Variance Swaps in the Presence of Jumps

Max Schotsman

July 1, 2013

Abstract

This paper analyses the proposed alternative of the variance swap, the simple variance swap. Its main advantage would be the insensitivity to jumps in the price of the underlying asset, making it possible to hedge even when that would not be possible using the normal variance swap. The data used for this will be that of the S&P 500 index. Both the SVIX and VIX will be calculated, the SVIX being an index made from the simple variance swap in the same way the VIX is made from the normal variance swap. Using regression analysis we find that the SVIX hold significant information over the return of the next period, and information that differs from the VIX. It is also determined that the dataset does not always contain enough observations, leaving us with only part of the intended range.

A variance swap is a financial derivative on the volatility of an asset. It can be used to hedge risk or speculate on the movement of the asset. Each variance swap has a strike price, the variance strike. The value of the variance swap is the difference between the realized variance and the variance strike, discounted until the present. The convention is to set the strike equal to the expected variance, so that no money switches hands at the time of purchase. A volatility index is an index of which the level under corresponds to the strike on a variance swap. The most well known of these indices is the VIX[1]. The definition of the VIX would mean that at any point in time the value of the VIX is equal to the expected variance on the S&P 500 in the next 30 days. This interpretation of the VIX is only valid if certain assumptions hold. A detailed definition of the assumptions will be given further on.

Chief subject of this paper will be the effect of price jumps in the price of the underlying asset. If there are price jumps in the price of the underlying asset the variance swap cannot properly be used for hedging. The fact that an asset is not allowed to have jumps is a very strict assumption, certainly since the financial crisis of 2008, which introduced an abundance of jumps in stockmarke. An example of problems that jumps hold for variance swaps is if a company goes bankrupt. The S_t would at some point t be zero, leading to an infinite payoff. This can be mitigated by using caps on the payoffs, but the hedging capabilities are broken in this case.

Ian Martin (2013) proposes a different variance swap called the simple variance swap.

[1]The VIX is the Chicago Board Options Exchange Market's volatility index, a measure of implied volatility of S&P 500 index options

This variance swap would still have a valid interpretation of the variance in the presence of jumps. In this paper this proposed simple variance will be analyzed, and its corresponding index will be compared to the normal variance swap's index.
This paper will start with a description of the data. For data this paper will focus on the S&P 500, which is used for the VIX. After that the methods used will be explained. After this the results will be shown and discussed.

1 Data description

To evaluate the VIX there is a need for a large amount of observations. Ideally we would have an infinite number of options, for each conceivable strike and date of maturity. Since this ideal does not exist discretization has to be used. For the pricing to be as accurate as possible the amount of observations has to be very high. This is why the S&P 500 as the underlying asset is chosen, as it contains the most actively traded options of the European type. The sample period of the data is from January 1^{st} 1996 to January 31^{st} 2012. The data was obtained using the optionmetrics database. This dataset contains 3.410.396 observations; both put and call options on the S&P 500. Before the data can be properly used it has to be filtered in several ways. First the observations with missing data are removed. After that options with more than 550 days to maturity are removed, as they are not of consequence for the calculations horizon that will be used. Options with a date of maturity shorter than 7 days were also removed, as their prices are too different from options with other maturity and are therefore not usable for interpolation. After these adjustments 2.489.200 observations remain. For the risk-free rate, r, T-bills with corresponding dates of maturity were used. If there was no corresponding date of maturity one was made through interpolation.

Table 1: Number of options per day

	Call	Put	Both
Min.	86	78	197
Max.	876	968	1742
Average	301,17	313,75	614,92
(Std.)	(152,83)	(222,90)	(365,76)
Total	1.219.122	1.270.078	2.489.200

Table 1 has a description of the data regarding the number of options per day. The number of call and put options is about even, and there are never less than 197 options in the dataset in a single day. This seems to be plenty, but this amount is for all maturities on one day and call options with no corresponding put options have to be discarded and vice versa. This leads to few options on a day that can be used. This issue is discussed further on.

2 Methods

This section details the methods used. For each model different assumptions must hold. The assumptions are :

1. The put and call options on the underlying asset can be traded for any strike.
2. The underlying asset does not pay dividends.
3. The continuously-compounded interest rate is constant.
4. The underlying asset can be traded in real time.
5. The underlying asset's price does not experience jumps.

Assumption one is very strict and hard to find in real life, so this assumption is relaxed to that options are available for a large number of strikes, and the calculations that were to be done using integrals will be performed using Riemann sums. For both the SVIX and the VIX only assumption 1 and 2 need to hold. For the variance swap all 5 assumptions need to hold. For the simple variance swap assumptions 1 through 4 need to hold, but these can all be relaxed except for assumption 2.

