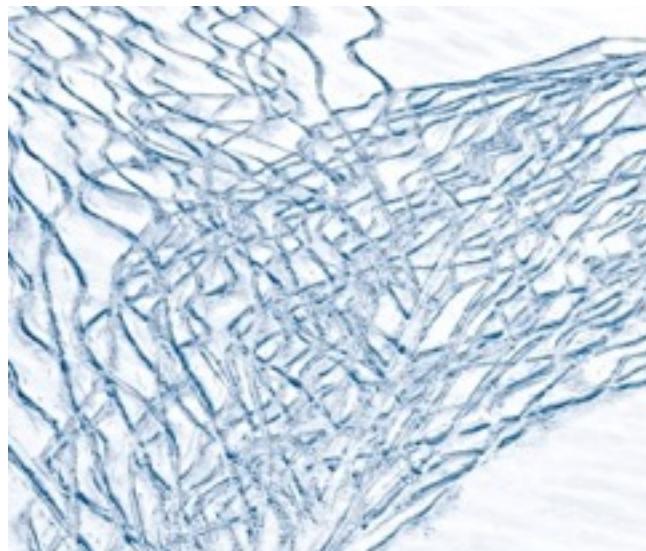


How do you
make your
paths?

A tail-focused asset path simulation technique



How do you make your paths?



An important feature of financial markets is that they are highly unpredictable. But despite market randomness, finance professionals spend considerable amounts of time and effort designing the best possible frameworks for reality. These frameworks, which help model and quantify the investment strategy, come in all shapes and sizes and are used in the investment decision process. The effectiveness of decisions taken in the light of these representations of reality will depend on how well they have been constructed.

One simple way of testing an investment strategy is to back-test it over many years of historical data. But in doing so, one assumes that the past performance is indicative of the future performance. A commonly used method is then to simulate mathematically several possible evolution paths for each asset in order to test the strategy under scenarios that differ from the past. This technique helps build performance and risk measures from each scenario and derive statistically significant data to quantify how well the strategy performs and how risky it is.

The advances in computing power have made simulations of asset paths faster. However, the relevance of the simulations depends on how realistic are the generated asset paths. One way of verifying their relevance is to check that they share some characteristics with the observed data series, such as distribution of returns or dynamics. Nowadays most finance professionals use Monte Carlo simulation methods for generating assets paths. Others recycle the range of observed returns to run Historical Simulations. We would like to discuss the pros and cons of those two simulation approaches and present a third one, Filtered Historical Simulations that overcomes some of the biggest drawbacks of Monte Carlo and Historical Simulations.

The relevance of simulations depends on how realistic are the generated asset paths

Content

Prologue	2
Monte Carlo Simulations using Wiener process	4
Historical Simulations	7
Filtered Historical Simulations	10
Conclusion	12
Appendix	13
References	14

Active Asset Allocation Research

Adina Grigoriu
Actuary Co-founder & CEO
adina.grigoriu@aaaic.com

Jacques-Yves Merelle
Risk Manager
jacques-yves.merelle@aaaic.com

Lawrence Johny
Quantitative Researcher
lawrence.johny@aaaic.com

Silvia Paris
Quantitative Researcher
silvia.paris@aaaic.com

Monte Carlo Simulations using Wiener Process

Monte Carlo simulations are repeated experiments to obtain numerical results. In the process, these simulations generate asset paths that fluctuate randomly based on the input parameters – distribution, volatility and return (drift).

In order to conduct the Monte Carlo simulations, we need a model to explain how the behavior of a random variable will change at every instant of time. One of the most common models in finance is the Wiener Process [1]. Wiener process is a basic model, which takes into account the drift and the volatility of the random variable. In this section, we consider Monte Carlo simulations using the Wiener Process.

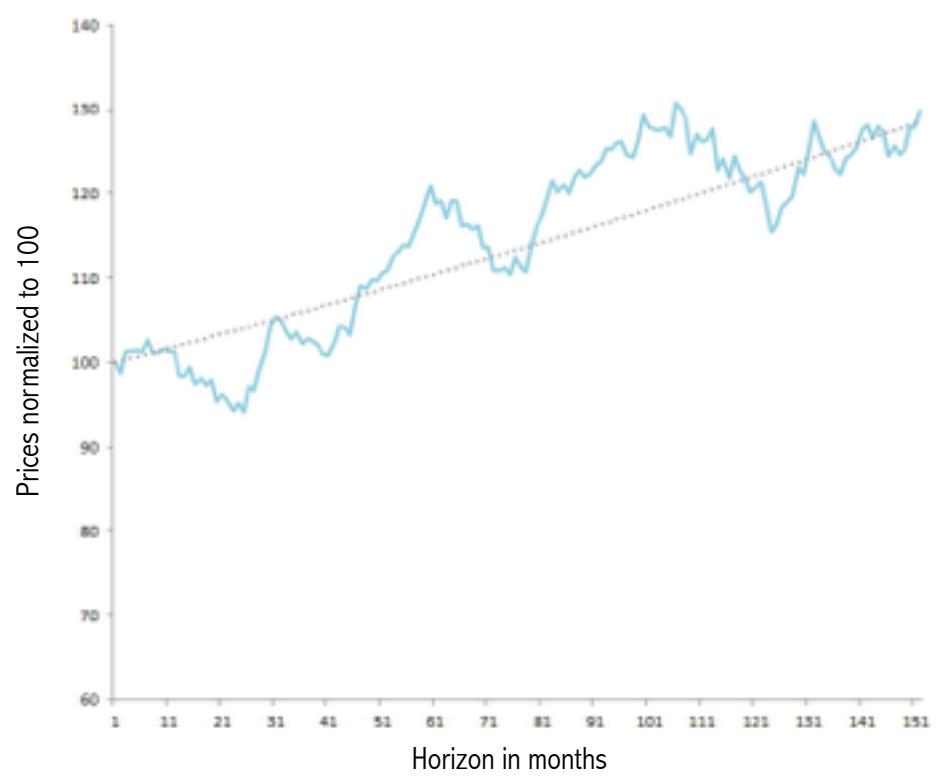
Figure 1 shows the path of an asset evolving only at the drift rate (line growing steadily) and another path evolving at the drift rate with uncertainty (line growing in a wavy fashion).

