Historical-simulation method

Dario Cintioli and Marco Marchioro explain how StatPro developed the historical 'simulation method' as a base for its risk management product, SRM, and outline the reasons why this model is rapidly gaining ground in the asset management industry

Risk management, in the asset management industry is facing new challenges. The growth of exchange-traded funds (ETFs) and the growing success of hedge funds are increasing the pressure on traditional asset managers. The success of ETFs limits the ability of asset managers to sell passive funds, while the growth of hedge funds forces asset managers to move from benchmark-driven products to absolute return strategies.

These trends are forcing asset managers to increase the component of 'active' risk in their funds and, subsequently, putting pressure on them to produce even greater returns. Today, many asset managers are achieving these objectives, by increasing the 'leverage' of their active bets through the use of derivatives and structured over-the-counter (OTC) products. The explosion of the credit-derivatives market and the birth of exotic instruments, such as CDOs are simply the tip of the iceberg of this phenomenon.

The consequent increase in the portfolio complexity and the corresponding pressure coming from regulators, who want to ensure that these instruments are correctly evaluated and managed, (see, for example, the UCITS III regulation in Europe), which are seriously challenging the existing risk management operations and the econometric models traditionally used by asset managers. The most popular risk management solutions for asset managers are based on the so-called 'factor' models. These models assume that the price of all financial assets can be explained by a limited set of statistical or macro-economic factors, by means of linear relationships. However, the introduction in the portfolios of complex and illiquid products, often based on asymmetric payouts, is reducing the ability of these models to capture the correct risk profiles. The factor models are not suited because:

- The linear relationships cannot explain the asymmetric behaviour of exotic and option-based products.
- Very often the complex products are not covered at all and ignored, failing to identify the risk where needed.

We believe that asset managers can respond to these emerging needs by adding the 'option-based' risk-management models, most widely used in the banking industry to the traditional factor approaches.

Among the existing different approaches to risk management, we believe that the best one is the historical-simulation model because it can consistently manage the portfolio complexity and can be adapted to the needs of asset managers.

In this article, we explain how StatPro has implemented the historical-simulation model in its StatPro Risk Management (SRM) product and we give some evidence of its flexibility and accuracy.

Historical simulation method for a simple stock

The goal of the historical-simulation method is to find a large number of possible scenarios for tomorrow's asset price – given today's price. These scenarios will build up the distribution, so that any risk measure can be computed. Since, for the most part, these scenarios are computed artificially, rather than observed on the market, we will name the set of such scenarios the 'simulated distribution'.

We will begin with the computation of the simulated distribution for a simple stock and then gradually extend it to more complex assets. For simplicity, it will be assumed that the horizon period for the evaluation of risk is one day.

The computation of the scenario distribution for a simple stock is straightforward and is built on the 'stationary assumption', where it is assumed that the return of the stock from today to tomorrow is distributed like the historical frequency of the daily returns for the same stock during a given past period, for example, two years.

However, it is possible to consider a shorter reference period, such as one year, or a longer period, such as five years. Experience has shown us that using two years of daily data is optimal: the period is long enough to have a statistically relevant number of samples – about 500 – and short enough to react adequately to market changes.

Simulated distribution for a zero-coupon bond

Let us consider the computation of the risk distribution for a zero-coupon bond, such as a bond that does not pay any intermediate coupon, but only the notional amount at maturity. Such a bond will be traded at a discount and its price can be computed bootstrapping an interest-rate term structure that uses, for example, an array of deposit and swap rates (see, for example, J. Hull [1] for an introduction, or QuantLib [2] for an implementation). StatPro, in order to correctly evaluate the price and risk of any bond, considers the credit risk associated with its issuer. However, to simplify the explanation that follows, we will neglect the influence of credit spread on the bond price. Therefore, let us assume that the price of the given bond is a function of the deposit and swap rates used to bootstrap the interest-rate curve:

$$B = f_B\,(t;\, d_1, \ldots, d_n, s_1, \ldots, s_m)\,, \qquad (1)$$

where B is the bond price at t=today, $d_1, \ldots, d_n$ are the n deposit rates, and $s_1, \ldots, s_m$ are the m swap rates observed on the evaluation date; finally, f_B is the bootstrapping function that allows the computation of the discount factor at bond maturity. In this case, we say that the deposit and swap rates are the 'risk factors' for the zero-coupon bond.