2.1 Variance swaps

As said before a variance swap is an agreement to exchange realized variance for an agreed upon strike price at the time of maturity. The strike is chosen so no money needs to change hands at time $t = 0$ or

$$V = E\left[\left(\log\frac{S_\Delta}{S_0}\right)^2 + \left(\log\frac{S_{2\Delta}}{S_\Delta}\right)^2 + \cdots + \left(\log\frac{S_T}{S_{T-\Delta}}\right)^2\right] \tag{1}$$

If all assumptions hold the strike on a variance swap is, with Δ in equation 1 approaching 0:

$$V = 2e^{rT}\left(\int_0^{F_{0,T}} \frac{1}{K^2} P_{put}(K)dK + \int_{F_{0,T}}^{\infty} \frac{1}{K^2} P_{call}(K)dK\right). \tag{2}$$

With T as the date to maturity, $F_{i,j}$ as the forward price form time i to time j, K as the strike price of the option, P the option price and dK the difference between two available strikes. Equation 2 is equivalent to that the strike of the variance swap is equal to the expected realized variance.
To replicate this financial derivative using only standard European call and put options we need to take the following positions:

1. For every available $K \leq F_{0,T}$ buy $2/K^2dK$ puts with a maturity at time T.
2. For every available $K \geq F_{0,T}$ buy $2/K^2dK$ calls with a maturity at time T.

3. A dynamic position in the underlying asset, with the amount, $2(F_{0,t}/S_t - 1)F_{0,t}$, changing over time, financed by borrowing.

Here is a sketch of the proof of this position: In the limit Δ to 0 equation 1 converges to

$$V = E\left[\int_0^T (d\log S_t)^2\right] \tag{3}$$

Since there are no jumps according to assumption 5, by Ito's lemma
$d\log S_t = (r - \frac{1}{2}\sigma_t^2)dt + \sigma_t dZ_t$ under the risk-neutral measure, so $(d\log S_t)^2 = \sigma^2 dt$ and

$$V = E\left[\int_0^T \sigma_t^2 dt\right] \tag{4}$$

$$V = 2E\left[\int_0^T 1/S_t dS_t - \int_0^T d\log S_t\right] \tag{5}$$

$$V = 2rT - 2E\log\frac{S_t}{S_0}. \tag{6}$$

This means that the strike of the variance swap is determined by pricing a contract that pays the log of the underlying asset's simple return. Using a method from Breeden and Litzenberger (1978) this contract can be priced in terms of the European call an put options:

$$P_{log} = e^{-rT}E\log R_t = rTe^{-rT} - \int_0^{F_{0,T}} \frac{1}{K^2}P_{put}(K)dK - \int_{F_{0,T}}^{\infty} \frac{1}{K^2}P_{call}(K)dK \tag{7}$$

When we substitute equation 7 in equation 6 we have the proof.
This result is called model-free since the underlying asset's price can follow any Itô process. The downside is that assumption 5 does not always hold, as there are almost always some jumps in the price of the underlying asset.

2.2 VIX

The VIX is constructed from variance swaps. The valuation of the variance swap and the VIX will be done using the standard procedure[2] of the CBOE. Assumption one does not hold for real life examples, so to calculate the VIX we have to use discrete data. If assumption one holds then VIX equals:

$$VIX^2 = 2e^{rT}\left(\int_0^{F_{0,T}} \frac{1}{K^2}P_{put}(K)dK + \int_{F_{0,T}}^{\infty} \frac{1}{K^2}P_{call}(K)dK\right). \tag{8}$$

[2]See http://www.cboe.com/micro/vix/vixwhite.pdf for a step by step calculation.

The VIX is often used as a proxy for the risk neutral expectation of the variations of the log returns (equation 1), but research[3] has shown that there is often a large gap between this number and the VIX index. If assumption 1 through 5 hold this would not be the case. What the VIX measures can be described as risk neutral entropy of the simple returns on the S&P 500. Entropy is a measure of chaos or disorder of a positive random variable. It is always nonnegative and measures variability by how much the expected value differs from the real value. Under Assumptions 1 and 2 VIX measures

$$VIX^2 = \frac{2}{T}L(R_t) \tag{9}$$

with $L(X) = \log E[X] - E[\log X]$. LOGNORMAL.

2.3 Simple variance swaps

A simple variance swap is similar to the variance swap, but the payoff is calculated by

$$V = \left(\frac{S_\Delta - S_0}{F_{0,0}}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_{0,\Delta}}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{0,T-\Delta}}\right)^2. \tag{10}$$

If assumption one holds, then the forward price, $F_{0,t}$, is the strike for which the put and call price are the same. If the asset pays no dividend then $F_{0,t} = S_0 e^{rt}$.
To replicate this financial derivative using only standard European call and put options we need to take the following positions:

1. For every available $K \leq F_{0,T}$ buy $(2/F_{0,T}^2)dK$ puts with a maturity at time T.
2. For every available $K \geq F_{0,T}$ buy $(2/F_{0,T}^2)dK$ calls with a maturity at time T.
3. A dynamic position in the underlying asset, with the amount, $2e^{-\delta(T-t)}(1 - S_t/F_{0,t})/F_{0,t}$, changing over time.

An obvious advantage of the simple variance swap is that no cap on the payoff needs to be imposed, which is necessary for single asset variance swaps.

