Estimating the Dynamics of Volatility

by

David A. Hsieh

Fuqua School of Business
Duke University
Durham, NC 27706
(919)-660-7779

October 1993

Prepared for the Conference on Financial Innovations: 20 Years of Black/Scholes and Merton, November 18-19, 1993, Fuqua School of Business, Duke University, Durham, NC.

Estimating the Dynamics of Volatility
by
David A. Hsieh

The volatility of financial markets has long been a favored subject of investigation for academics and market participants. Since volatility is not observed, there has been no agreement on how to measure it. However, one conclusion appears to have emerged, namely, that volatility is volatile. This paper examines various measures of volatility, and proposes a diagnostic to test which of these measures of volatility best captures the dynamics of volatility of daily price movements.

The paper has five sections. Section 1 discusses the various measures of volatility, including three price-based measures of volatility (historical volatility, close-to-close volatility, and intraday volatility) and two option-based measures of volatility (implied volatility of at-the-money call and put options). Section 2 examines the properties of these five volatility measures. Section 3 estimates the dynamics of volatility. Section 4 proposes a diagnostic to test for the best measure of volatility. Section 5 provides concluding remarks.

1. Measures of Volatility

In this section, we define the various measures of volatility. While this methodology applies to analysis of volatility in all financial markets, we restriction our attention to the foreign currency market, in particular, the U.S. Dollar/Deutsche Mark exchange rate. Like the U.S. government bond market, the foreign exchange (FX) market is an over-the-counter market where transactions are generally conducted through interbank networks. The liquidity of the FX market is by far the highest of all financial markets, estimated to be around $1 trillion per day, with Dollar/Mark being the most widely traded currency.

Due to the nature of the interbank market, transactions data are not available. While it is possible to examine quotations obtained through information agencies such as Reuters or Telerate, quotes (which are solicitations to trade) appear to have substantially different characteristics than transactions prices. Thus, we focus our attention on the Deutsche Mark (DM) futures contract on the Chicago Mercantile Exchange (CME), which also trades options on these futures contracts.

The tick-by-tick (also called quote capture or time-and-sales) data contain the time and price of every transaction in which the price has changed from the previous transaction. In addition, a bid price is recorded if it is above the previous transaction, and an ask price is recorded if it is below the previous transaction. Since these bid and ask prices do not represent actual transactions, we eliminated them from our sample. Note that there is no information on the number and volume of transactions at any given price. Our data began on February 25, 1985, when daily price limits were removed on currency futures, and ended on June 28, 1991, spanning 1605 trading days. Since futures contracts expire 4 times per year, we use the contract which is nearest to maturity, switching to the next nearest to maturity on the Friday preceding the second Wednesday of each expiration month.

We begin our analysis by defining the term 'volatility.' Let F_t be the settlement price of the DM futures contract at date t. Let $x_t = \ln[F_t/F_{t-1}]$ be the continuously compounded rate of change, where "ln" denotes natural logarithm. The volatility of the DM futures contract, denoted by σ_t, is the standard deviation of x_t.

As σ_t is not observable, we proxy it in different ways. If we are willing to assume that x_t is normally distributed with mean zero and variance σ_t, then the expected value of the close-to-close volatility,

$$av_t = (\pi/2)\ |x_t|,$$

is σ_t. Unfortunately, this is a very noisy measure of σ_t, because it uses only one observation per day.

Next, we consider a popular measure, called historical volatility, which is the standard deviation of past observations of x_t. In this paper, we use a

20-day rolling measure:

$$hv_t = \{ \Sigma_i [x_{t-i} - \Sigma_j x_{t-j}/20]^2 / 20 \} .$$

While hv_t is less noisy than av_t because it uses more data, the rolling window induces a moving average process of order 19 in hv_t.

Instead of using close-to-close returns, as in av_t and hv_t, we can make use of tick-by-tick information on the DM futures contract. In particular, the intraday volatility is the standard deviation of the 15-minute rates of change of the nearby futures contract, denoted as iv_t. It is appropriate to discuss the choice of a 15-minute interval. In tick-by-tick data, as in most transactions data, there are bid-ask bounces, which induces a large and negative first-order serial correlation in the data. We need a sufficiently long time interval, such as 15 minutes, to remove this effect. We note that, while the volatility is likely to be changing over the course of a trading day, we are interested in the cumulative volatility from close to close. As long as daily "seasonals" in volatility are not time varying, the intraday volatility is reasonable proxy of the close-to-close volatility.

