



Research Brief

March 2016

Are markets becoming more unstable?

Readers of financial news may believe that ‘market corrections’, or ‘shocks’, or ‘five-sigma events’ are more common than they used to be. We look at the historical data for a number of financial markets and find that there is no evidence for increasing instability over the past 60 years.

It’s a cliché that this is no longer “your grandfather’s market”, and we tend to think that our grandparents had a quieter time. We remember recent crises such as the turmoil of August 2015, reactions to central bank decisions and the unpegging of the Swiss franc from the euro. Often high-frequency traders, or other algorithmic investors, are blamed for increasing instability¹. Has there really been an increase in the frequency of these events, or are we allowing the fresh, painful memory of recent shocks to influence us too much?

Our data suggest that financial markets have *not* become more unstable. Figure 1 shows the cumulative count of extreme events (‘shocks’) across a range of markets over 60 years. The increase of the count over any period measures the number of shocks *per market* over that time, so the rate of increase can be compared across decades as more markets become available.

A shock is defined as a daily return so large that, if returns were normally distributed, it would occur less than once in a thousand years. This is roughly equivalent to a ‘five-sigma event’. The top chart shows results across 25 markets, using two different periods, 33-day and five-year, to estimate the volatility. If market instability were increasing, we would see the lines curving upwards with the cumulative count of shocks increasing more quickly over time. The bottom chart shows shocks per market in each of the nine sectors.

Because the five-year volatility is an average over a long period, it tends to lag behind when a market becomes more volatile, leading to more events labelled as shocks. The financial crisis of 2008 is the most extreme example: volatility increased suddenly in many financial markets across the world. But there is no evidence of a general trend of increasing instability over time, whichever volatility measure we use. In this research brief, we describe why we use this measure of instability and how it is constructed.

¹ For example: Gavin Jackson, *Financial Times* 14 September 2015, ‘Why market volatility is growing more intense’; D. Sornette and S. von der Becke, ‘Crashes and high-frequency trading’, UK government Foresight report August 2011.