For each data point, the method draws random asset returns from a probability distribution that has been previously calibrated to fit past data series or given hypotheses.

Monte Carlo simulation method gained its popularity due to the fact that it is simple to understand and easy to implement.

But it has some important drawbacks that it is important to analyze before using it.

Fig. 1: Monte Carlo simulated path (blue line) and constant 5% drift (grey line)



[1] Hull, J. C. (2000). Options, Futures and Other Derivatives, 4th ed. Prentice Hall, NJ.

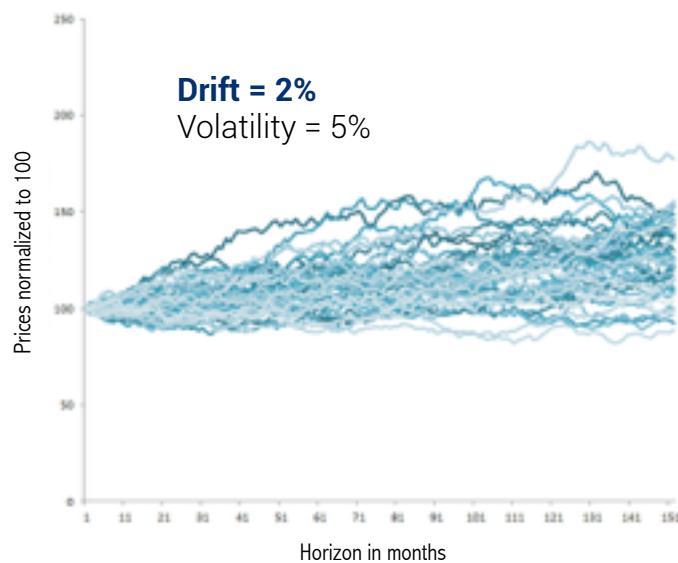
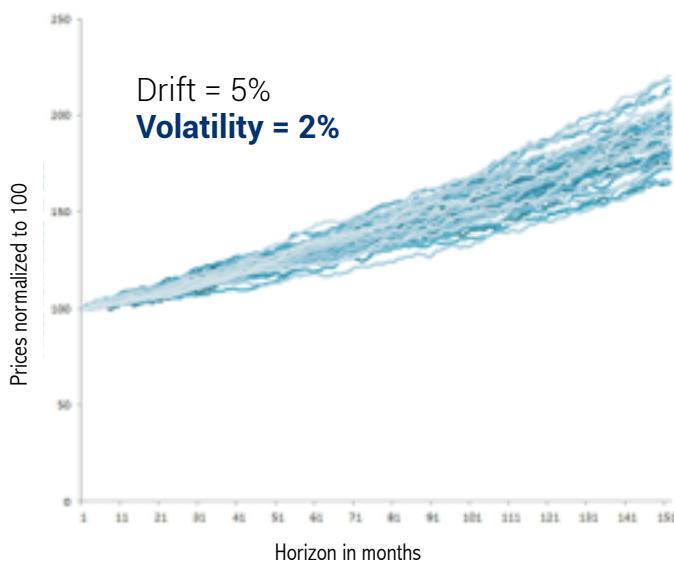
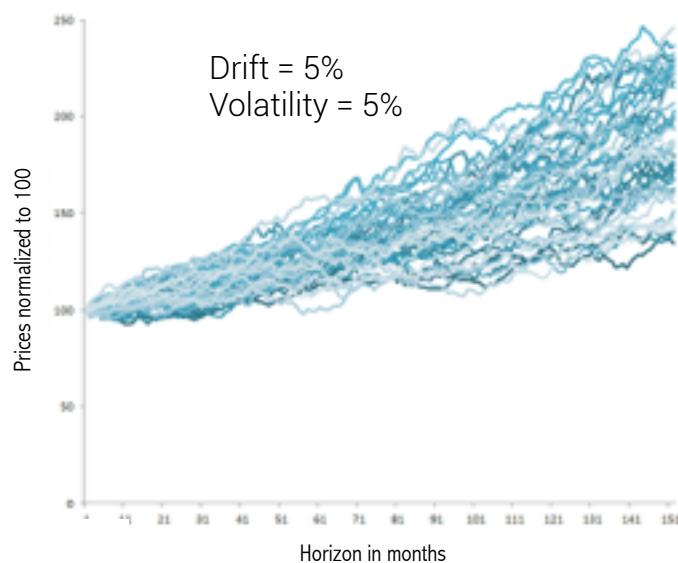
First of all, it is extremely sensitive to the value of the parameters that have been chosen: the probability distribution, the mean (which represents the price drift) and the volatility. Figure 2 illustrates this sensitivity with three sets of data simulated with different mean and volatility parameters.

