



Doomed to repeat it?

November 2008

Christopher C. Finger
chris.finger@riskmetrics.com

We can track the continuing deterioration of the financial markets by how far back we go to find comparable events: we started with the bursting of the technology bubble in 2000, then moved to the LTCM crisis of 1998, to the Scandinavian banking crisis of the early 1990s, to the recession of the 1970s and finally to the Great Depression of the 1930s. It has kept raining, and a ten-year flood turned into a twenty-year flood and finally into a one hundred-year flood. Accompanying the economic arguments that go with these comparisons have been citations of the recent market movements in the context of those over the last hundred years. But what about risk?

With all of the criticisms of models we have heard recently, we first need to establish that we can meaningfully quantify risk in the first place. Once we have done this, then we can start to ask questions about history. How volatile have markets actually gotten? How long do we have to wait, typically, for this level of volatility to abate? It may have stopped raining, but how long will it take the floodwaters to recede?

Volatility matters

We begin with an examination of the Dow Jones Industrial Average (DJIA). Though by no means a broad indicator of the market (as it contains only 30

US equities), it is attractive to examine for two reasons: first, it is the index most utilized by the popular press as representing the market; second, it has been calculated, mostly uninterrupted, for over one hundred years, and as such gives us the opportunity for a long historical perspective.

The simplest risk models rely on three basic tenets: volatility is relevant; volatility changes; and changes in volatility are (at least somewhat) predictable. The framework for the model is that the return to come is the product of the volatility, which we forecast using the information available at the time, and the *residual*, which we do not know, but which comes from a defined statistical distribution. Intuitively, each day's return can be thought of as an n -sigma event, where sigma is the standard deviation, or volatility, that we have forecast, and n is the size of the residual.


We work with a volatility forecast that is a simple weighted average of prior days' squared returns, with the weighting scheme from one of our standard risk models.¹ Importantly, the volatility we consider on any given day is a forecast that a risk manager could have made (had the techniques been invented yet) at the time. We plot the DJIA along with its volatility in Figure 1. Here and for the remainder of this article, we will refer to volatility in annualized terms. For the volatility, we also indicate the average over 1900–45 (20%) and over 1945 to the present (15%).

©2008 RiskMetrics Group, Inc. All Rights Reserved.

¹Specifically, the long-memory model described in Zumbach (2006), with standard parameter settings

²More precisely, the t -distribution with five degrees of freedom, scaled to have unit variance. As discussed in

Figure 1: Dow Jones Industrial Average, 1900–2008. Index level (top) and volatility (bottom)



Source: www.analyzeindices.com, Yahoo Finance, own calculations

For the residual, we assume the t -distribution.² Moreover, we assume that each day's residual is distributed this way, independent of whether volatility is high or low, of whether we are in the early or latter stages of our history and of whether there has been a large surprise (or lack thereof) in the immediate past.

Cynics will note that the word *assume* appears frequently in the preceding paragraphs. With so much data at our disposal, we have the luxury of testing these assumptions empirically. To do so, we calculate the historical residuals, dividing each day's return by the volatility we would have forecast just prior to the return. Again, we can interpret these residuals as how many standard deviations each

day's return represents, or in other words, how surprising the return was, based on our knowledge of volatility at the time.

With this large set of residuals, we can assess the assumptions we have made about their distribution. In a simple test, we choose a threshold, and compare how many residuals we actually see of this magnitude to what our assumed statistical distribution predicts. For instance, the t -distribution predicts that over the history in question (about 30,000 trading days), we should see between 29 and 49 days³ on which the market loss is a five-standard deviation event or greater; there are in fact 32 such days. Similarly, the t -distribution would predict four or fewer

Zumbach (2006), this number of degrees of freedom provides a good fit in general to most liquid financial time series.

³Based on a 90% confidence band around the expected value

losses of nine standard deviations or greater; there are four such days. The oft-maligned normal distribution would predict not a single five-standard deviation event, even over such a long historical period.⁴

So as we have emphasized in this space in the past, bad models do exist. But useful models exist as well, and the combination of a good volatility forecast and the *t*-distribution does fit the data here quite well. We cannot say with certainty this will always be the case; and the residual distribution is a statistical fit, rather than the result of a more fundamental modeling of how prices evolve. Still, one hundred and eight years is a good track record.

