

An Information-Based Model of Market Volatility

A model for volatility, based on the relation between volatility and information flows, leads to the specification of a stochastic process for volatility. This allows one to compute potentially useful properties, such as the probability that volatility will change from one level to another within a specific time period.

The model relies on three characteristics of information. (1) Information arrives in discrete "packets," and the probability of its arrival is a function of time. (2) Different pieces of information have different degrees of impact on the market, hence on the market's volatility. (3) It takes time for the market to digest information; the greater the impact of the information, the longer its effect on volatility will last.

Comparison of actual daily volatility for the Treasury bond yield with the volatility given by the information-based model shows clear similarities between the frequencies of variability, the ranges of volatility, the speeds with which spikes dampen and the longer-term waviness of the processes. The model also did well in explaining equity market volatility following the 1987 crash.

The model suggests that volatility is mean-reverting. The tendency is for above-average volatility to decline and for below-average volatility to increase. The nature of the information flows underlying the model suggests that volatility can be expected to be reasonably stable over long time periods; the mean level of volatility may not change dramatically, even over five or 10 years.

THE ANALYTICAL SOPHISTICATION of the financial markets has lent increasing importance to volatility. Volatility is central to many investment decisions in new product areas and is the critical variable in options. Just as bond traders make quotes in terms of yield rather than price, option traders typically make option quotes in terms of volatility rather than dollar value. Option-related products such as interest rate caps and floors vary in price according to volatility. The pricing and invest-

ment attractiveness of mortgage-backed securities, callable bonds and other fixed income instruments with embedded options depend on market volatility. The cost of many sophisticated hedging strategies, such as portfolio insurance, also depends on the level of market volatility.

Despite its central importance to investment analysis, to date there has been no model that describes the nature of volatility. Much has been written on estimating volatility, using past security prices or the volatility implied in option prices. And attempts have been made to specify volatility as a time-series process. But the questions central to understanding volatility—what determines volatility, what accounts for the short-term variability and the long-term stability of volatility, how volatility differs from one market to another—remain unanswered. Vola-

Richard Bookstaber is a Principal of Morgan Stanley. Steven Pomerantz is a Vice President of Citibank.

This paper was presented at the 1989 AMEX Options Colloquium and the 1989 meetings of the Institute of Management Science. The authors thank Jeremy Evnine, Sandy Grossman, Ravi Myneni and Zvi Schechner for their comments.

Table I Historical Volatilities of Selected Markets

Market	Annualized Volatility
S&P 500	15%
Major Market Index	17
Nikkei 225 Index	15
Gold	25
Copper	24
Platinum	32
Oil	22
Japanese Yen	11
German Mark	13
Treasury Bonds (Yields)	15

tility analysis today depends on the same rudimentary estimation methods, based on historical and implied volatility, used in the mid-1970s.

This article develops a model for volatility, based on the relation between volatility and information flows. The model leads to the specification of a stochastic process for volatility that allows us to compute potentially useful properties, such as the probability that volatility will change from one level to another within a specified time period. The model also provides a basis for deeper understanding and analysis of volatility, because it establishes a relationship between information and volatility.

Volatility and Information

The volatility of an asset indicates the variability of its returns. Conventionally, volatility is measured as the annual standard deviation of the asset's returns.¹ It thus expresses variability around a trend growth rate.

An important aspect of volatility is its emphasis on the variability, rather than the direction, of prices. Each path illustrated in Figure A, for example, has the same 18 per cent volatility, even though the direction of prices differs in each case. To give a sense of the range of volatility, Table I presents the average historical volatilities of a number of markets.


The Impact of Information

Information leads to changes in expectations, which in turn lead to changes in prices. Because volatility is the product of unanticipated price movements, it is closely related to information. In fact, we can think of price volatility as nothing more than a manifestation of information in the market. If information never came along to cause a reevaluation of the markets, prices

would be set according to the present value of expected future value and would grow smoothly.

The information that affects asset prices comes primarily from financial and economic announcements, such as announcements of the Consumer Price Index, money supply, corpo-

Figure A Same Volatility, Different Price Trends



1. Footnotes appear at end of article.

rate earnings, Gross National Product or trade numbers. The extent to which the numbers contained in such announcements differ from the values expected by the consensus can be used as a gauge of the information content of the announcements. If the government announces a trade deficit of \$14 billion, for example, and the consensus had expected a deficit of \$14 billion, then the announcement has little information content and will have little effect on volatility. If the consensus had expected a deficit of \$20 billion, however, then the announcement of a \$14 billion deficit would have substantial information content and a large impact on price volatility.

The types of information mentioned above are conveniently measurable. Their arrival and the consensus expectation of their arrival, their content and the consensus expectation of their content are easily monitored and quantified. The timing and impact of other sources of information, such as surprise political events (wars or assassinations) or natural phenomena (earthquakes or floods) are not as easily anticipated or quantified. But their overall effect is not as pervasive, either.