Because of this sensitivity, the choice of the parameters' value is extremely subtle. Indeed making assumptions on the mean and volatility of the distribution severely biases the simulations and therefore, the conclusions of a strategy's assessment. These parameters are calibrated based on past data, or based on expectations about the future. In both cases, it is difficult to get them right. Moreover, these calibrated parameters are assumed constant in Monte Carlo simulation models, while they are not in real datasets.

Another limitation of this method is that it generates scenarios based on the assumption that asset returns are normally distributed. While this is a widely used assumption, the distribution of real financial assets returns is not normal as it has a greater density at the extremes than the normal distribution predicts.

This is not just a mathematical consideration. It actually means that the returns simulated have very limited chance to be extreme, while the past has shown that extreme moves are possible and recurrent. According to a normal distribution, the likelihood to have a stock return similar to the one observed at the peak of 2008 crisis would be roughly 1 in 33 trillion trading days [2]. This shows that none of the scenarios generated with Monte Carlo simulation methods would have integrated the possibility of a scenario that resembles the 2008 crisis.

Fig. 2: Asset paths simulated with Monte Carlo simulation, with different parameters' values



[2] Gibson et al. (2008). Black Swan Dive: Fat Tails, Stochastic Volatility and 10-Sigma Corrections, Desjardins Securities Research Comment, November 20, 2008

The last point concerning the simulation of asset paths is the importance of simulating paths for different assets at the same time, also called “joint simulations”. Monte Carlo techniques allow for the joint simulation of asset paths, given a certain correlation between the assets. The correlations are usually established using historical data or hypotheses and considered to be fixed over time. Again, this last point does not hold true in real life. Indeed, during market stresses the correlations between the different assets tend to increase dramatically. As a result when using multi-asset Monte Carlo simulated scenarios to test the robustness of a portfolio, one largely underestimates the losses that the portfolio could incur.

In order to overcome limitations of the Monte Carlo simulations arising from the assumption of normality, Efron and Tibshirani developed Historical Simulations, which aimed to preserve the original distribution of the given asset.



Historical Simulations

Developed in early 90s by Efron and Tibshirani [3], Historical Simulations are used to estimate Value at Risk (VaR). Historical Simulations eliminate the limitations of Monte Carlo simulations described earlier to a large extent by not relying on returns simulated based on a theoretical distribution. Instead it generates multiple scenarios by reordering or looping observed past returns. While different methods exist with different processes to sample the past returns, the logic remains the same: past returns represent the full range of possible future returns. Historical Simulations are sample dependent and therefore, they are realistic with respect to the sample.

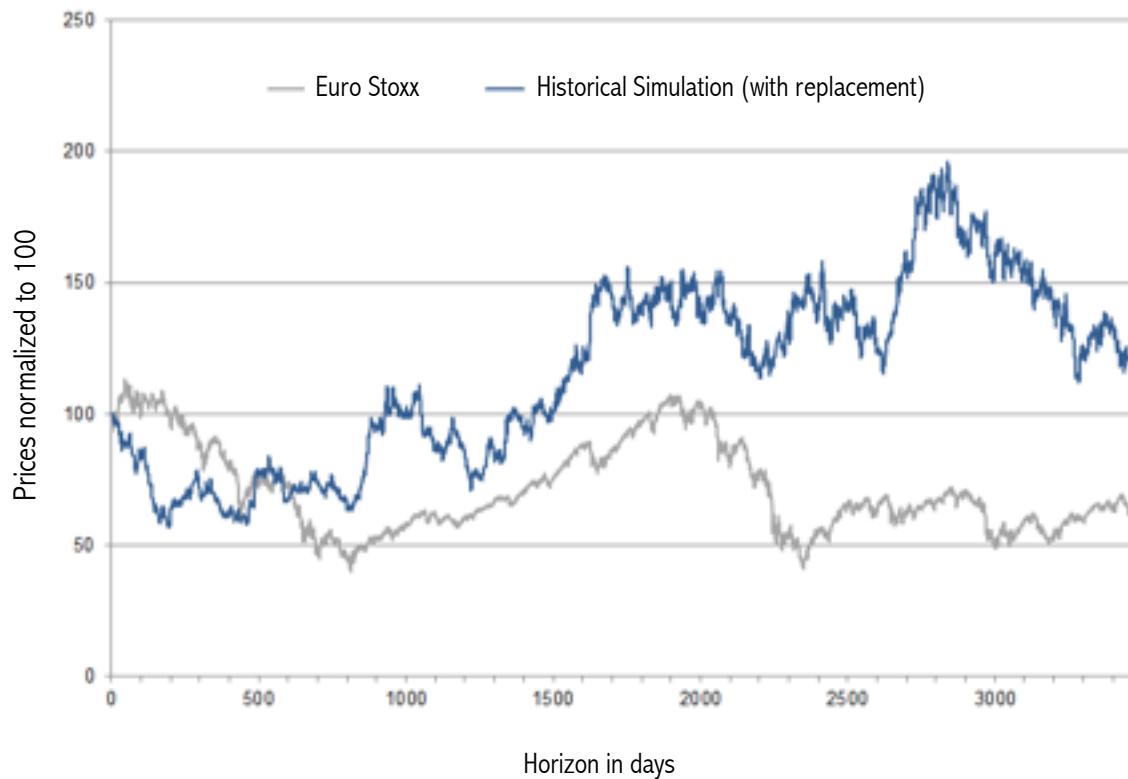
Among the different sampling methods is the “bootstrap resampling”, in which the returns are randomly reordered. The returns can be reordered either by using each return exactly once or by having the possibility of using returns more than once. Another approach in Historical Simulations is “looping time-path” method in which one can generate asset paths by looping entire segments of historical asset returns with different starting points.