Aside from the three volatility measures using price data alone, we can use information from options on the DM futures contract, which are also traded on the CME. In particular, we calculate the implied volatilities of at-the-money (ATM) calls and puts, denoted cv_t and pv_t, respectively. They are obtained as follows. For each day, we choose the nearby DM futures contract and the options on that contract that matures in the same month with at least 10 days to maturity. We match futures and options prices using the tick-by-tick data from the CME, selecting the strike price closest to the futures price at the close of the trading day. The interest rate is taken to be the Treasury bill rate that matures nearest to the options expiration data. The implied volatility of the option is then calculated using the Barone-Adesi and Whaley [1987] approximate solution to American options.

2. Properties of Volatility

These measures of volatility provide some insights on the properties of volatility. First, they confirm the general impression that volatility is time varying and serially correlated. Table 1 provides the autocorrelation coefficients of these various measures of volatility. The standard error of these correlation coefficients is 0.025. Since the coefficients themselves are typically many times larger than this standard error, there is good evidence that volatility is not only volatile, but also autocorrelated. In the case of the historical volatility, which is a 20-day rolling measure, it is not surprising that the first 19 autocorrelation coefficients are large. However, the next 10 autocorrelation coefficients are more than two times larger than their standard errors, indicating a fair amount of persistence. Even in the cases of iv_t and av_t, which use non-overlapping data to construct a daily measure of volatility, the correlation of the 20-th lag is still statistically different from zero.

Second, the degree of volatility persistence depends on the measure of volatility. Three measures (hv_t, cv_t, and pv_t) indicate that volatility is highly persistent because they have large first-order autocorrelation coefficients, which are close enough to unity that volatility appear to be a nonstationary, unit-root like, process. On the other hand, the remaining two measures (av_t and iv_t) indicate that volatility is much less persistent because they have much lower first-order autocorrelation coefficients, which are far enough away from unity that volatility is a stationary process. On closer examination, hv_t is much more stationary that cv_t and pv_t. The autocorrelation coefficients of hv_t are similar in size to those of av_t and iv_t, while the autocorrelation coefficients of cv_t and pv_t remain substantially higher, even out to the 40-th lag. The price-based measures of volatility (hv_t, iv_t, and av_t) indicate that volatility is a stationary process, while the option-based measures of volatility (cv_t and pv_t) indicate that volatility may be a nonstationary process with unit-root type behavior.

Our economic intuition rules out the possibility that volatility is a unit-root process, since such a process leads to arbitrarily high volatilities with certainty. In fact, the Dicky-Fuller test indicates that cv_t and pv_t are stationary processes. However, we are still faced with the fact that the option-based measures of volatility find much more persistence in volatility than do the price-based measures.

A potential explanation of this disagreement is the presence of two different components of volatility: a short term component which is fast moving, and a long term component which is slow moving. Both components are stationary. The price-based measures of volatility are capturing only the short term component. The amount of noise in high frequency data masks the slow moving, long term component of volatility. Option-based measures of volatility, on the other hand, is capturing more of the long term component, since the option is forecasting the average volatility over its life time. As we constrain the option maturity to be longer than 10 days (but typically shorter than 110 days), the implied volatility is dominated by the slow moving long term component of volatility. If this is the explanation, the "correct" way to measure and predict volatility will depend on the horizon. To the extent that we are interested in short term (e.g. one day) volatility, the price-based measures are more appropriate. The option-based measures would be more appropriate for longer term (e.g. one month) volatility.

Another explanation of the disagreement in volatility persistence between price-based and option-based measures is that the latter is the result of a misspecification of the option pricing model. The option pricing model may have incorrectly assumed a log normal distribution for the underlying asset's price. Or the option pricing model may have omitted important variables, such as the price of volatility risk, in the case that volatility is stochastic and so an option cannot be replicated by arbitrage. The persistence in volatility is a result of the systematic mispricing of the options by the (misspecified) pricing model.

3. Estimating Dynamics of Volatility

As we pointed out in the previous section, the appropriate measure of volatility depends on the time horizon. For the purpose of this paper, we assume that the horizon is one trading day. This choice is not entirely random. Many interesting questions in financial risk management concern price distributions from the close of one trading day to the close of the next trading day. For example, futures exchanges typically collect margins and mark the positions of traders to market once a day at the close. These futures exchanges set their prudential margins to protect their clearing members from an extreme price move over the course of a trading day. The amount of margin is therefore related to the daily volatility of the futures price in question.

In this section, we will estimate the dynamical properties of volatility. As all five measures of volatility are stationary processes, we describe them by simple autoregressive time series models, of the following form:

$$y_t = a + \sum_{i=} b_i \, y_{t-i} + e_t,$$

where p is the lag length and y_t is the variable of interest. Using the Schwarz [1978] information criterion, we determine p to be 1 for av_t, 21 for hv_t, 7 for iv_t, 3 for cv_t, and 2 for pv_t. This is taken to be the minimal value of p. Then, we increase p until the regression residuals, e_t, are no longer serially correlated. This yields p to be 7 for both av_t and iv_t, 3 for cv_t, and 2 for pv_t. We are, however, forced to abandon hv_t because the serial correlation of e_t persists even when we increase p to 30 lags. The regressions are reported in Table 2. In all cases, past volatility is useful in predicting current volatility. Since the price-based volatility measures, av_t and iv_t, have low degrees of autocorrelation, the R^2's of their regressions are low. On the other hand, the option-based volatility measures, cv_t and pv_t, have high degrees of autocorrelated, so the R^2's of their regressions are much higher.