Another source of volatility is the trading brought about by the liquidity and asset allocation needs of institutions and the market-timing decisions of investors. For the purposes of this article, we can think of this trading as the result of information—changes in the parameters of market timing models or in the liability needs of pension plans, for example. Because it is private information, however, it cannot be monitored the way public economic information can.

Although not all sources of information can

Figure B Information Arrival and Assimilation





Figure C Magnitude of Arriving Information (cumulative probability distribution)



be easily monitored, they can in principle be identified and quantified. Information is thus not an ethereal concept. The frequency with which information arrives in the market, its average impact and the time it takes the market to adjust to it and discount it can all be analyzed.

An Information-Based Model of Volatility

We use three characteristics of information to develop a volatility model. First, information arrives in discrete units or "packets," and the probability of its arrival is a function of time.² Second, different pieces of information have different information contents and therefore different effects on market volatility. Third, once information comes into the market, it takes time for the market to "digest" it fully.

The assertion that a packet of information has an extended, though diminishing, impact on volatility does not conflict with allowing information to be immediately reflected in the market price. The impact of information on volatility is due to the price being revised as the implications of the information packet continue to be analyzed by the market. The model suggests that the extent of these revisions will diminish, the longer the information is analyzed by the market, and that the bigger the impact of the information, the longer its overall effect on volatility will last.

Figure B illustrates the process of information arrival into the market. A jump in the line represents new information. We can denote the average arrival time of information in the market by λ . The impact or amplitude of the information can

be measured by its height on the vertical axis, h . The decay of the information's effect is illustrated by the curving line sloping downward from the information's arrival point. The half-life of the curve (the time it takes for the curve to drop to one-half its original height) is determined by alpha, α .³

Figure C shows the distribution of information amplitude—the probability that information packets with given information content will arrive within a given time span. The probability of information of a given size arriving decreases exponentially with the amplitude of that information.

We can now envision volatility as a process that is a function of three parameters—the average arrival time of information, λ ; the magnitude, or importance, of information, h ; and the speed with which information is analyzed by the market, α . As we will see, these parameters differ across markets and may differ within a market over time.⁴

By assuming that volatility is linearly related to information, we can specify volatility as a compound Poisson process. The volatility in the market is given by the following equation:

$$\sigma(t) = \sum_{k=1}^{N_t} H_k \exp[-\alpha(t - t_k)],$$

where N_t is the number of information packets that arrive in the interval from 0 to t ($0, t$). The H_k are independent exponential random variables that represent the magnitude of the information.⁵

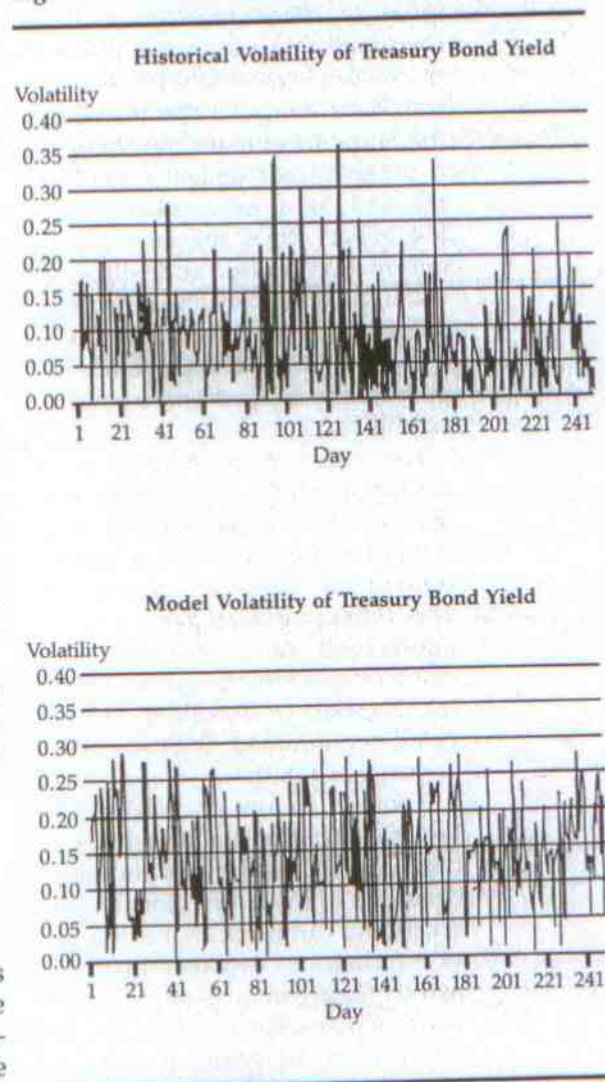
Does the information-based model of volatility provide a reasonable estimate of actual volatility? The examples below provide some indication of how the model works in various markets.