Despite the advantage that Historical Simulations preserve the original distribution, one important feature missed by the bootstrap resampling method is the “volatility clustering”, the phenomenon by which “large changes tend to be followed by large changes” [3]. But, due to the randomized resampling, Historical Simulations assume that returns are independent and identically distributed (“i.i.d.” meaning that each return has no relation to past returns and that all returns have the same likelihood). It therefore ignores the volatility dynamics that exist for the returns of a given asset. A direct consequence is that simulations do not exhibit realistic periods of persistently high volatility, which leads to underestimating risks.

Quantifying this fact using statistical tests (see Appendix), it is observed that more than 90% of the scenarios simulated with historical simulations did not exhibit volatility clustering. One of such scenarios is illustrated in Fig. 3.

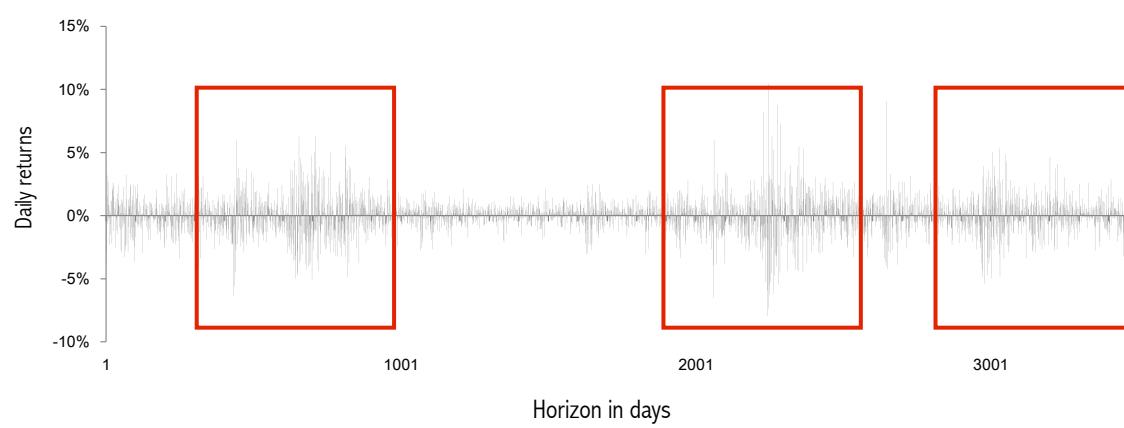
[3] Efron, B., & Tibshirani, R. (1993). An introduction to the bootstrap (Vol. 57). CRC press.

Fig. 3.a: Historical Euro Stoxx price series and a simulated scenario bootstrapped from this series



In Figure 3.a Euro Stoxx historical price series is plotted with a grey line and a simulated new price series using the bootstrap resampling method is plotted with a blue line.