2.4 SVIX

Now that we have a way to calculate the simple variance swap, we can make an index similar to the VIX. This index is based on the simple variance swap instead of the variance swap and will be called the SVIX, or simple VIX. The SVIX is calculated by the following formula:

$$SVIX^2 = \frac{2e^{rT}}{TF_{0,T}^2}\left(\int_0^{F_{0,T}} P_{put}(K)dK + \int_{F_{0,T}}^{\infty} P_{call}(K)dK\right). \tag{11}$$

[3] See Aït-Sahalia, Karaman and Mancini (2012).

with both the put and call options purchased at t=0 and expiring at t=T. This is equal to the strike of the simple variance swap.
Under assumptions 1 and 2 the SVIX can be interpreted as

$$SVIX^2 = \frac{1}{T}\text{var}\left(\frac{R_t}{R_{f,T}}\right) \tag{12}$$

Proof of this statement is in appendix 1.

2.5 The effect of discretization

Assumption 1 does not hold in real life: it is not possible to buy options for every strike and date of maturity. To evaluate the VIX without the necessary data this assumption is relaxed. To calculate the area under the hypothetical function Riemann sum is used. If the step size of the Riemann sum would go to zero, we would have a fair approximation of the integral, but the S&P 500's options have a step size of 5. Further analysis concerning the discretization is in appendix 2.

3 Results

To calculate the SVIX the data of the S&P 500 was used. The data was sorted according to: call/put flag, date, expiration date, strike price and option price. Using the same method as the CBOE does for their VIX the forward price was calculated for every call option of every day that had a put option counterpart at the same strike price. With this forward price for different times the SVIX was calculated.

To have a consistent interpretation of this SVIX despite the different lengths of maturity available each day, interpolation was used to create a SVIX with a 30 day forward price, 60 day forward price etc. If interpolation was not possible extrapolation was used.

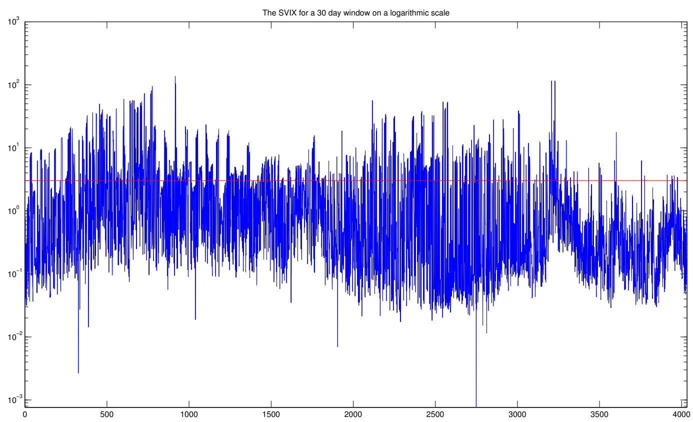

The graph above shows the results of the 30 day expected variance. The horizontal line is the average SVIX of 3,0238. Some large spikes can be noticed at time of great unrest, the most notable being the spike around n=3200, during the financial crisis of 2008. Also around n=1600 where the dotcom bubble had just burst is a period of relatively extreme results. Among the highest peaks is also n=700, when Long-Term Capital Management received its bailout following the Russian financial crisis. There are some extremely low values for the SVIX. These points are the result of having few options available, and the options that were available were far from the $F_{0,T}$, giving a smaller SVIX than should be. There are also more high values than should can be explained with economic distress. This is discussed more in section 3.2.

3.1 What information does the SVIX hold

In this section we will use regression analysis to see if the SVIX holds (extra) information when forecasting the returns of the S&P 500. We will use several variables to predict the returns of the S&P 500. The variables are:

1. R_t, the return at time t
2. R_{t-1}, the return at time t-1
3. C a constant
4. r, the risk-free rate at time t, calculated using T-bills
5. $SVIX_t$ the level of the SVIX at time t
6. VIX_t the level of the VIX at time t.

Two models are used for the regression:
model A

$$R_t = \alpha_1 + \alpha_2 R_{t-1} + \alpha_3 r_{t-1} + \alpha_4 VIX_{t-1} + \alpha_5 SVIX_{t-1} \tag{13}$$

and model B

$$R_t = \alpha_1 + \alpha_2 R_{t-1} + \alpha_3 r_{t-1} + \alpha_4 VIX_{t-1}. \tag{14}$$

Model B is a control model, to see what the results are if the SVIX would not be used. Table 2 shows the covariance/correlation matrix. Above the diagonal the correlation is give, the diagonal is the variance and under the diagonal is the covariance. Important

Table 2: Covariance and correlation of the variables (all data)

	SVIX	VIX	r	R_t
SVIX	68.17892	0.90137	0.15575	-0.02995
VIX	0.06194	0.00007	0.19858	-0.05132
RF_RATE_ST	0.02770	0.00004	0.00046	-0.12434
returns	-0.00330	-0.00001	-0.00004	0.00018

to note in this table is the high correlation between the SVIX and the VIX, which was expected. The large difference in variance between the VIX and the SVIX is due to the fact that they are standardized in different ways, here the VIX is several factors smaller as can be seen in table 3. Table 3 gives a description of the data. Table 3 is divided in two parts, one part from the end of 2007 until January 2012, the other with the entire dataset. This split is also in the regression. The reason for this split is there was not enough data to properly calculate the SVIX or VIX available. This problem is discussed more in depth in the next section.