Enquiries

researchpapers@wintoncapital.com
+44(0)20 8576 5800

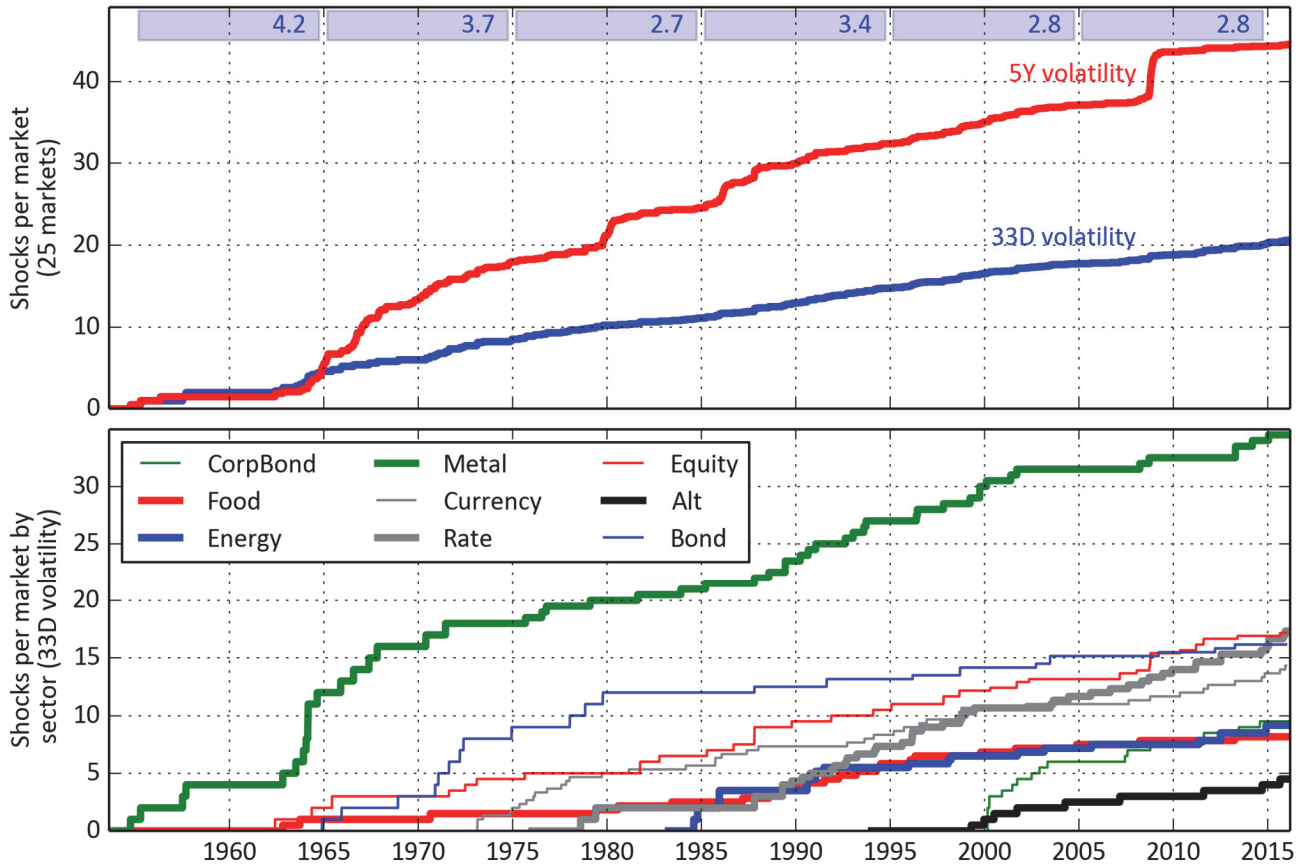


Figure 1: Shocks per market across 25 markets in 9 sectors. Each curve increases by $1/n$ when a shock occurs for a single market in a sector containing n markets (or for the global total, when data for n markets are available). Shocks are daily returns that are so large compared to the volatility that they would occur once in 1000 years for a market with normally distributed returns. Totals across all sectors are shown in the top chart using 33-day and five-year volatilities. Numbers in blue boxes are mean shocks per market over each ten-year period across all sectors using 33-day volatility. The 'Alternatives' sector (Alt) includes private equity and real estate. The curve for each sector in the bottom chart starts when daily data from the first market becomes available and is calculated using 33-day volatility. Markets and sectors are defined in Appendix 1.



Figure 2: Futures price for the EUR/USD exchange rate in autumn 2015. The increase on 3 December was particularly extreme against the background of recent returns.

What is a shock?

Defining a shock is not simple because there is no fixed threshold which a market must pass for its performance to be shocking. For example, a daily gain or loss of 1% is unexceptional in the equity markets. In short-term US interest rate ('Eurodollar') futures, it would be a truly shocking 67-sigma event. It is more useful to compare a day's return with the size of previous returns, measured using the volatility.

A good example is the 4% increase in the EUR/USD exchange rate on December 3 last year. It was particularly damaging for trend followers because it came after a month of low volatility accompanied by a downwards trend (Figure 2).

To assess the shock value of a daily return, we divide the return by the volatility over the last few weeks, to get 'volatility-adjusted' returns². For euro futures, December 3 was exceptional by this measure. It appears on the right-hand end of the histogram of volatility-adjusted returns for the last fifteen years, as a 'five-sigma' event (Figure 3).

We don't find five-sigma events particularly shocking, because it is common knowledge that market returns are not normally distributed. If they were, we would expect one event like this in a thousand years. In fact, for euro futures, there have been two in the last fifteen years.