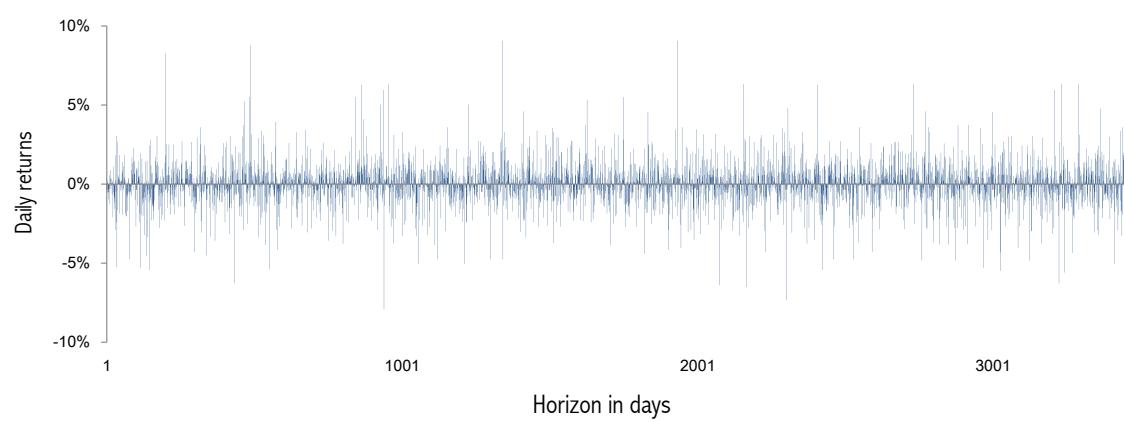
Fig. 3.b: Historical returns corresponding to fig 3.a



In Figure 3.b and 3.c the returns of each series are plotted.

The returns of the historical data are clearly clustered (which is outlined in red), while the returns from the simulated series do not respect any specific dynamic (see appendix for the measures). It is important to note that Monte Carlo simulations also suffer from this drawback.

Fig. 3.c: Historical Simulation returns corresponding to fig 3.a



While simulations generated by bootstrap resampling cannot be considered as “good” simulations because they do not respect the underlying dynamics of assets, this is not the case for the looping time-path method. It overcomes this flaw by looping entire segments of price series. A new problem appears however: it severely limits the diversity of the scenarios tested, since all the scenarios are sampled from the same blocks and they reproduce more or less the same patterns. The looping-time path method is therefore a poor test to conclude that a portfolio strategy is robust.

Lastly, another significant shortcoming that characterizes Historical Simulations is that they do not allow for price movements that have not been observed in the sampling period. All the methods of Historical Simulations limit the output range by having scenarios extracted from the historical set. Thus the simulations will not be able to consider future returns that could be higher or lower than the historically observed highs or lows. This can lead to a dramatic underestimation of the risks if the sampling period does not include periods of high stress.

Choosing between Monte Carlo and Historical simulations is mainly about choosing the approach that is less inappropriate for the utilization one wants to have. To avoid this trade-off, another approach has been developed to encompass the advantages of both techniques, known as Filtered Historical Simulations (FHS).

While most of historical simulations preserve the original distribution of returns, either they do not respect the dynamics of assets, or they severely limit the diversity of scenarios