Table 3: Descriptive statistics of the variables

12/2007-1/2012	SVIX	VIX	r	R_t
Mean	1.374	0.002	1.009	0.001
Std. Dev.	6.835	0.006	0.010	0.018
Skewness	14.168	13.273	1.312	-0.574
Kurtosis	229.277	252.276	3.090	8.871
1/1996-1/2012	SVIX	VIX	r	R_t
Mean	3.024	0.005	1.033	0.003
Std. Dev.	8.257	0.008	0.022	0.013
Skewness	7.273	4.657	-0.122	-0.339
Kurtosis	76.928	42.659	1.464	9.905

Table 4 shows the results of the regression. In brackets are the HAC standard errors. Again there is a divide, one part using data form December 2007 to January 2012, the other using data from January 1996 and on. It is easy to see the problem with the models using the full dataset: the SVIX and VIX are not significant at any reasonable level. The problem earlier on is the lack of sufficient options to properly discretize. For this reason we will only be looking at the results from December 2007 and on. We see at model B that the VIX is insignificant.

The R^2 is comparable to what others have found such as in the paper of Bollerslev, Tauchen and Zhou (2009). In their regression the VIX and realized variance are used to explain returns. They found an R^2 of 0,0107 for variables with a one month horizon and an R^2 of 0,0682 using variables with a three month horizon. For regressing with the SVIX there was no significant gain when using a three month horizon, leading to an R^2 of 0,058 for model A.

The coefficient of the SVIX is positive while the coefficient of the VIX is now negative. Since the VIX and SVIX themselves are always positive, this means that the VIX has a negative contribution to the predicted return of tomorrow, while the SVIX has a positive contribution to the predicted return of the next period. This interpretation is of course shaky, but it is the closest fit of the data the model gives. For a proper interpretation we have to look at the difference between the SVIX and the VIX. The following graph shows the difference between the SVIX and the VIX. For a larger graph see the appendix. For this figure the SVIX and the VIX have been standardized in the same way. N=0 is on 21 December 2007, or in the entire dataset N=0 would correspond to N=3000. At N=4030 the date is the 31st of January 2012. Three interesting spikes remain in the

Table 4: Regression results using the SVIX

	12/2007-1/2012		1/1996-1/2012	
Model	A	B	A	B
C	0.251	0.252	0.087	0.088
	(-0.058)	(0.059)	(0.009)	(0.009)
R_{t-1}	-0.165	-0.153	-0.084	-0.083
	(0.044)	(0.049)	(0.023)	(0.023)
VIX	-0.489	0.086	-0.039	0.061
	0.134	0.119	0.144	0.025
SVIX	0.0007		0.0001	
	0.0001		0.0002	
r	-0.249	-0.251	-0.087	-0.088
	0.058	0.059	0.009	0.009
R-squared	0.061	0.039	0.025	0.024
S.E. of regression	0.018	0.018	0.013	0.013

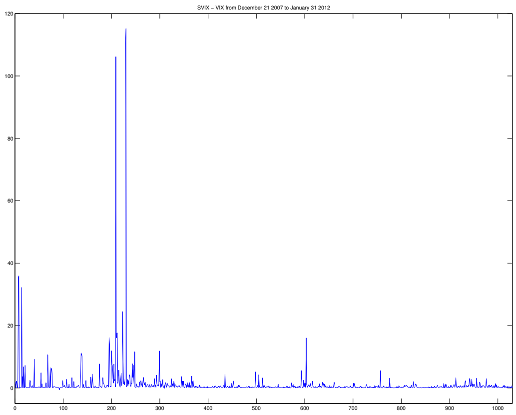

data, the first is where N is between 6 and 13. For these days there were few options available, so the result is biased. In this case the model thinks there are only options of high value and so it overestimates the SVIX. The other spikes are right at the 2008 financial crisis, or around N=200. This means that the SVIX reacts more heavily than the VIX to extreme changes in the index level. This can again be seen at the third spike, at N=600. This was the European debt crisis that is still lingering. Again the SVIX has a more extreme reaction to a crisis.