Surprises

With confidence that our volatility forecast does in fact provide useful information, we turn now to a historical examination. Returning to Figure 1, we see that historically, volatility has tended to stay below its average level for extended periods, with occasional spikes and shorter periods of higher volatility. Against this backdrop, the recent peak of 70% is extremely high, though lower than the spikes after the crashes in 1929 and 1987, when volatility jumped to 95% and almost 120%. Today's level is comparable to what we saw in the early 1930s, which was also the longest sustained period of volatility on record.

The volatility forecast allows us to assess which days produced the largest surprises in history. It is easy to find lists of the largest losses on the DJIA, both in absolute index points and in percentage moves.

We complement these lists with a list of the largest surprises—that is, losses on the index relative to the volatility at the time—in Table 1.

From the table, we see that the twenty largest surprises have been 5.6 standard deviations or greater. Under the *t*-distribution, we would expect 23 events of such magnitude. The magnitude (13.3) of the largest surprise is in the range of what we would expect; an event in the realm of twenty standard deviations would have in itself cast doubt on our assumptions. Providing additional validation to our model, the surprises seem to be distributed evenly throughout history, and to have occurred as often in times of low volatility as high.

Our basic assumptions also imply that surprises should not cluster: that is, the likelihood of a surprise on one day should have nothing to do with whether a surprise has occurred in the recent past. In particular, this model property is what we rely on to forecast over periods longer than a single trading day. We do see one cluster in the table: the two events in July 1914. These events came at the outset of World War I, when the fear that investors would retreat from equities and into gold bullion (the flight-to-quality trade of the early 20th century) compelled the New York Stock Exchange to close for four and one half months. July 30, the date of the second event, was the last day the exchange was open until December. The market has never been closed for such an extended period since, nor have we again seen six-standard deviation surprises so close together.

The first date on the list of surprises is in itself surprising, as 1955 is not a year we remember as a

⁴For the DJIA, though it does poorly in capturing the extreme losses, the normal distribution does in fact describe the data well in the one- to two-standard deviation range.

Table 1: Twenty largest surprise losses on the Dow Jones Industrial Average, 1900–2008, with volatility forecast prior to the event

Date	Residual	Return (%)	Volatility (%)	Date	Residual	Return (%)	Volatility (%)
26-Sep-1955	-13.3	-6.5	8.1	13-May-1940	-6.5	-5.0	12.4
19-Oct-1987	-12.6	-22.6	32.4	3-Nov-1948	-6.5	-3.8	9.6
29-Jul-1927	-10.1	-5.2	8.3	28-Oct-1929	-6.4	-12.8	34.2
13-Oct-1989	-10.0	-6.9	11.4	27-Oct-1997	-6.3	-7.2	18.8
26-Jun-1950	-8.1	-4.7	9.3	14-Mar-1907	-6.1	-8.3	22.7
27-Feb-2007	-7.8	-3.3	6.8	7-Dec-1904	-6.0	-5.0	13.4
20-Jan-1913	-7.0	-4.9	11.4	28-May-1962	-6.0	-5.7	15.4
30-Jul-1914	-6.7	-6.9	16.9	17-Sep-2001	-6.0	-7.1	19.6
28-Jul-1914	-6.7	-3.5	8.5	3-Sep-1946	-5.9	-5.6	15.5
15-Nov-1991	-6.6	-3.9	9.6	1-Feb-1917	-5.6	-7.2	21.3

rough one for the markets, and indeed volatility was low just prior to the 6.5%⁵ loss on September 26. In fact, this loss was precipitated by the heart attack suffered by President Eisenhower the previous weekend. Though his health situation was initially quite serious, Eisenhower recovered, and was released from the hospital in November. By that time, the market had recovered all of its losses. Volatility jumped to 30% after the initial shock, but was back down to 10% by the end of the year.