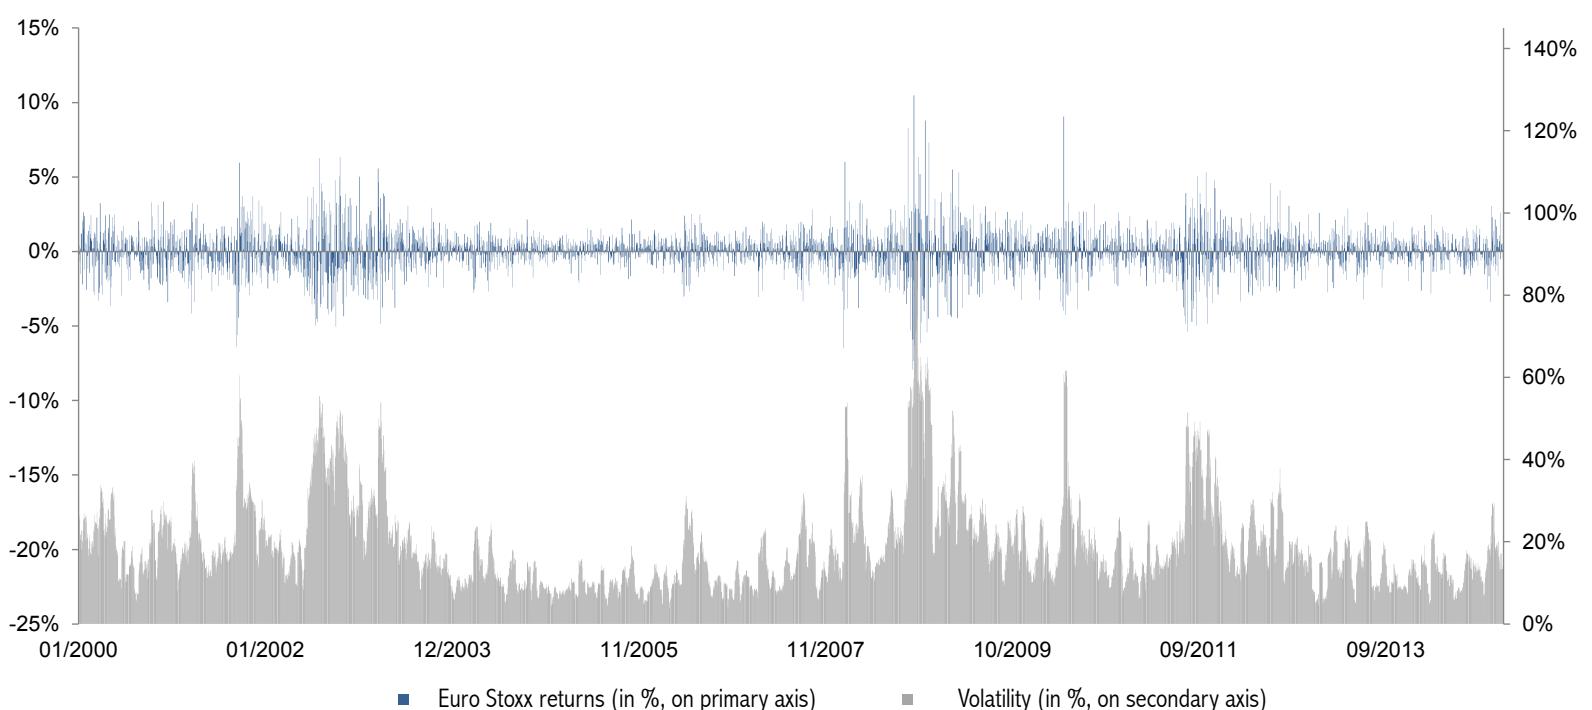
Filtered Historical Simulations

Barone-Adesi, Bourgoin and Giannopoulos (1998) [6] and Barone-Adesi, Giannopoulos and Vosper (1999, 2000) [5,7] introduced Filtered Historical Simulation (FHS) while working on a new generation of VaR models. FHS take into account the changing dynamics of past and current volatilities of historical returns of an asset (see fig. 4) and make as few assumptions as possible about the statistical properties of future returns.

In order to capture the changing dynamics of volatility, also known as heteroscedasticity, we use Generalized AutoRegressive Conditional Heteroscedasticity (GARCH) models. Nowadays one has vast amount of literature on GARCH models at disposal to help choose a GARCH model for an asset.

Fig. 4: Comparison between Eurostoxx returns and realized volatility

Figure 4 plots the return series (in blue) and the volatility of the returns series (in grey). The figure clearly illustrates that the volatility of the returns series varies dynamically over time.



[5] Barone-Adesi, G., Giannopoulos, K., & Vosper, L. (2000). Filtering historical simulation. Backtest analysis. Manuscript, March, 2000

[6] Barone-Adesi, G., Bourgoin, F., & Giannopoulos, K. (1998). Don't look back. Risk, 11, August, pp. 100-104

[7] Barone-Adesi, G., Giannopoulos, K., & Vosper, L. (1999). VaR without correlations for Non-Linear Portfolios. Journal of Futures Markets. 19, August, 583-602

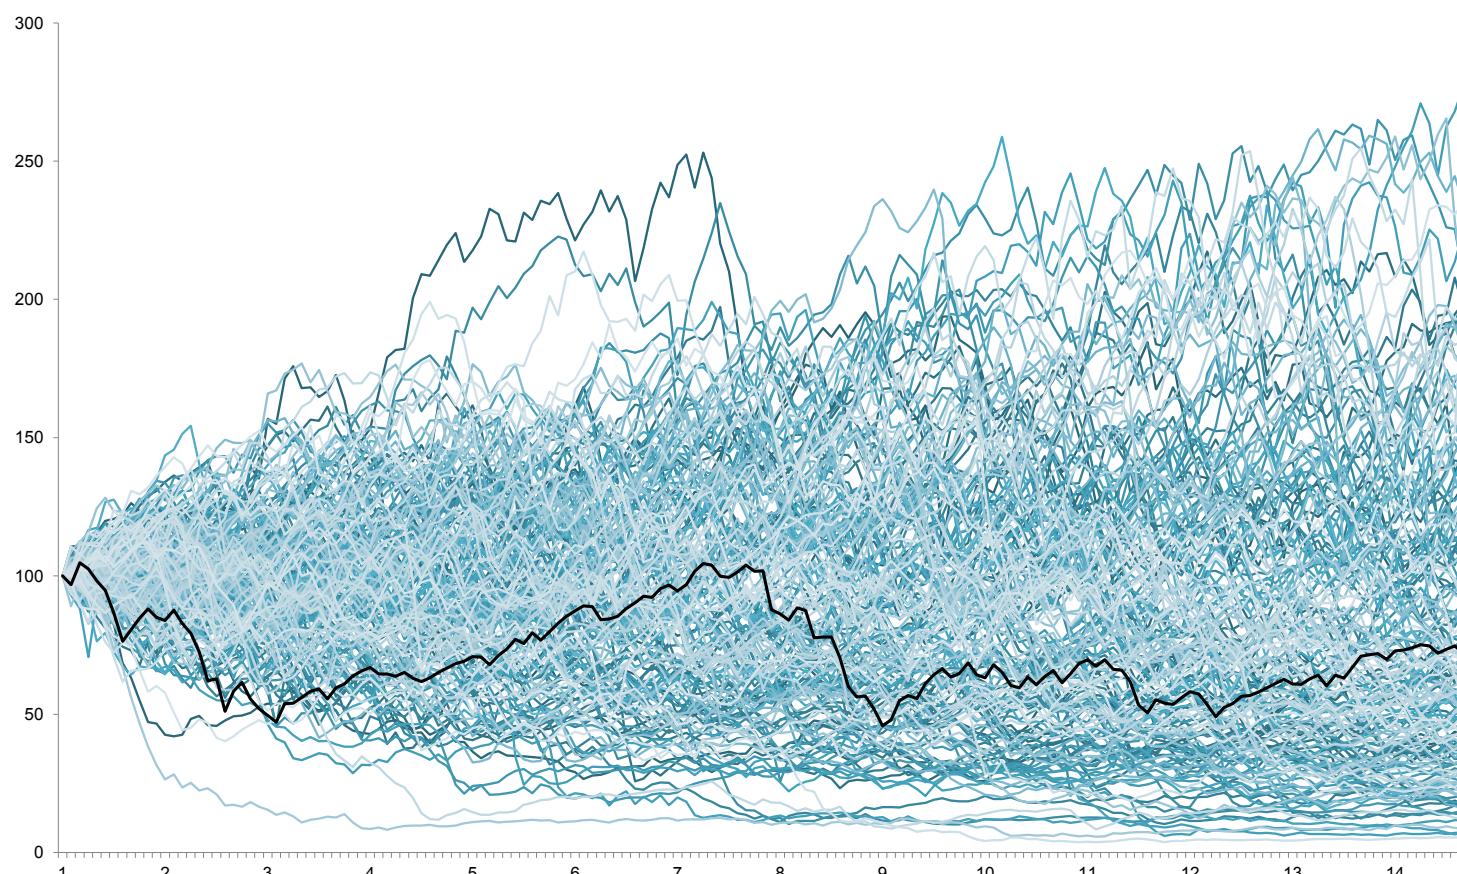
Using this literature, it is easy to choose and confirm a model for a given asset based on the statistical significance of the model parameters. Once a model is chosen for the asset, we apply it to the observed data and identify how accurately the model explains the data. Since we do not fit the model specifically to a given sample, there are differences between the observed data and that predicted by the model. These are also known as residuals. In order to use the residuals meaningfully, we standardize them by the corresponding estimates of the volatility from the model. Once standardized (adjusted to the level of current volatility), we pick them randomly using a bootstrap matrix, and then reintegrate them back into the selected model to simulate a number of new asset returns.