3.2 The effect of volume

We have talked before about the problem with the number of options. Graph 3 shows the number of options for the smallest maturity each day. The blue dots show call options, the red crosses show put options. The graph shows that there is an enormous increase in the number of options available from the end of 2007 on. The second part of the graph shows why the data before the end of 2007 has not been used, and gives poor regression results. Under the number of options is a plot of the SVIX. This shows a lot

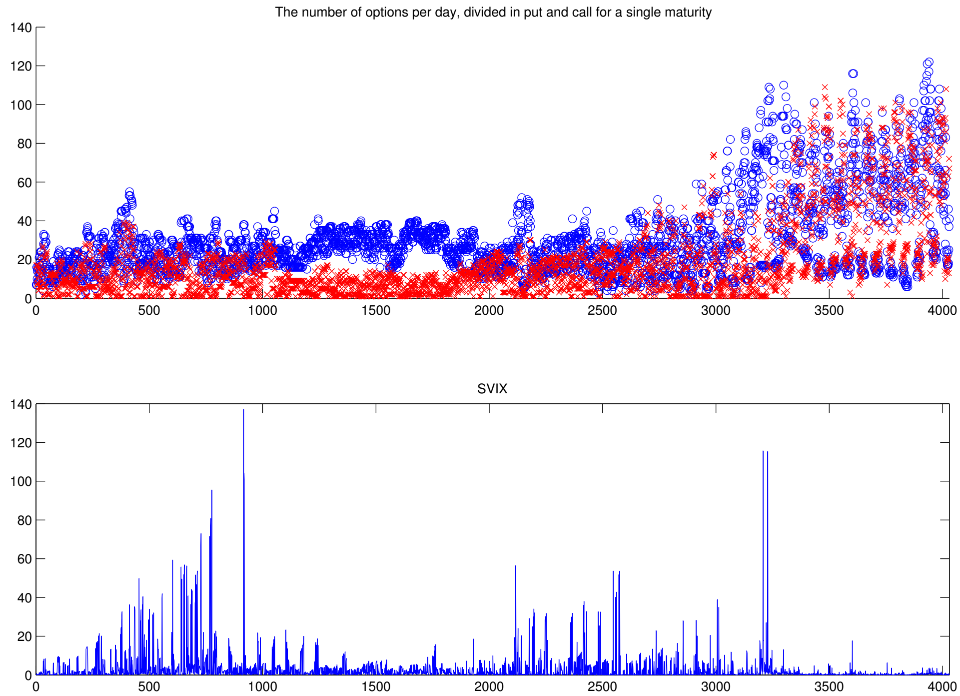

of abnormal results, enormous spikes where they should not be. This can be directly linked to the number of options available. The data is directly received from the optionmetrics database, but before the end of 2007 there are just not a sufficient amount of observations available. We theorized that this lack of data could be compensated for using spline interpolation or a similar interpolation method. The problem however was that the missing options were inconsistent, sometimes it were those that had a low value and sometimes those that had a high value. When interpolating with missing high information the SVIX would be undervalued, when there is missing low price information the SVIX/VIX would be greatly overvalued. This led to similar problems as discretization has, so it offered no solution. Ultimately the data before the end of 2007 was not taken into account for the regression.

4 Conclusion

This paper focuses on what the effects of a jump in the price of the underlying asset are on the valuation of options. This is done by analyzing the VIX and SVIX, a proposed measure for variance by Martin (2013). We have shown that the SVIX holds information over the return of the S&P 500, and that this information does not correspond with the already existing VIX. From this we can conclude that the SVIX and its way of dealing with jumps is of significant importance for analyzing the performance of an index. In evaluating the SVIX there were issues concerning the number of options available. Without a sufficient amount of options both the VIX and SVIX are impossible to properly evaluate, which led to a smaller window of the data being usable then expected. This could be solved by having a more complete dataset, one which the CBOE itself uses to evaluate the VIX. The number of observations per day per maturity per option type would have to be around 500 since a small difference between strikes is needed and a lot of the options will be useless if the strike of a call option is too low or that of a put option to high. This leads to a need for 1000 options for each date of maturity, while for the calculation several different maturities are needed. The availability of this data is a problem however, as often only the (most) traded options were taken into account in the optionmetrics datasets. Another problem would be the amount of data, having tens of millions of observations for the same window used in this paper. Calculations with these amounts would have to be performed on a proper programming language e.g. C++, Java, which is time extensive to program, or it would have to be done on a Matlab type mathematical package, which was already having difficulties handling 3 million observations.
The SVIX also reacted heavily to times of great volatility, more so than the VIX did. This is useful for when a portfolio calls for a more conservative approach, using the SVIX as an indicator for financial stress or using it to hedge risks with a risk minimizing approach.

5 Appendix

5.1 Proof SVIX

$$\mathrm{var}R_T = E\left[\left(\frac{S_t}{S_0}\right)^2\right] - \left[E\left(\frac{S_t}{S_0}\right)\right]^2 = \frac{e^{rT}\Pi(T)}{S_0^2} - e^{2rT} \tag{15}$$

with

$$S_t^2 = 2\int_0^\infty max(0, S_T - K)dK \tag{16}$$

in a equal weighted portfolio of call options of all strikes. No arbitrage opportunities implies that

$$\Pi(t) = 2\int_0^\infty P_{\mathrm{call}_{0,t}}(K)dK \tag{17}$$

To express this in out of the money options we us the put call parity:

$$P_{\mathrm{call}_{0,t}}(K) = P_{\mathrm{put}_{0,t}} + e^{-rt}(F_{0,t} - K) \tag{18}$$

to make

$$\Pi(t) = 2\int_0^{F_{0,T}} P_{put}(K)dK + 2\int_{F_{0,T}}^\infty P_{call}(K)dK + e^{-rt}F_{0,t}^2. \tag{19}$$