The Eisenhower story of course begs the question of where the Kennedy assassination, on November 22, 1963, fits in among market surprises. Like the Eisenhower heart attack, Kennedy's assassination came at a time of very low volatility—around 8%—and did spark a sell-off in the market. The loss that day, 2.9%, was 4.5 standard deviations; as surprises go,

this ranks as the 40th largest. Similar to the Eisenhower event, volatility jumped to the still benign level of 30%, and came down quickly thereafter.

Other presidential scares do not appear to have registered. Neither the attempt on Reagan in 1981, nor the assassination of McKinley in 1901 nor the health-related (but suspicious) death of Harding in 1923 produced so much as a two-standard deviation move.

Larger scale attacks, of course, also hold the potential to surprise. Viewed as a one-day move, the return from September 10, 2001 to September 17 (the day the exchange reopened after the September 11 attacks), was a six-standard deviation event, and volatility jumped afterwards to around 35% before settling down to 20% by year-end.

The attack on Pearl Harbor in 1941 came at a time of lower market volatility—just 11%—and the loss

⁵For the particularly keen reader, the returns we reference in the text and table are expressed in conventional percentage form. Our volatility and residual calculations were performed on logarithmic returns.

that day registered as a five-standard deviation event, the 34th largest surprise in history. Again, volatility jumped after the event, but the jump was short-lived, and market volatility remained low for most of World War II. Volatility over the course of the war averaged only 12%, lower than the historical average volatility, and significantly lower than the 18% volatility experienced during World War I.

Volatility and crises

Among the surprises are dates associated with three significant financial crises: the Panic of 1907, the Great Depression and the current crisis. In 1907, the surprise came in the form of a market fall in March amid general economic unease (new regulatory powers, reconstruction after the 1906 San Francisco earthquake, weakness in railroad stocks). Volatility rose to the elevated (for the time) level of 15% in early March, was at 22% by the time of the largest drop and rose to almost 50% the week after the surprise. Volatility remained over 15% for most of the remainder of 1907. The true financial crisis, with bank failures and J.P. Morgan's famous intervention, occurred in October: volatility did stay over 20% during this time, but never again reached 30%.

The Great Depression followed a similar pattern, though over a longer timeframe. The initial surprise was of course October 28, 1929, when the market fell 12.8%. Volatility was already nearly 35%, but the crash still registered as a large (over six standard deviations) surprise. Though volatility did rise to over 90%, it had fallen back into the twenties by early 1930. In 1931 the real volatility returned, beginning

the longest stretch of sustained high volatility in history. Interestingly, large surprises were absent in this period: the most surprising loss was a 2.8-standard deviation event in September 1931.

Like these other crises, the current one seemed to have begun with a surprise—the 7.8-standard deviation loss in February 2007—though this surprise came in a period of particularly low volatility. The loss was triggered primarily by slowing demand from China, but also served as a symbolic beginning to the subprime crisis, with the troubles at the Bear Stearns hedge funds surfacing shortly thereafter. Still, volatility remained below 25%, even through the turbulence of August 2007. It was only after the failure of Lehman Brothers in September 2008 that volatility rose above 30%. Since February 2007, however, despite all that has happened and the historic run-up in volatility, there have been no large surprises: the largest was the fall on September 29, 2008, the day the US Congress rejected the first bank bailout plan. This loss was one of the twenty largest ever, yet registered as only a 3.7-standard deviation event amid the already high volatility.

We make similar conclusions examining weeks rather than days. As residuals, consider weekly returns relative to the volatility forecast at the beginning of the week.⁶ By this standard, the greatest surprise of the recent turbulence was the first week of October, whose 18% loss was a 3.4-standard deviation drop; the last week of February 2007 produced a 4.2-standard deviation fall. No other week in the last two years has been as much as a three-standard deviation surprise. The last five-standard deviation (weekly) surprise came in 1946.

⁶(Recall that volatility forecasting over this time horizon relies on our assumption that surprises do not cluster.)

Other warnings

With all of this discussion of surprises, it is important to ask whether the surprises are a result of our expectations (that is, our forecasts) missing something, rather than the fates simply acting. We should not pretend that a forecast from historical data is the only view of future market volatility, and ask then whether the options markets have given fair warning, especially in cases where our historical models were surprised. To this end, we turn to the VIX index.

