



Title	How unlucky is 25-Sigma?
Authors(s)	Dowd, Kevin, Cotter, John, Humphrey, Christopher, Woods, Margaret
Publication date	2008
Publication information	Dowd, Kevin, John Cotter, Christopher Humphrey, and Margaret Woods. "How Unlucky Is 25-Sigma?" Euromoney Institutional Investor, 2008. https://doi.org/10.3905/jpm.2008.709984 .
Publisher	Euromoney Institutional Investor
Item record/more information	http://hdl.handle.net/10197/1625
Publisher's version (DOI)	10.3905/jpm.2008.709984

Downloaded 2026-05-02 00:27:29

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

How Unlucky is 25-Sigma?

KEVIN DOWD, JOHN COTTER, CHRIS HUMPHREY,
AND MARGARET WOODS

KEVIN DOWD

is a professor of financial risk management at the Centre for Risk and Insurance Studies, Nottingham University Business School in Nottingham, UK.

kevin.dowd@nottingham.ac.uk

JOHN COTTER

is an associate professor and the director of the Centre for Financial Markets at the School of Business, University College Dublin in Dublin, IE.

john.cotter@ucd.ie

CHRIS HUMPHREY

is a professor of accounting in the School of Accounting and Finance at the University of Manchester in Manchester, UK.

chris.humphrey@mbs.ac.uk

MARGARET WOODS

is an associate professor in accounting and finance at the Centre for Risk and Insurance Studies, Nottingham University Business School, in Nottingham, UK.

margaret.woods@nottingham.ac.uk

One of the more memorable moments of last summer's credit crunch came when the chief financial officer of Goldman Sachs, David Viniar, announced in August that Goldman's flagship GEO hedge fund had lost 27% of its value since the start of the year. As Mr. Viniar explained, "We were seeing things that were 25-standard deviation moves, several days in a row (Larsen [2007])." One commentator wryly noted:

That Viniar. What a comic. According to Goldman's mathematical models, August, Year of Our Lord 2007, was a very special month. Things were happening that were only supposed to happen once in every 100,000 years. Either that...or Goldman's models were wrong (Bonner [2007b]).

But sadly Goldman was not alone. In 2007, massive losses were announced by Bear Stearns, UBS, Merrill Lynch, and Citigroup, and in early 2008, Société Générale and Bear Stearns, once again, were making headlines, with rumors of more losses yet to come. Of course, these followed on the heels of the earlier financial disasters—1987, Daiwa, Barings, Long-Term Capital Management, the dotcoms, Russia, East Asia, and so on.

Citi's case was particularly interesting. To quote from the same commentator:

Gary Crittenden, Citi's chief financial officer, claimed...that the firm was simply a victim of unforeseen events... No mention was made of the previous five years, when Citi was busily consolidating mortgage debt from people who weren't going to repay..., pronouncing it 'investment grade'..., monogering it to its clients..., and stuffing it into its own portfolio...while paying itself billions in fees and bonuses. No, according to the masters of the universe, downgrades by Moody's and Fitch's were completely unexpected..., like the eruption of Vesuvius; even the gods were caught off guard. Apparently, as of September 30th, Citigroup's subprime portfolio was worth every penny of the \$55 billion that Citi's models said it was worth. Then, whoa, in came one of those 25-sigma events. Citi was whacked by a once-in-a-blue-moon fat tail. Who could have seen that coming (Bonner [2007c])?

Be this as it may, one thing is for sure—there are certainly a lot of very unlucky financial institutions around.

HOW UNLIKELY IS A 25-SIGMA EVENT?

The once-in-a-100,000-year figure was quoted in a number of places and suggests that Goldman, Citi, and others must have been *very* unlucky indeed.¹ But *exactly* how unlikely is a 25-sigma shock?²

To begin, let's assume that losses are normally distributed, that is, that losses obey the classic bell curve, then ask the question: What is the probability of a loss that is, say, two standard deviations or more away from the mean (i.e., what is the probability of a 2-sigma loss event)?

The answer is given in Exhibit 1. The probability associated with a 2-sigma event is equal to the mass of the right tail of the distribution demarcated at the point where the number of standard deviations from the mean is equal to 2; this probability is equal to 2.275%.³ We might therefore expect to see a 2-sigma loss event on one trading day out of $1/2.275\% = 43.956$ days (i.e., on approximately 1 day out of 44 days). A 2-sigma event is unlikely to occur on any given day, but we would expect to see a few of them in any given year.⁴

We now use the same approach to estimate the probabilities of losses that are 3 sigmas or more away from the mean, 4 sigmas or more away from the mean, and so on. The results of this exercise for 3-, 4-, 5-, 6-, and 7-sigma events are given in Exhibit 2.