So far, we have only computed the bond price. To obtain a distribution of prices for the bond at t=tomorrow, it is possible to proceed in one of the following two ways:

1. Consider the historical values of the bond price B, compute the daily returns and assume that the simulated distribution is given by these returns.
2. Assume that the distribution of $d_1, \ldots, d_n$, and $s_1, \ldots, s_m$ at tomorrow is given by their historical returns in a way similar to the simulated distribution for a simple stock. Using f_B and the simulated values for the deposit and swap rates, it is possible to obtain a distribution for the bond price B. This is the method used in SRM.

The first method is much simpler to implement, since it requires the historical values of the past bond prices, only. However, it does not account for the changing of the bond duration during its life. The second method is more computationally intensive (every day, one has to bootstrap a large number of interest-rate curves), but correctly keeps track of the bond duration. The observation of actual price variations shows that this makes a big difference in assessing the actual risk of a bond.

Figure 1 shows the daily variations for a zero-coupon bond issued by the European Investment Bank expiring in 2011, together with the 99th, 95th, 5th, and 1st percentile of the simulated distribution for the two years from August 2003 to August 2005. The 99th and 95th percentile of the distribution directly correspond to the measure of value-at-risk (VaR) at the respective percentiles. The 1st and the 5th percentile correspond to the so-called 'potential upside' of the distribution. From the figure, we notice that the variance or covariance methods, the distribution is not symmetric and the downside is greater than the upside. It can also be noticed that the decreasing duration of the bond, during the observed two-year period, results in a corresponding decrease of the VaR and potential upside of the bond, as confirmed by the behaviour of the recorded price series. This can be a rather large effect, where in the above figure, the 95th percentile roughly halves in two years, and consequently would not be captured if the distribution of the actual historical returns were used. Using this method would result in an overestimated VaR figure for this bond. However, it is not a difficult task to envision instruments, where the risk would be underestimated, for example, because today's market makes them more sensitive to variations of their risk factors than they were in the past. This shows the superiority of using the historical-simulation method, rather than other historical analyses.

1. Back-testing chart for a zero-coupon bond

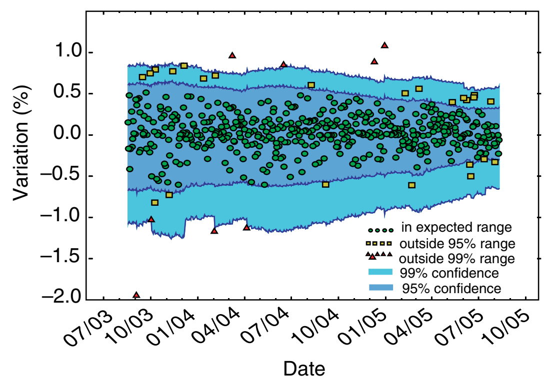

General financial assets and their risk factors

As mentioned earlier, the portfolios of asset managers increasingly hold positions on all sorts of asset types. These assets could be simple stocks, bonds of all types or more complex instruments such as Asian options, mortgage-backed securities or any other type of derivatives. In general, we can assume that any asset is a derivative on a certain number of underlying financial variables or risk factors. For example, for the futures on the S&P 500 index expiring in three months, the risk factors to be considered are the value of the index and all the components of the interest-rate curve for the US dollar.

We assume that for each asset in the portfolio it is possible to write its price as a deterministic function of n risk factors:

$$P_j = f\,(t;\, r_1, r_2, \ldots, r_n)\,, \qquad (2)$$

where P is the price of the asset, t = today, f is a deterministic function depending on n variables, and $r_1, r_2, \ldots, r_n$ are the risk factors for that particular asset. In the simplest case, an asset depends on a single risk factor, which is the asset itself. This is the case for equities, mutual funds, and indexes. Another example of equation (2) is the Black-Scholes-Merton pricing

2. Simulated distribution of daily returns for a double-performance bond

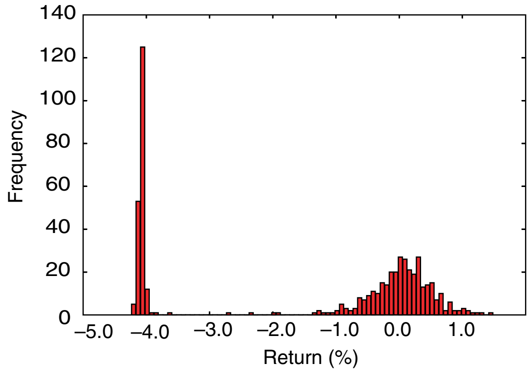

formula for equity options, where the risk factors are the underlying stock prices together with their volatilities and the interest rates.