Using this definition for $\Pi(t)$ we get

$$\mathrm{var}R_T = \frac{2e^{rT}}{S_0^2}\left(\int_0^{F_{0,T}} P_{put}(K)dK + \int_{F_{0,T}}^\infty P_{call}(K)dK\right) \tag{20}$$

Which was the definition of the SVIX^2 under assumption 1 and 2:

$$SVIX^2 = \frac{2e^{rT}}{TF_{0,T}^2}\left(\int_0^{F_{0,T}} P_{put}(K)dK + \int_{F_{0,T}}^\infty P_{call}(K)dK\right). \tag{21}$$

5.2 Discretization

First we define the step size as $\Delta K_i = (K_{i+1} - K_{i-1})/2$. The set of options that hold usable information (puts with strikes smaller than $F_{0,T}$ and calls with strikes larger than $F_{0,T}$) we define $\Omega_{0,T}$. The part of the formula that is influenced by the discretization is only the integral over the set of options used, or

$$\int_0^\infty \Omega_{0,T}(K)dK \tag{22}$$

This is by the CBOE in the calculation of the VIX approximated by

$$\sum_{i=1}^{n} \Omega_{0,T}(K_i)\Delta K_i. \tag{23}$$

How well does this Riemann sum approximate the integral?

$$\frac{2e^{rt}}{TF_{0,T}^2}\sum_{i=1}^{n} \Omega_{0,T}(K_i)\Delta K_i \leq \frac{2e^{rt}}{TF_{0,T}^2}\int_0^\infty \Omega_{0,T}(K)dK + \frac{(\Delta K_j)^2}{4TF_{0,t}^2} \tag{24}$$

or in words: the discretization is smaller than or equal to the ideal SVIX plus a small amount. This is also true for the VIX:

$$\frac{2e^{rt}}{T}\sum_{i=1}^{n} \frac{\Omega_{0,T}(K_i)}{K^2}\Delta K_i \leq \frac{2e^{rt}}{T}\int_0^\infty \frac{\Omega_{0,T}(K)}{K^2}dK + \frac{(\Delta K_j)^2}{4TF_{0,t}^2} \tag{25}$$

where again the discretized version is smaller than or equal to the ideal VIX plus a small amount. Equation 24 and 25 only hold if the appropriate assumptions hold. It is easy to see that if the step size goes up, so when there is a small amount of options, the error goes up rapidly as then $(\Delta K_j)^2$ increases.

5.3 Graphs

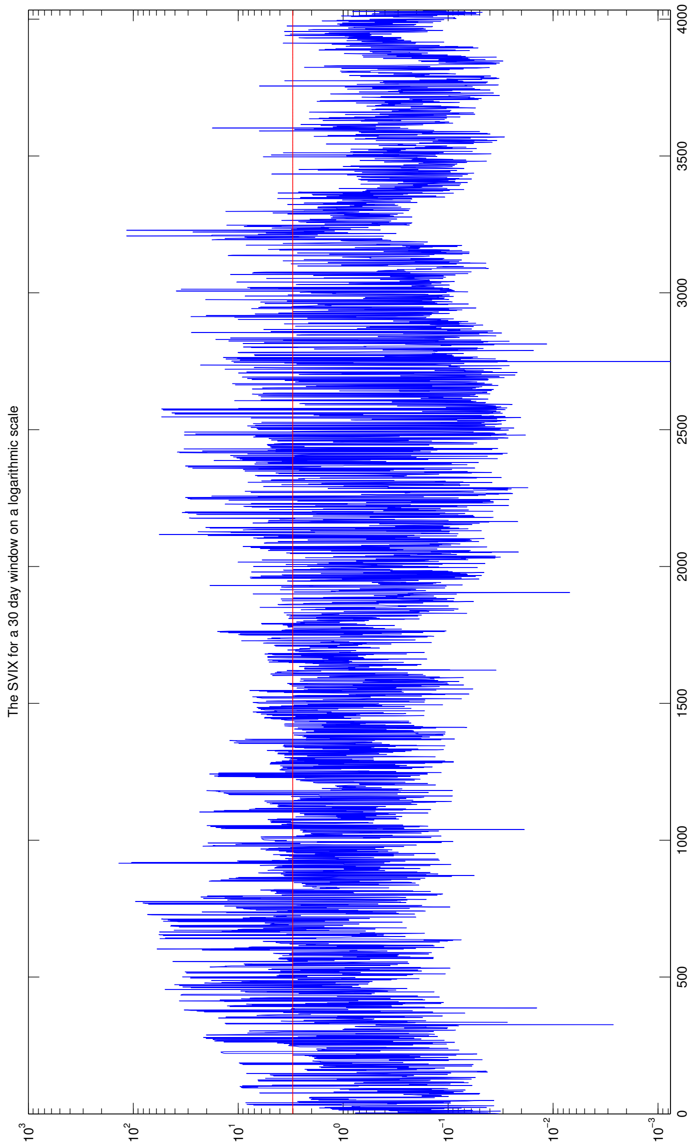
The SVIX for a 30 day window on a logarithmic scale
0
500
1000
1500
2000
2500
3000
3500
4000
10^{3}
10^{2}
10^{1}
10^{0}
10^{-1}
10^{-2}
10^{-3}