The reader will note, based on the following, as k gets bigger the probabilities of a k -sigma event fall extremely rapidly:

- A 3-sigma event is to be expected about every 741 days or about 1 trading day in every three years.
- A 4-sigma event is to be expected about every 31,560 days or about 1 trading day in 126 years (!).
- A 5-sigma event is to be expected every 3,483,046 days or about 1 day every 13,932 years (!!).
- A 6-sigma event is to be expected every 1,009,976,678 days or about 1 day every 4,039,906 years.
- A 7-sigma event is to be expected every $7.76e+11$ days. The number of zero digits is so large that Excel now reports the number of days using scientific notation; this number is to be interpreted as 7.76 days with the decimal point pushed back 11 places. This frequency corresponds to 1 day in 3,105,395,365 years.

These results are breathtaking. To give them some perspective, consider the following: a 5-sigma event corresponds to an expected occurrence of less than just one day in the entire period since the end of the last Ice Age; a 6-sigma event corresponds to an expected occurrence of less than one day in the entire period since our species, Homo Sapiens, evolved from earlier primates; and a 7-sigma event corresponds to an expected occurrence of just once in a period approximately five times the length of time that has elapsed since multi-cellular life first evolved on this planet.⁵

At this point, there are so many decimal points in the numbers involved that Excel is unable to handle values of k bigger than 7, and we are still a very long way from the 25-sigma events that Goldman Sachs experienced. To overcome the limitations of Excel, we now switch to MATLAB and use the MATLAB command “1-normcdf(8,0,1)” to estimate the corresponding probability of an 8-sigma event, which corresponds to an expected occurrence once every $6.429e+012$ years. To put this into perspective, it is a period (considerably) longer than the entire period that has elapsed since the Big Bang.⁶ If we observe a profit or loss once a day, then a mere 8-sigma event should occur less than once in the entire history of the universe.

We then moved to the analysis of a 9-sigma event, but ran into a problem because the MATLAB “normcdf” function gives 9 sigma and larger events a flat value of zero; the probabilities are so small they now fall under the function's radar. To get around this problem, we wrote a specially designed MATLAB function to estimate the

EXHIBIT 1 Probability of a 2-Sigma Event

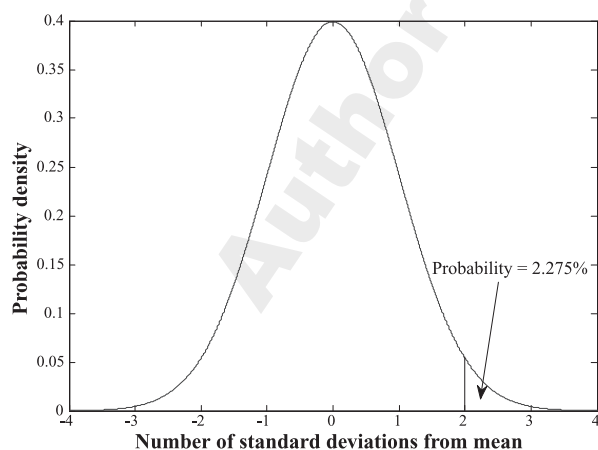


EXHIBIT 2**Probabilities of k -sigma events: $k = 3, 4, 5, 6$ and 7**

K	Probability in Any Given Day	Expected Occurrence: Once in Every:	
3	0.135%	740.8	Days
4	0.00317%	31,559.6	Days
5	0.000029%	3,483,046.3	Days
6	0.000000099%	1,009,976,678	Days
7	0.00000000129%	7.76e+11	Days

probabilities and expected-occurrence periods associated with larger losses; the details involved are reported in the appendix. We then used this function to produce the high-sigma results shown in Exhibit 3.

These numbers are on a truly cosmological scale so that a natural comparison is with the number of particles in the universe, which is believed to be between $1.0e+73$ and $1.0e+85$ (Clair [2001]). Thus, a 20-sigma event corresponds to an expected-occurrence period measured in years that is 10 times larger than the high end of the estimate range for the number of particles in the universe. A 25-sigma event corresponds to an expected-occurrence period that is equal to the high end of the estimate range, but with the decimal point moved 52 places to the left!