From a computational standpoint, even if it requires more expertise than Monte Carlo or Historical Simulations (particularly in the selection and calibration of the GARCH model), using FHS has a number of distinct advantages. First of all, the approach requires very few assumptions about the statistical properties of future price changes. This avoids relying on potentially risky assumptions that bias conclusions about a potential investment strategy.

Second, owing to the flexibility of the volatility models it uses, FHS is appropriate to simulate a great variety of assets by direct use or by a generalization of the technique: equity indexes, bond indexes, commodities, real assets, currencies, etc. Finally, FHS can be extended to joint simulations: each portfolio's asset simulated return is produced from standardized residuals extracted on the same date. It helps integrate the cross-asset dynamics and thus produce realistic co-movements. This is particularly noteworthy while assessing multi-asset strategies and large portfolios.

Using the advances in econometrics, FHS not only factor in the volatility structure and generate various returns that are not limited to past observations but also integrate scenarios that are more extreme than those the normal distribution generates.

Figure 5 plots 200 scenarios for Euro Stoxx generated using FHS for 14 years. The dark line represents the actual path of Euro Stoxx.



Conclusion

To ensure investment strategies meet their objectives, one needs to stress test them using simulated asset paths. Being a fundamental element of modeling, simulations help investors go beyond the « user defined » stress tests to make sure they can cope with market events they had not envisaged.

In this paper we compared three different approaches: Monte Carlo simulations using Wiener Process, Historical Simulations and finally Filtered Historical Simulations.

Monte Carlo simulations are easy to implement but they require input parameters that can significantly bias the scenarios they generate. Owing to the dependence on normal distribution assumption, the generated simulations exhibit few extreme moves, which underestimates asset's risks.