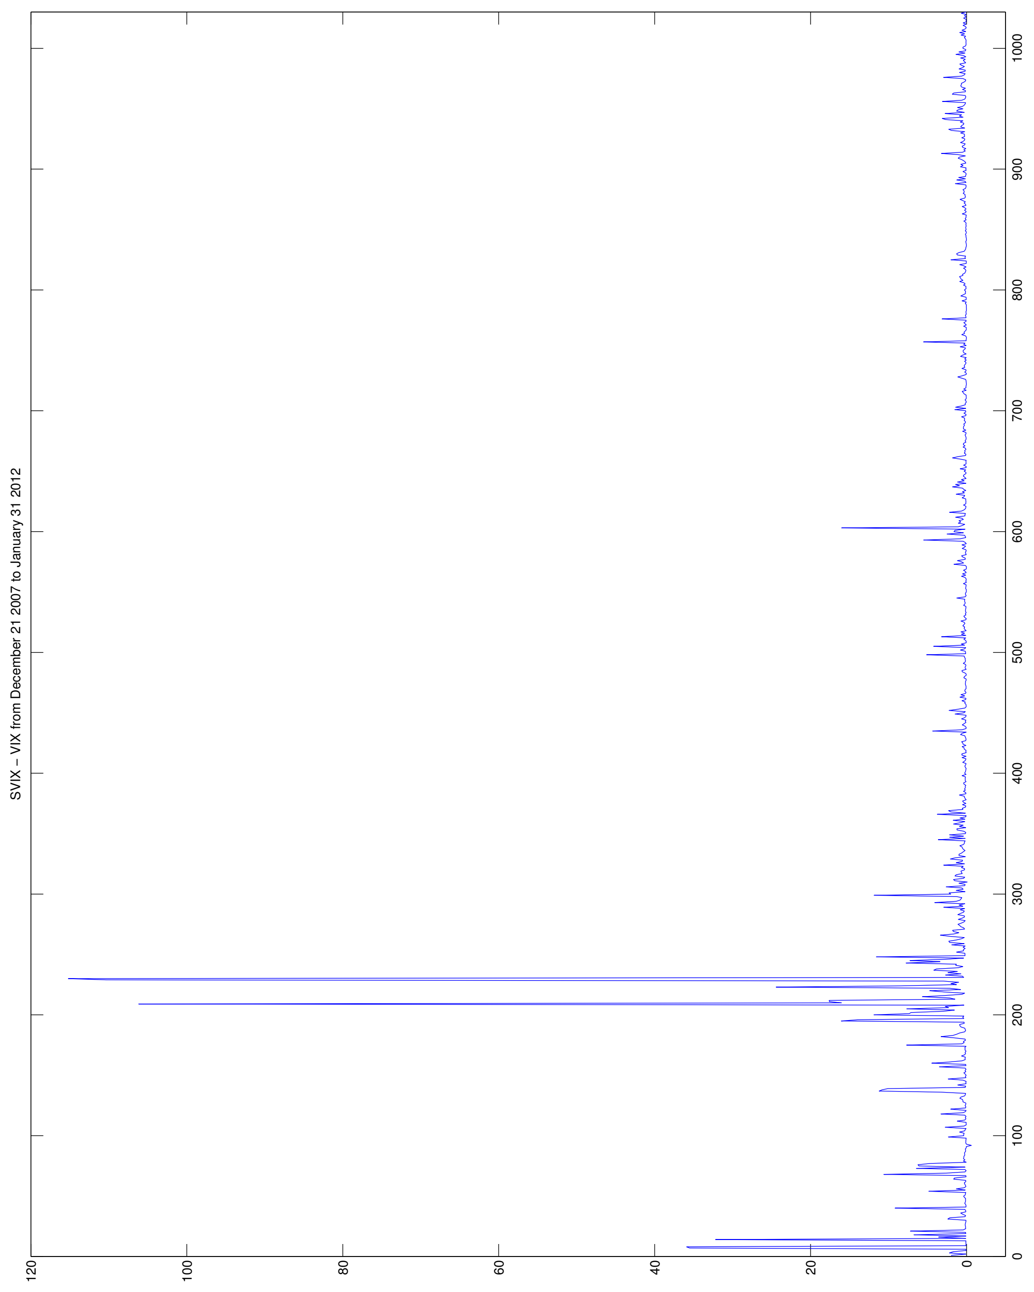
SVIX – VIX from December 21 2007 to January 31 2012
0
20
40
60
80
100
120
0
100
200
300
400
500
600
700
800
900
1000