To give a more down-to-earth comparison, on February 29, 2008, the United Kingdom National Lottery was offering a prize of £2.5 million for a ticket costing £1. Assuming it to be a fair bet, the probability of winning the lottery on any given attempt is 0.0000004, and the probability of winning the lottery n times in a row is 0.0000004^n . Thus, the probability of a 25-sigma event is comparable to the probability of winning the lottery 21 or 22 times in a row.⁷

We should not forget Goldman's losing streak—Goldman did not experience a single 25-sigma event, but several in a row—nor forget that other institutions also experienced 25-sigma events. If the probability of a single

25-sigma event is low, the odds of two or more such events are truly infinitesimal. For example, the odds of two 25-sigma events on consecutive days are equal to $3.057e-136$ squared, which is $9.3450e-272$. This is as likely as winning the lottery about 42 times in a row. The corresponding expected-occurrence period is the square of $1.309e+135$ years—that is, $1.713e+270$ years—a number so vast that it dwarfs even cosmological figures. As Oscar Wilde might have put it, to experience a single 25-sigma event might be regarded as a misfortune, but to experience more than one does look like carelessness.

It is pretty clear by now that a 25-sigma event is much, much, *much* less likely than is suggested by an expected occurrence every 100,000 years. In fact, if we take the true figure ($1.309e+135$) and divide it by 100,000 we get a figure equal to $1.309e+130$, which tells us that the estimate of 100,000 is out by more than 130 decimal points. We suspect, therefore, that the estimate of a 25-sigma event being on a par with hell freezing over is probably about the correct characterization.

BAD LUCK OR JUST INCOMPETENCE?

However low the probabilities and however frequently 25-sigma or similar events actually occur, it is always *possible* that Goldman and the other institutions which experienced such losses were just unlucky, although

EXHIBIT 3**Probabilities of High-Sigma Events**

k	Probability in Any Given Day	Expected Occurrence: Once in Every:	
10	7.620e-022%	5.249e+020	Years
15	3.671e-049%	1.090e+048	Years
20	2.754e-087%	1.453e+086	Years
25	3.0570e-136%	1.309e+135	Years

to an extent that strains credibility.⁸ But if these institutions are really *that* unlucky, then perhaps they should not be in the business of minding other people’s money. Of course, those who are more cynical than we are might suggest an alternative explanation—namely, that Goldman and its ilk are simply incompetent. Heaven forbid!

All this poses an interesting dilemma for investors: Would investors prefer that the people looking after their money be incredibly unlucky or just plain incompetent? Maybe truth is found in both explanations. Benjamin Franklin once aptly observed that “[d]iligence is the mother of good luck,” and it stands to reason that the opposites—bad luck and incompetence—might also be related.

As for Mr. Viniar: he promised that Goldman would take a more “robust approach” in future, which will no doubt be welcome news to hard-hit investors and the millions facing foreclosure on their mortgages. Mr. Viniar further acknowledged that the experience “makes you reassess how big the extreme moves can be (Larsen [2007]).” Indeed.

Funny things, these 25-sigma events—and surprisingly common, too.

APPENDIX

Estimating the Probabilities of High-Sigma Events

Following Zeghbroeck [2007], let the Gaussian probability density be $G(x)$ and the error function be $erf(x)$. These are related via

$$\int_{-x}^x G(x)dx = erf\left(\frac{x}{\sigma\sqrt{2}}\right) \quad (A1)$$

If we set $\sigma = 1$, then x is the number of standard deviations away from the mean. Hence,

$$\int_{-x}^x G(x)dx = erf\left(\frac{x}{\sqrt{2}}\right) \quad (A2)$$

Let p be the tail probability associated with some value x . It follows from rearranging Equation (A2) that

$$p = 0.5 - 0.5 \times erf\left(\frac{x}{\sqrt{2}}\right) \quad (A3)$$

However, it is also known that for large y ,

$$erf(y) \approx 1 - \frac{e^{-y^2}}{y\sqrt{\pi}} \left(1 - \frac{1}{2y^2} + \frac{1 \times 3}{(2y^2)^2} + \frac{1 \times 3 \times 5}{(2y^2)^3} + \dots\right) \quad (A4)$$

It then follows from further rearranging that

$$p \approx \frac{e^{-y^2}}{2y\sqrt{\pi}} \left(1 - \frac{1}{2y^2} + \frac{1 \times 3}{(2y^2)^2} + \frac{1 \times 3 \times 5}{(2y^2)^3}\right) \quad (A5)$$

where $y = x / \sqrt{2}$.

