.

References

Anderson, T.G., Bollerslev, T., Christoffersen, P.F., and Diebold F.X. (2006). Volatility and Correlation Forecasting. In G. Elliot, C.W.J. Granger and A.Timmerman (Eds.) *Handbook of Economic Forecasting* (pp. 778-878). Amsterdam: Elsevier B.V.

Bates, D. (2000). Post-'87 crash fears in S&P 500 futures options. *Journal of Econometrics, 94,* 181-238.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31,* 307-327.

Black, F., Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy, 81,* 637-654.

Busch, T., Christensen, B.J., Nielsen, M.O. (2006). The role of implied volatility in forecasting future realized volatility and jumps in the foreign exchange market, stock, and bond markets. *Working Paper Series No. 223*. University of Aarhus.

Christensen, B.J., Prabhala, N.R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics, 50,* 125-150.

Day, T., Lewis, C. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics, 52,* 267-287.

Engle, R.F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. *Econometrica, 50,* 157-168.

Fan, R., Gupta, A., Ritchken, P. (2006). On pricing and hedging in the swaptions market: How many factors, really?. *Working Paper.* Case Western Reserve University.

Hull, J. (2006). *Options, futures, and other derivatives.* New Jersey: Pearson/Prentice Hall.

Lamoureux, C.G., Lastrapes, W. (1993). Forecasting stock return variance; towards understanding stochastic implied volatility. *Review of Financial Studies, 6,* 293-326.