Across a range of markets, we will define a 'shock' as a day with a volatility-adjusted return (positive or negative) so large that we would expect it to happen only once in a thousand years in a normally distributed market. This

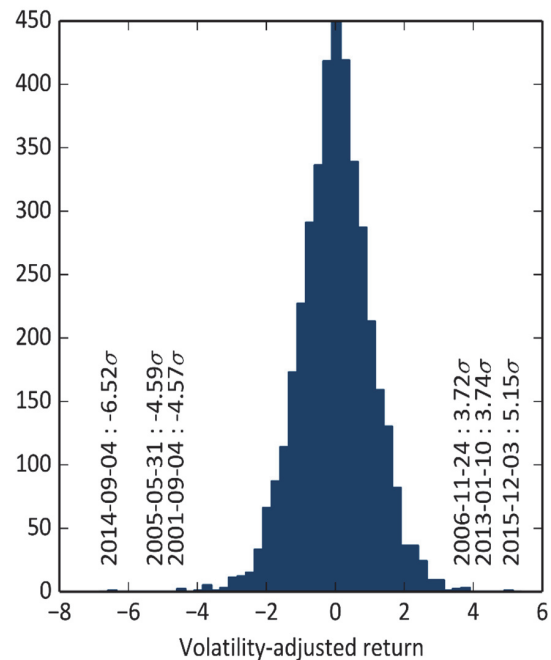


Figure 3: Daily returns of euro futures, 2000-2015, divided by the previously estimated 33-day volatility. The three largest negative and positive returns are labelled.

corresponds to a 4.97σ event when we use a 33-day volatility³.

We are interested in understanding whether the frequency of these extreme events has changed over time. In other words, have markets become more unstable?

The S&P 500

We first look at the S&P 500⁴, using futures returns where they are available and the excess returns from holding the index stocks before the futures market started. The

² We use '33-day volatility' to mean the rolling standard deviation of daily futures returns using an exponentially weighted moving average, where the centre of mass of the decay function is at 33 days. Volatility adjustment is done by dividing by the *previous* day's volatility, so we assess the return against our knowledge of the market before that day. The five-year volatility introduced later is similarly defined (again using daily returns).

³ We might have expected a 4.62σ event to occur once in a thousand years for normal data, since allowing for 262 business days in a year, we find $2\Phi(-4.62) \times 262 \approx 10^{-3}$, where Φ is

the normal cumulative distribution function. But if we estimate the volatility using a 33-day moving average the frequency of recorded shocks is a little higher than we expect because of fluctuations in the estimated volatility. For a five-year volatility the frequency is closer to the expected one.

⁴ For financial assets such as stock indices, the futures returns are very close to the excess total returns from holding the asset ('excess' meaning over the risk-free rate, and 'total' meaning including dividends for stocks and coupon payments for bonds). We therefore use futures returns back to the start of futures trading, and the excess returns from holding the index assets before this date.

curves in Figure 4 move up by one step every time a shock is observed.

The results depend on the period used to measure volatility. If we use a 33-day volatility, as in Figure 3, then we see some quiet periods, and other periods where shocks happen about once per year. The last decade is one of these active periods, but no more active than the 1990s or the early 1960s. If we use a longer volatility period, then the financial crisis of 2008 stands out. The short-period volatility reacts more quickly to changes in the market and gives a better estimate of the size of future returns than the long-period volatility. This is also the reason that the five-year measure records more shocks overall.

Either way, there is no evidence that recent years (except 2008) have been exceptionally turbulent. But over 50 years there are at least fourteen shocks, so the frequency is much higher than the one shock per thousand years we would expect from normal data.⁵

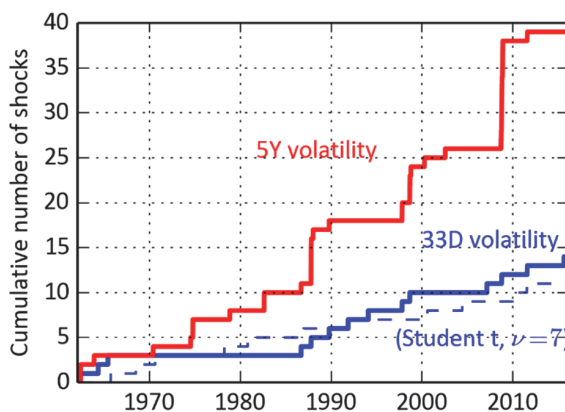


Figure 4: Running total of shocks, defined relative to five-year and 33-day volatility, for the S&P 500. Each curve increases by 1 when a shock occurs. A shock is defined as a volatility-adjusted return large enough to occur only once in a thousand years for normally distributed returns. For comparison, we also show (dashed) the total for simulated returns with a student's t distribution having 7 degrees of freedom (using 33-day volatility).