Bond Market Volatility

Figure D compares the actual daily volatility of the Treasury bond yield with volatility given by the information-based model.⁶

The model volatility graph is based on the stochastic process defined by the three parameters λ , h and α . We estimated the first two moments of the actual volatility distribution as well as the first-order serial correlation, and used these three values to solve for the three parameters.⁷ Placing these values into the model generated the stochastic process that

Figure D The Stochastic Process of Volatility



yielded the time series of instantaneous volatility shown in Figure D.

A visual comparison of the two graphs shows a clear similarity in their stochastic processes—the frequency of variability, the range of variability, the speed with which spikes dampen and the longer-term waviness of the process masked by the spikes. The most noticeable difference between the two processes is that the actual volatility occasionally ranges above 30 per cent, while the model process is truncated on the upside. This difference suggests that the exponential distribution we are using for the magnitude of information may not capture the number of high-information events that were actually observed in the market. That is, the

distribution we are using may not provide a thick enough tail.

Autocorrelation of Volatility

One tool for looking at the behavior of volatility is the correlation between current volatility and the volatility of previous periods. This is analyzed using a correlogram. Figure E compares the correlogram for the model with the correlation of actual Treasury bond volatilities.

The model correlogram has a negative exponential shape, starting high and rapidly decaying toward zero correlation. That is, the correlation between current and past volatilities (autocorrelation) drops off rapidly, the farther into the past we go. The speed of the decay is a function of the parameter values. In particular, the decay rate, α , determines the time span over which current volatility will be related to past volatilities. The arrival of major information, for example, will cause volatility to spike, but (everything else being equal) that spike will gradually dissipate as the information is assimilated into the market. From the time of the arrival of the information until the information is fully assimilated, volatility will exhibit autocorrelation.

It is also apparent from Figure E that the autocorrelation of actual bond volatilities has the same characteristic shape as the model autocorrelation. The autocorrelation of actual bond volatilities, however, tends to decay more rapidly than the model autocorrelation. This suggests that the exponential decay rate we use in the model may understate the speed of decay, especially immediately following the arrival of the information.

Figure E Autocorrelation of Volatility




Table II Volatility Decay for 1987 Market Crash

Volatility	Day Number
220%	0
110	2.87
55	5.73
27.5	8.60
13.75	11.46

Interpreting the Model Parameters

The model, by tying volatility to information flows, gives volatility an economic and financial interpretation. An important question is whether the model parameters make sense financially. What do the parameters say about the assumed rate of information arrival and the time it takes for information to be assimilated into the market?

Solving for the three parameters using the bond data, we obtained an estimate of 0.035 for the height parameter, h , which measures the expected magnitude of the information's impact. This value represents the average volatility that a new piece of information must "cause" if it is to offset the decay of previous information impacts. We obtained for α an estimate of 50; this means that roughly one-fifth of the volatility decays each day.⁸

If past information is assimilated into the market at this rate, then an amplitude of information of 0.035 in each information packet is sufficient to maintain the average 0.14 volatility of the actual bond data. This can be thought of as a "background radiation" of information, which is affecting the market continuously. An amplitude of information above this level leads to higher-than-average volatility, while momentary lulls in the information flow cause volatility to drop below the mean.

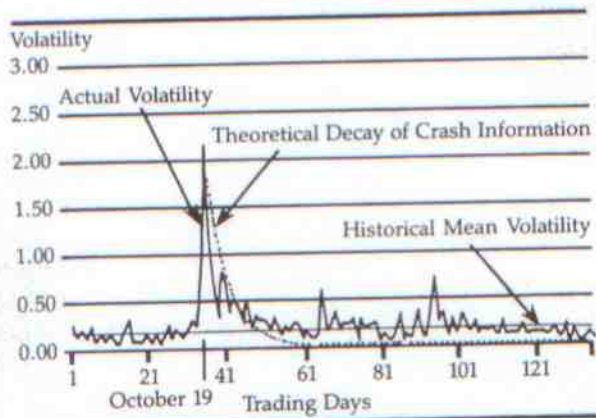
The frequency of information arrival is given by λ . For the bond data, we estimated a value of 197 for this parameter. The expected waiting time between information arrivals is the reciprocal of λ —0.5 per cent of a year, or 1.3 trading days.⁹

Volatility During Equity Market Crashes

The stock market crash of October 19, 1987 provides a particularly interesting test of the volatility model. Market volatility that day was over 10 times the historical level.