4. Diagnostic Test

As the time series properties of these measures of volatility are quite different, we now investigate which is a better measure. Our criterion is as follows. Based on the regression in Table 2, we obtain the (in-sample) fitted values of volatility for, say av_t, denoted by fav_t. Then, we construct the standardized variable zav_t:

$$zav_t = x_t / fav_t.$$

Under the assumption that x_t has mean zero and standard deviation σ_t, if fav_t is a good estimate of the volatility σ_t, then zav_t should have mean zero, and standard deviation 1. In addition, if fav_t correctly measures the dynamics of σ_t, then $|zav_t|$ should not be serially correlated. Similarly, we construct the fitted values of iv_t, cv_t, and pv_t, denoted as fiv_t, fcv_t, and fpv_t, respectively, and the corresponding standardized variables ziv_t, zcv_t, and zpv_t.

Table 3 provides the diagnostics for these standardized variables. All four standardized variables have means which are not statistically different from zero. In addition, there appears to be little autocorrelation coefficients of $|zav_t|$, $|ziv_t|$, $|zcv_t|$, and $|zpv_t|$. This means that the autoregressive models for all four measures are correctly capturing the dynamics of daily volatility. However, the standard deviation of zav_t is statistically greater than 1; that of ziv_t less than 1; only those of zcv_t and zpv_t are not statistically different from 1. This means that fav_t tends to underestimate daily volatility. The opposite is true of fiv_t. Only fcv_t and fpv_t are unbiased estimates of daily volatility. On the basis of this in-sample test, we consider fcv_t and fpv_t to be the best estimates of one-day ahead volatility.

5. Concluding Remarks

This paper measures the daily volatility of DM futures prices using both price-based methods and option-based methods. All volatility measures indicate that volatility is volatile. Except for historical volatility, the other four measures (av_t, iv_t, cv_t, and pv_t) indicate that volatility can be described as a stationary, autoregressive process. Autoregressive models were identified and estimated, and fitted values of volatility are obtained. These fitted values indicate that the autoregressive models were able to capture the dynamics of volatility. However, only the option-based measures (cv_t and pv_t) were unbiased predictors of volatility. This indicates that option-based measures of volatility can be valuable in providing accurate forecasts of daily volatility.

References:

Barone-Adesi, G. and R. Whaley, 1987, Efficient Analytic Approximations of AMerican Option Values, *Journal of Finance*, 42, 301-320.

Schwarz, G., 1978, Estimating the Dimension of a Model, *The Annals of Statistics*, 6, 461-464.

Table 1
Autocorrelation of Measures of Volatility

Lag	hv_t	iv_t	av_t	cv_t	pv_t
1	0.943	0.341	0.034	0.965	0.959
2	0.887	0.287	0.036	0.938	0.934
3	0.835	0.269	0.066	0.916	0.912
4	0.788	0.295	0.058	0.893	0.890
5	0.752	0.292	0.079	0.873	0.872
6	0.716	0.262	0.083	0.858	0.855
7	0.678	0.285	0.079	0.840	0.837
8	0.642	0.238	0.090	0.822	0.819
9	0.608	0.243	0.057	0.806	0.800
10	0.576	0.242	0.146	0.787	0.784
11	0.546	0.195	0.024	0.768	0.766
12	0.517	0.214	0.043	0.750	0.748
13	0.485	0.196	0.088	0.732	0.730
14	0.451	0.203	0.051	0.716	0.713
15	0.418	0.221	0.099	0.700	0.696
16	0.385	0.178	0.035	0.683	0.681
17	0.352	0.171	0.041	0.667	0.663
18	0.318	0.137	0.068	0.651	0.650
19	0.285	0.197	0.032	0.635	0.634
20	0.247	0.180	0.055	0.622	0.620
21	0.251	0.176	0.001	0.609	0.607
22	0.256	0.142	0.032	0.595	0.593
23	0.258	0.136	0.033	0.580	0.578
24	0.254	0.148	0.006	0.565	0.563
25	0.232	0.114	0.028	0.551	0.547
26	0.210	0.091	-0.021	0.539	0.535
27	0.191	0.114	0.054	0.527	0.523
28	0.170	0.140	0.023	0.517	0.515
29	0.150	0.117	0.024	0.506	0.505
30	0.131	0.118	0.015	0.494	0.495
31	0.111	0.087	-0.015	0.482	0.485
32	0.091	0.109	-0.027	0.472	0.475
33	0.073	0.095	0.057	0.463	0.467
34	0.057	0.075	-0.036	0.455	0.455
35	0.041	0.090	0.030	0.447	0.447
36	0.025	0.084	0.024	0.439	0.439
37	0.004	0.085	0.006	0.431	0.431
38	-0.017	0.070	-0.025	0.424	0.424
39	-0.038	0.091	0.022	0.416	0.418
40	-0.054	0.077	0.049	0.409	0.411