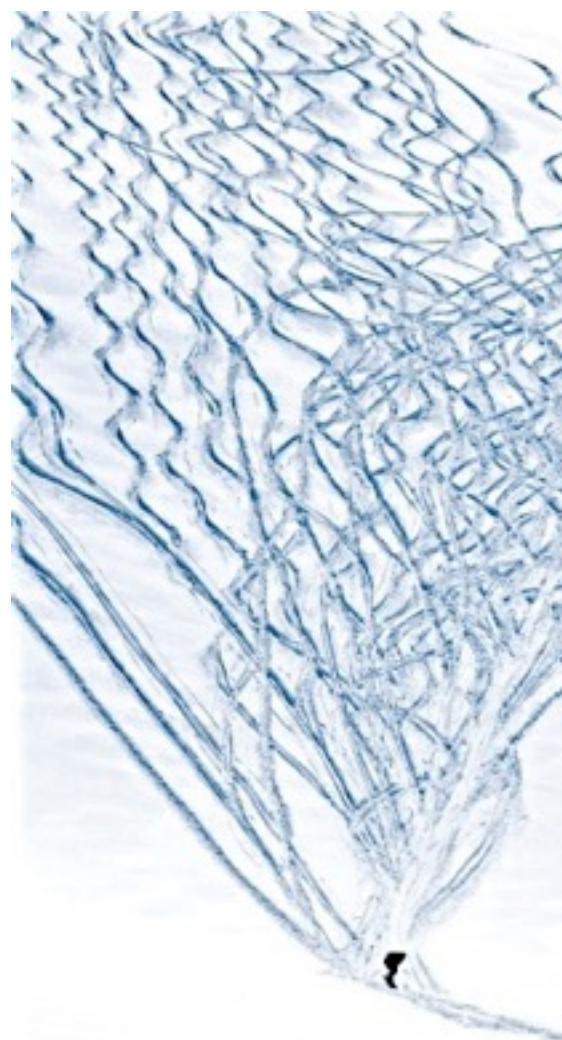
Historical simulations are designed to overcome the limitation of normality of returns by reordering past returns. While different methods exist to reorder returns, none of them manage to respect the dynamics of the assets and to generate sufficiently diverse market scenarios simultaneously.

To overcome these drawbacks, one may prefer to simulate asset paths using Filtered Historical Simulations. This modeling process has a number of distinct advantages. First, it makes as few hypotheses as possible about the statistical properties of future returns. Secondly, it generates deviations that exceed those found in the original return data.

Finally, by construction, it reproduces volatility clusters observed in empirical data and allows generating realistic cross-assets simulations.

Although less widespread than other simulation techniques, we believe that FHS could respond to the needs of the industry appropriately.

Visit active-asset-allocation.com to see how we use Filtered Historical Simulations to design customized, risk-based dynamic asset allocation solutions.



Appendix

Statistical tests employed to determine if the scenarios display volatility clustering are ARCH test and Ljung Box Q test. ARCH test assesses the null hypothesis that a series of returns exhibits no volatility clustering (conditional heteroscedasticity). Ljung-Box Q test assesses the null hypothesis that a series of returns exhibits no large changes followed by large changes (autocorrelation) for a fixed number of lags. These tests were conducted across all the scenarios generated using Historical Simulation for lags from 1 to 10, 15 and 20. In these tests, there were only a small percentage (less than 8%) of the scenarios where the null hypothesis was rejected.

Lag tested	Scenarios where null hypothesis was rejected using ARCH test	Scenarios where null hypothesis was rejected using LBQ test
1	4.09 %	5.08 %
2	5.21 %	4.97 %
3	5.89 %	5.02 %
4	6.25 %	4.93 %
5	6.65 %	4.98 %
6	7 %	5.06 %
7	7.46 %	5.1 %
8	7.68 %	5.04 %
9	7.87 %	4.96 %
10	8.04 %	5.03 %
15	8.47 %	5 %
20	8.86 %	5.21 %

References

- [1] Hull, J. C. (2000). Options, Futures and Other Derivatives, 4th ed. Prentice Hall, NJ
- [2] Gibson et al. (2008). Black Swan Dive: Fat Tails, Stochastic Volatility and 10-Sigma Corrections, Desjardins Securities Research Comment, November 20, 2008
- [3] Efron, B., & Tibshirani, R. (1993). An introduction to the bootstrap (Vol. 57). CRC press.
- [4] Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, Vol. 36, No. 4 (Oct., 1963), pp. 394-419
- [5] Barone-Adesi, G., Giannopoulos, K., & Vosper, L. (2000). Filtering historical simulation. Backtest analysis. Manuscript, March, 2000
- [6] Barone-Adesi, G., Bourgoin, F., & Giannopoulos, K. (1998). Don't look back. *Risk*, 11, August, pp. 100-104
- [7] Barone-Adesi, G., Giannopoulos, K., & Vosper, L. (1999). VaR without correlations for Non-Linear Portfolios. *Journal of Futures Markets*. 19, August, 583-602



Active Asset Allocation provides institutional investors with cutting-edge risk-controlled asset allocation solutions.

Our clients are portfolio managers, pension funds and foundations that seek to make the best use of assets they select, while keeping their downside risk under control. We offer them bespoke asset allocation models based on proprietary technologies, and support them throughout the life of their portfolio.

Find more about our solutions at active-asset-allocation.com

Disclaimer

© 2015, Active Asset Allocation. All rights reserved.

This document has been prepared by Active Asset Allocation, as thoroughly as possible and with extreme caution. Information, descriptive analyses or statistics contained herein have been obtained from sources believed to be reliable, but not guaranteed and they do not engage the responsibility of Active Asset Allocation. This material has been distributed for informational purposes and for the recipients' use only. It should not be considered as investment advice or a recommendation of any particular security, strategy or investment product and it does not release readers from the need to form their own judgment. Active Asset Allocation shall not be held responsible either for any direct or indirect damage incurred as a result of the use of information provided in this document. No part of this material may be reproduced in any form, without express written permission from Active Asset Allocation. Violation of this provision may lead to legal action.

February 2015