The MATLAB function used to produce the high-sigma results reported in the article is reproduced below:

```
function p=high_sigma_prob(x)
% Function estimates p, the probability of a high-sigma loss event.
% Input:
% - x: size of loss in terms of number of sigmas away from mean
% Written by Kevin Dowd March 1, 2008.
*****
*****
y=x/sqrt(2);
p=0.5*(exp(-y^2)/(sqrt(pi)*y))*...
    (1-1/(2*y^2)+1*3/(2*y^2)^2+1*3*5/(2*y^2)^3);
```

ENDNOTES

¹See, for example, Bonner [2007b, c], *Randomwalk* [2007], or Hutchinson [2008]. A more colorful interpretation, although sadly one that we are unable to verify—was that a 25-sigma event was as likely as hell freezing over (Bonner [2007a]). Another commentator, Seth Jayson, suggested that a 25-sigma event is as likely as catching an asteroid in the hand (Jayson [2007]). Mr. Jayson also made another entertaining analogy, but one that we are sadly unable to repeat in a professional publication.

²A straw poll of students and colleagues at Nottingham University Business School conducted by one of the authors of this article indicated that not a single person had the remotest idea of the true probability of a 25-sigma event, and most seemed to think it was likely to occur more often than once in 100,000 years.

³For readers who wish to verify it, this number can be obtained using the command “=1-NORMSDIST(2)” in Excel.

⁴Again, in the interests of verifiability, if we go by the common market convention of 250 trading days in a year, then we would expect to see about 250/44 or 5.68 such events in any given year.

⁵The end of the last Ice Age was about 10,000 years ago, Homo Sapiens evolved within the last million years, and multi-cellular life on Earth originated about 600 million years ago.

⁶The NASA website reports that Big Bang occurred between 12 and 14 billion years ago.

⁷Note that $0.0000004^{21} = 4.3980e-135$ and $0.0000004^{22} = 1.7592e-141$, so the probability of a 25-sigma event lies somewhere between the two figures.

⁸Defenders of the “bad luck” scenario might argue that the probabilities are only as low as they are because we have assumed normality, whereas real world loss distributions are fat tailed and fat-tailed distributions lead to much higher probabilities of extreme events. This is true but irrelevant; so what if the true probability of a 25-sigma event is $3.057e-36$ rather than $3.057e-136$? Moving the decimal point a hundred places or more makes no practical difference. We would invite anyone who disagrees to do the calculations themselves.

REFERENCES

Bonner, B. “Goldman Sachs Fund Loses 30%, Wall Street Math Fails to Predict Future.” August 16, 2007a, www.dailyreckoning.com.au/wall-street-math/2007/08/16.

———. “25 Standard Deviations in a Blue Moon.” November 11, 2007b, www.moneyweek.com/file/37640/25-standard-deviations-in-a-blue-moon.html.

———. “Goldman Sachs Was Wrong and 2 Million Families May Lose Their Homes.” November 14, 2007c, www.dailyreckoning.com.au/goldman-sachs-3/2007/11/14.

Clair, B. “The Biggest Numbers in the Universe.” *Strange Horizons*, April 2, 2001, www.strangehorizons.com/2001/20010402/biggest_numbers.shtml.

Hutchinson, M. “Scrap Heap for Financial Models.” *The Bear's Lair*, February 11, 2008, www.prudentbear.com.

Jayson, S. “Goldman’s CFO: Short Truth or Long Ego.” *The Motley Fool*. August 15, 2007, www.fool.com/investing/value/2007/08/15/goldmans-cfo-short-truth-or-long-ego.aspx.

Larsen, P.T. “Goldman Pays the Price of Being Big.” *Financial Times*, August 13, 2007.

Randomwalk. “The Day the Gammas Lost Sight of the Betas and Alphas.” September 30, 2007.

Van Zeghbroeck, B. *Principles of Semiconductor Devices*, 2007. Published on the web at <http://ece-www.colorado.edu/~bart/book/book/index.html>.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@ijjournals.com or 212-224-3675