For the parameter α —the time it takes for information to be assimilated into the market,

Figure F S&P 500 Instantaneous Volatilities, 9/1/1987 to 3/8/1988



and therefore the time over which it will have a volatility impact—we estimated a value of 60.95, based on S&P 500 data. The number of trading days in the half-life of the impact of information on volatility can be computed as the value for t that sets the expression $\exp(-\alpha t)$ equal to 0.50. We thus derived a half-life of information of 2.87 trading days.

On the day of the crash, the instantaneous volatility (measured by the high/low prices) was 220 per cent. A half-life of 2.87 days suggests the volatility pattern shown in Table II. The model thus predicts that it would take approximately 14 trading days for the full impact of an event as tumultuous as the crash to be absorbed by the market.

Figure F shows the daily volatility of the S&P 500 over the period from September 1, 1987 to March 8, 1988. Superimposed on the volatility graph is the decay rate predicted by the volatil-

Figure G Stock Market Volatility




Table III Volatility Parameter Values for Various Markets

	λ/α	h
Equity		
S&P	3.94	0.031
MMI	4.33	0.032
NIKKEI	2.17	0.050
FTSE	5.13	0.024
Commodities		
Gold	1.82	0.125
Copper	1.91	0.125
Silver	1.28	0.125
Platinum	2.82	0.111
Foreign Exchange and Interest Instruments		
Yen	3.23	0.028
Deutschmark	3.17	0.033
Treasury Bond	5.31	0.027

ity model. The decay of volatility after October 19 followed the model's prediction, dampening over time at the prescribed rate and returning to the long-term mean after 12 days.

These results may come as a surprise to those who followed the option markets after October 19. The implied volatilities in the option markets remained at double their historical levels for several months after October 1987, and the general impression was one of a long-term, persistently high volatility. As Figure F shows, while there were several days over the course of the following months when high volatility recurred in the equity market, the impact of the volatility spike from October 19 dissipated in several weeks.¹⁰

1928-35 vs. 1980-87

Comparing the model results for the same market over different time periods helps to verify the stability of the model parameters. We estimated the model for the U.S. equity market over the 1928-35 and 1980-87 periods.¹¹

The 1928-35 period included three years of unprecedented high volatility, while the 1980-87 period was a fairly average time for the market. Figure G graphs market volatilities over these two periods. Given the differences in the volatility time series, we should expect differences in the model parameters. In fact, we arrived at values for λ/α and h of 3.44 and 0.064, respectively, for the 1928-35 period and 3.28 and 0.041 for the 1980-87 period.¹²

The ratio of frequency to decay rate of information (λ/α) in the two periods is virtually the same. The difference in the volatilities can thus be ascribed completely to the value of h , which

is the average impact of the information on prices. The impact of information on prices was roughly 50 per cent higher in the 1928-35 period than in the 1980-87 period.

This might have been expected, given the economies of the two periods. Many more firms were close to default during the Depression years; they would have had higher leverage, so any information would have been expected to lead to greater price movement. That the other parameters remained constant over these highly disparate times provides additional support for the robustness of the model.

Commodities, Foreign Bonds and Foreign Equities

To test the robustness of the model and provide some insight into information flows and volatility dynamics, we estimated the parameter values for other markets. In all cases, the data used covered January 1980 to October 1987. The volatility estimate was based on close-to-close prices for weekly periods. Table III gives the results.

By comparing the values in this table, we can conclude that the information flow in the foreign exchange and interest rate markets is similar to that in the equity markets. For the commodities, though, there are substantial differences. The height parameter is much larger, indicating that news events have a greater impact on the markets. This may be attributable to the less diversified nature of these markets. The ratio of arrival frequency to decay (λ/α), by contrast, is smaller, perhaps because news specific to these markets arrives less frequently than in the more economically driven markets.

Figure H Weekly Volatility of Silver, September 1979-September 1987




Figure I The Volatility Time Curve




Figure H shows the weekly volatility in the silver market. A visual comparison of this volatility with the Treasury bond yield volatility (Figure D) or the stock market volatility (Figure G) helps to illustrate these differences.

Applications of the Model

The immediate applications of the model are to show important stochastic properties of volatility and to compute the probability of volatility changing within a specified time period.

Mean Stability and Mean Reversion

A key characteristic that comes out of the model is that volatility is mean-reverting. A high volatility will tend to decline and a low volatility will tend to increase to a mean level. We can also draw some conclusions about the time expected for such a reversion to occur, but first the property of mean reversion itself deserves some comment.

Our model states that volatility is a function of the arrival rate, amplitude and decay rate of information in the market. The nature of information flows suggests that, over reasonably long time periods, these three parameters, hence volatility itself, can be expected to be stable. This is because the nature of information flows is tied to the foundation of the economic system.

The rate of information arrival depends on fundamental things such as the reporting structure and accounting periods of regulatory and financial institutions. The amplitude of information depends on the speed with which policy is set and changed by the Federal Reserve System and the legislature, and on the corporate struc-