⁵ In fact, a t-distribution with 5-10 degrees of freedom is a better model of the returns of real markets, as shown in Figure 4.

Across markets and sectors

For financial assets like stocks, futures returns are very close to the excess total returns (including dividends, or coupon payments for bonds) of the underlying asset. 'Excess' means the returns over and above the risk-free rate. For real assets such as gold or oil, futures markets are generally the most liquid way to take positions, either to invest or hedge. To get a general view of investable assets, we repeat the calculation with 21 futures markets across seven sectors. For four other asset classes where no futures markets exist, we calculate daily excess returns from a total return index. The details are given in Appendix 1.

We want a view of the frequency of shocks across a wide range of financial markets and over a long period of time. Data for some markets does not extend back as far as for others. We therefore plot the cumulative count of shocks per market. So if a shock occurs in one market during a period when there are ten markets, we increment the counter by 1/10. The results are shown in Figure 1.

Across all nine sectors, and for the grand total, we see the same picture as for the S&P 500. There has been no increase in instability across the decades. Some sectors have been more prone to shocks than others. Short-term interest rates average one shock every three years per market, while government bonds are particularly quiet since 1980, averaging nine years between shocks.

These conclusions are quite robust: they don't change if we use a different volatility measure, or use quarterly returns rather than daily ones, except for the difference we noted earlier for the S&P 500: the 2008 financial crisis stands out if we use a longer timescale to measure volatility.

Markets may feel more unstable than they used to be, but 'five-sigma events' are a regular feature of financial returns. There is no evidence that crashes or corrections happen more often than they did in previous decades.

Appendix 1: Markets and sectors

Sector	Markets
Equity	S&P 500, Dow Jones EuroStoxx, Nikkei 100, NASDAQ Composite
Bond	10Y US Treasury note, German Bunds, Japanese 10Y
Currency	EUR/USD, GBP/USD, JPY/USD
Energy	WTI crude oil, RBOB unleaded gasoline, natural gas
Food	Coffee ('C'), Wheat (Chicago), Sugar (11)
Metals	Gold, Copper (London Metal Exchange)
Rates	USD LIBOR 3-month ('Eurodollar'), euro LIBOR 3-month, yen TIBOR 3-month
Corporate bonds*	US AAA, US high-yield (Bloomberg)
Alternatives*	US real estate (FTSE), US private equity (LPX50).

(*: daily excess returns calculated using total asset returns and US treasury-bill rate. Other markets' excess returns are calculated from back-adjusted futures returns, generally using the 'front' contract (closest to expiry) and rolling to the next a few days or weeks before expiry.)