The VIX is essentially a composite of the volatility implicit in the prices of options referencing the S&P 500, with roughly one month to maturity, over a variety of strike prices. The index is constructed to represent the market's implicit expectation of the volatility to be experienced (that is, the square root of the sum of the actual squared returns) over the next month.

Though they introduced futures on the VIX only in 2004, the CBOE began publishing the index in 1993, and has reconstructed the history of the index back to 1986. In fact, it is the new version of the index that references the S&P 500. We focus on the old version of the VIX,⁷ which referenced the S&P 100, and for which we have the longest historical data sample.

To assess whether the VIX can complement our forecasts, we examine its behavior just before and just after a sample of surprise days. We apply the same methodology as before: compute volatility forecasts on the S&P 100 using the historical data, and divide each day's return by the volatility we would have forecast the day before. Over the period for which we have VIX data (6645 days), we observe thirteen days

on which the actual loss is greater than four standard deviations;⁸ we choose these as our surprise days.

For a selection of these surprise days (plus the September 2001 event, which was a 3.9-standard deviation event for the S&P 100), we examine the evolution of the volatility (both our historical estimate and the VIX level) from four weeks before the surprise until two weeks after. Since we are mostly interested in how volatility moves, we normalize by the volatility on the surprise day, that is, the last volatility forecasts we would have made before the surprise. We plot the evolution of this normalized volatility for our seven surprise days in Figure 2.


The evolution of our volatility forecast is on the left of the figure. Since we have chosen days on which this forecast was surprised, it is not shocking that the forecasts in general did not increase markedly prior to the event. Exceptions are the 1987 crash, where the forecast doubled in the week prior to the event, and the 2008 event, where volatility increased markedly from two weeks beforehand. After the event, the forecasts shot up significantly, with the smallest reactions in the cases of September 2001 and the most recent event. In the events where volatility reacted most, it then decayed, while in the 2001 and 2008 cases, volatility stayed elevated.

On the right of the figure, we display the evolution of the implied volatility for the same events. Again, most of the violent reactions come after the events, and again, there was a marked increase in the VIX prior to the 1987 and 2008 events. Curiously, there was also an increase in the VIX in the 2001 case (fodder perhaps for someone's conspiracy theories).

⁷Whose index values are now associated with the ticker symbol VXO

⁸Slightly fewer than the t -distribution would predict

Figure 2: Evolution of historical volatility forecast (left) and VIX volatility index (right) around selected surprise days for the S&P 100



Source: CBOE website, Yahoo Finance, own calculations

After the events, implied volatility jumped, though 1987 stands out for its reaction. From there, volatility decayed for the 1987, stabilized in most of the other cases but continued to increase after the 2008 event, as the market continued to deteriorate.

So what now?

From the historical perspective, we see two ways in which volatility has risen in the past. At times, volatility has spiked, typically as the result of a surprise. When the surprise comes at a time of already elevated volatility, volatility rises to historically high levels, but then comes back down relatively quickly.

In other cases, volatility has risen in an orderly way, with no true surprises. The run-up in volatility in 1931 is the best example of this phenomenon, and in

that case, volatility stayed elevated for quite a long time: it spent more sixteen months over 35%, during which time the index fell by 50%.

Somewhat concerning is that the next best example of this orderly run-up in volatility is what we have just seen, with no large surprises but volatility now having spent (as of this writing) 29 consecutive trading days over 35% and 21 days over 50%. There have been few runs in history as long of such high volatility. Though not all such runs coincided with market losses, the market's performance over such periods has been generally poor.

The pattern in the VIX has also been different for the 2008 crisis, with a sustained increase without precedent as far back as the VIX has been calculated. It is not often that we take solace in a lack of data, but we are relieved not to be able to make any more comparisons with earlier periods.

Further reading

- Chicago Board Options Exchange (2003). VIX White Paper. Available at www.cboe.com/micro/vix
- Wikipedia—Dow Jones Industrial Average, Panic of 1907
- Zumbach, G. (2006). The RiskMetrics 2006 Methodology. RiskMetrics Group. Available at www.riskmetrics.com/publications/techdoc.html