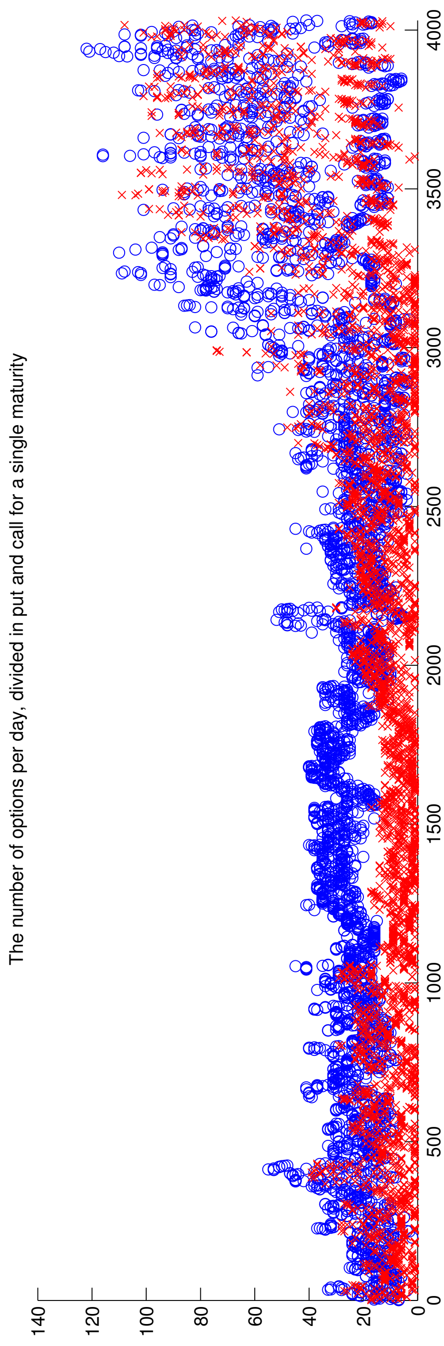
The number of options per day, divided in put and call for a single maturity
0
500
1000
1500
2000
2500
3000
3500
4000
0
20
40
60
80
100
120
140

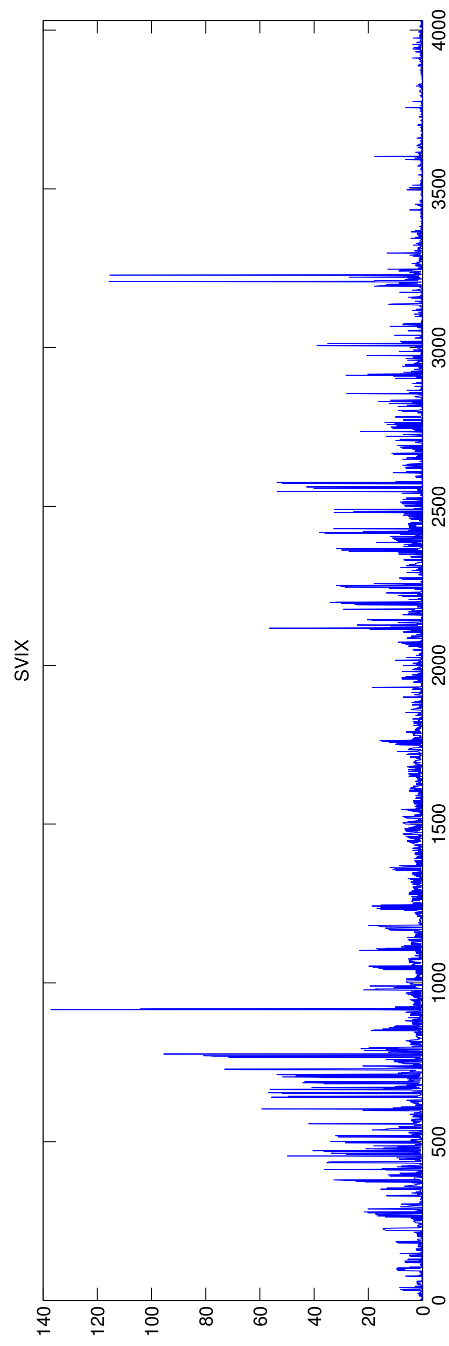
SVIX
0
500
1000
1500
2000
2500
3000
3500
4000
0
20
40
60
80
100
120
140

References

[1] T. Bollerslev G. Tauchen ans H. Zhou. Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492, 2009.

[2] T. Bollerslev and V. Todorov. Tails, fears, and risk premia. *The Journal of Finance*, 6(11):2165–2211, 2011.

[3] D. Backus M. Chernov and I. Martin. Disasters implied by equity index options. *The Journal of Finance*, 66(6):1969–2012, 2011.

[4] F. N. Fritsch and R. E. Carlson. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980.

[5] P. Friz and J. Gatheral. Valuation of volatility derivatives as an inverse problem. *Quantitative Finance*, 5(6):531–542, 2007.

[6] I. Martin. Simple variance swaps. *Working paper*, 2013.