Appendix 2: Changes in volatility

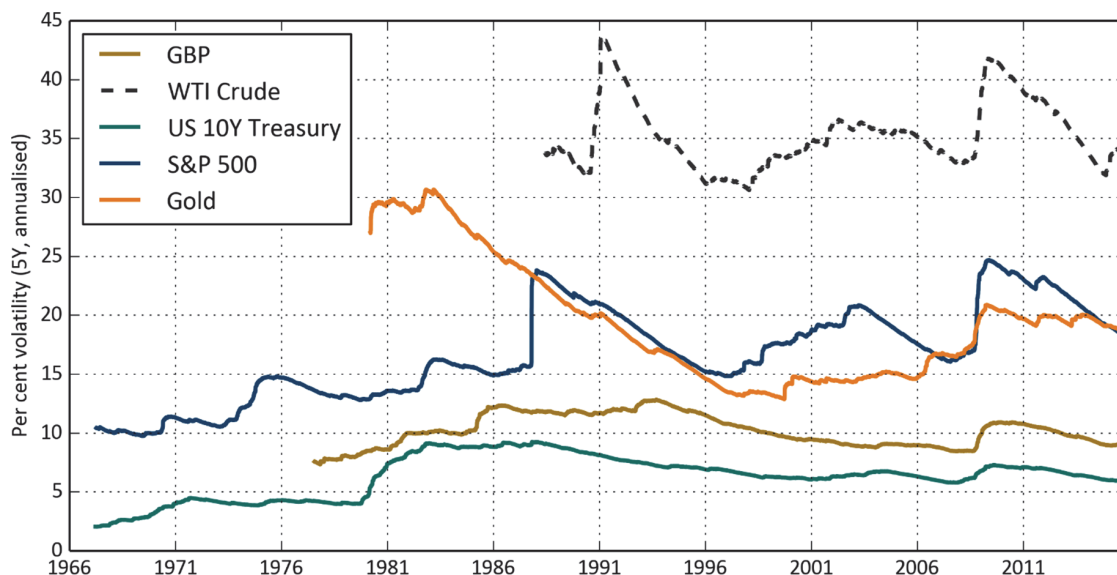


Figure 5: Changes in volatility over time. For five markets (chosen for their long history), we show the volatility, calculated as an exponentially weighted average with a five-year period, using daily returns in futures markets, and (for stocks and bonds) using excess returns from holding the asset before the futures markets were opened.

LEGAL DISCLAIMER

This document is communicated by Winton Capital Management Limited (“WCM” together with its affiliated companies, “Winton”), which is authorised and regulated by the UK Financial Conduct Authority, is registered as an investment adviser with the US Securities and Exchange Commission (“SEC”), is registered with the US Commodity Futures Trading Commission and is a member of the US National Futures Association.

The information herein does not constitute an offer to sell or the solicitation of any offer to buy any securities in a fund established, advised or managed by Winton.

No guarantee or representation is made that an Investment Program will achieve its investment objective and the value of investments made in accordance with an Investment Program may go down as well as up. Past performance is not indicative of future results. This document may contain simulated or hypothetical performance results that have certain inherent limitations. Unlike the results shown in an actual performance record, these results do not represent actual trading.

The information herein may be approximate, is subject to updating and further verification and may be amended at any time and WCM is under no obligation to provide you with an updated version. The information contained in this document is believed to be materially correct but WCM makes no representation or warranty as to its accuracy or completeness and accepts no liability for any inaccuracy or omission. Information obtained from third parties is believed to be reliable but has not been independently verified by WCM.

The MSCI information may only be used for your internal use, may not be reproduced or disseminated in any form and may not be used as a basis for or a component of any financial instruments or products or indices. None of the MSCI information is intended to constitute investment advice or a recommendation to make (or refrain from making) any kind of investment decision and may not be relied on as such. Historical data and analysis should not be taken as an indication or guarantee of any future performance analysis, forecast or prediction. The MSCI information is provided on an “as is” basis and the user of this information assumes the entire risk of any use made of this information. MSCI, each of its affiliates and each other person involved in or related to compiling, computing or creating any MSCI information (collectively, the “MSCI Parties”) expressly disclaims all warranties (including, without limitation, any warranties of originality, accuracy, completeness, timeliness, non-infringement, merchantability and fitness for a particular purpose) with respect to this information. Without limiting any of the foregoing, in no event shall any MSCI Party have any liability for any direct, indirect, special, incidental, punitive, consequential (including, without limitation, lost profits) or any other damages. (www.msci.com).

The S&P 500 Index is a product of S&P Dow Jones Indices LLC (“SPDJI”), and has been licensed for use by Winton Capital Management Limited. Standard & Poor’s®, S&P® and S&P 500® are registered trademarks of Standard & Poor’s Financial Services LLC (“S&P”); Dow Jones® is a registered trademark of Dow Jones Trademark Holdings LLC (“Dow Jones”); and these trademarks have been licensed for use by SPDJI and sublicensed for certain purposes by Winton Capital Management Limited. Winton Capital Management Limited’s product is not sponsored, endorsed, sold or promoted by SPDJI, Dow Jones, S&P, their respective affiliates, and none of such parties make any representation regarding the advisability of investing in such product(s) nor do they have any liability for any errors, omissions, or interruptions of the S&P 500 Index.