Notes:
hv_t: 20-day historical volatility, hv_t.
iv_t: intraday volatility, iv_t.
av_t: $(\pi/2)\,|x_t|$, av_t.
cv_t: at-the-money call option implied volatility, cv_t.
pv_t: at-the-money put option implied volatility, cv_t.
One standard error of the autocorrelation coefficients is 0.025.

Table 2
Estimating Volatility Dynamics

Regression: $y_t = a + \sum_{i=} b_i\, y_{t-i} + e_t$

$y_t =$	av_t	iv_t	cv_t	pv_t
a	0.0758 (0.0071)	0.0400 (0.0052)	0.0042 (0.0009)	0.0042 (0.0009)
b_1	0.0110 (0.0281)	0.1824 (0.0250)	0.8531 (0.0352)	0.7617 (0.0436)
b_2	0.0135 (0.0271)	0.0898 (0.0244)	0.0295 (0.0379)	0.2040 (0.0436)
b_3	0.0483 (0.0271)	0.0681 (0.0258)	0.0875 (0.0310)	
b_4	0.0449 (0.0252)	0.1153 (0.0246)		
b_5	0.0696 (0.0274)	0.1027 (0.0279)		
b_6	0.0728 (0.0297)	0.0583 (0.0274)		
b_7	0.0668 (0.0258)	0.1110 (0.0280)		
R^2	0.0236	0.2175	0.9355	0.9264
Test of $\sum_i b_i = 1$				
χ^2 (dof)	107.04 (6)	60.37 (6)	16.72 (1)	20.85 (2)

Notes:
Standard errors in parentheses.

Table 3
In-Sample Diagnostics of Volatility Dynamics:

	zav_t	ziv_t	zcv_t	zpv_t
Mean	0.0375	0.0219	0.0277	0.0282
Std Dev	1.0677	0.8236	0.9756	0.9724
t(Mean=)	1.40	1.06	1.13	1.15
t(Std Dev=1)	2.69	-7.01	-0.97	-1.10
Autocorrelation Coefficients of Absolute Values:				
Lag 1	-0.000	-0.025	-0.023	-0.021
2	-0.005	-0.035	-0.018	-0.019
3	-0.023	-0.019	0.001	0.001
4	-0.010	-0.011	0.010	0.011
5	-0.008	0.007	0.031	0.031
6	-0.024	0.000	0.020	0.021
7	-0.022	-0.007	0.020	0.020
8	0.054	0.022	0.033	0.035
9	0.035	0.013	0.025	0.024
10	0.106	0.080	0.094	0.094
11	-0.008	-0.032	-0.016	-0.018
12	0.000	-0.012	-0.006	-0.006
13	0.046	0.019	0.032	0.032
14	0.025	-0.001	0.005	0.003
15	0.063	0.047	0.045	0.045
16	0.005	-0.004	0.009	0.007
17	0.016	-0.002	0.004	0.004
18	0.044	0.032	0.035	0.034
19	0.016	-0.002	0.002	0.003
20	0.023	0.008	0.018	0.019
21	-0.014	-0.027	-0.015	-0.017
22	0.015	-0.001	0.007	0.009
23	0.000	-0.007	-0.008	-0.007
24	-0.004	-0.024	-0.011	-0.011
25	0.015	-0.010	0.007	0.006
26	-0.049	-0.068	-0.057	-0.059
27	0.048	0.035	0.044	0.041
28	0.005	-0.010	-0.006	-0.007
29	0.017	-0.008	-0.000	0.001
30	-0.008	-0.024	-0.022	-0.022
31	-0.015	-0.039	-0.029	-0.028
32	-0.025	-0.029	-0.040	-0.043
33	0.049	0.035	0.037	0.036
34	-0.047	-0.063	-0.055	-0.055
35	0.031	0.019	0.024	0.023
36	0.015	-0.005	0.010	0.008
37	0.005	-0.006	0.007	0.006
38	-0.028	-0.040	-0.035	-0.037
39	0.025	0.015	0.017	0.018
40	0.048	0.028	0.031	0.031

Note:
One standard error of the autocorrelation coefficients is 0.025.