It should be noticed that, in some cases, the choice of a pricing formula for an asset might be arbitrary. For example, in the case of a callable bond, you may decide to price the bond using different models, such as the one proposed by Hull and White [3] or the one proposed by Brace, Gatarek and Musiela [4]. In these cases, it is important to choose a model that is sophisticated enough to capture the correct dependency from the risk factors.

At StatPro, special care is given to the development of these pricing functions, by using many different numerical schemes such as finite differences (see, for example, Tavella and Randall [5]) or state-of-the-art models, such as the prepayment model for mortgage-backed securities, which were developed by Kalotay, Yang and Fabozzi [6]).

To compute the simulated distribution of returns for the generic asset with the pricing formula given by equation (2), it is necessary to obtain the distribution for each of its underlying risk factors $r_1, r_2, \dots, r_n$. Suppose that for $i = 1, \dots, n$ the simulated values for the risk factor, as computed in the case of a single stock, is $r_i^{(k)}$, then the simulated price for the asset j at the scenario k is given by:

$$P_j(k) = f(t; r_1^{(k)}, r_2^{(k)}, \dots, r_n^{(k)}) , \qquad (3)$$

where t = tomorrow, and $k = 1, \dots, s$, being s the total number of scenarios available. To obtain a distribution of returns for asset j to each historically-simulated scenario, we assign them a weight of $1/s$.

Note that, in this way, we capture all the non-linear dependencies of risk from the underlying financial variables and therefore gamma and vega risks are automatically included in the model; correlations among risk factors are also accounted for.

For example, consider a double-performance bond that expires in 2007. This bond pays coupons related to a basket of diversified stocks. The bond pays coupons every February with the following variable rate: 7.5%, if all the stocks in the underlying basket do not fall under 70% of their value at the end of February 2004 in the year prior to the payment of the coupon; 0%, if any do. Figure 2 shows the computed distribution of simulated returns on August 19, 2005. It can be seen that the price does not follow any well-known, bell-shaped distribution. The shape of the distribution can be justified as follows: at present, one of the assets of the basket is just above the 70% threshold. In the scenarios, in which it goes below this threshold, the price goes down 4%, whereas in the other scenarios, it fluctuates around the current price. The distribution shown in figure 2, clearly presents itself as a challenge for any analytical method, however, it is perfectly handled by the historical-simulation framework.

3. Risk factors used by StatPro Risk Management

equity prices	**fund prices**
equity-index prices	**interbank-deposit rates**
commodity prices	**interest rate swaps**
inflation swaps	**foreign-exchange rates**
corporate asset-swap	**indexes implied volatilities**
emerging market credit spreads	

Different types of risk factors

See figure 3 for the risk factors that were considered at the time of writing. Given the variety of risk factors, it is clear that the accurate estimate of all the correlations among these financial variables cannot be made without introducing any simplification. The historical-simulation framework automatically takes into account all these dependencies, without simplifying the complex relations among them with a variance or covariance matrix.

Back testing for different asset classes

The final word on any risk management framework should always be given to a back testing procedure. Therefore, we gathered two years of data from August 2003 to August 2005, across all the asset classes available in StatPro Risk Management. Figure 4, shows the percentage of assets outside the tail of the distribution – both at 99% and 95%. The last column shows the total number of samples considered to compute the percentages.

We notice that the number of events, in which the return exceeds the tail – the exceptions – is always less or about 1% for a percentile of 99% and less than 5% for a percentile of 95%.

In the asset classes, which are directly related to the stock markets, we notice that the percentage of exceptions is lower than the expected value. This is because the average volatility of the stock markets has decreased considerably during the past few years. Furthermore, it can be noticed that the percentage of exceptions for a percentile of 99% is always relatively higher than those for a percentile of 95%. This behaviour suggests that the presence of so-called ‘fat tails’ in the distribution of returns.

The presence of fat tails is seen even more clearly in the case of warrants, where the percentage of exceptions is close to 1% for a percentile of 99%, but only about 3% for a percentile of 95%. This effect can be attributed to the leverage that warrants have compared with the underlying stocks.

We notice that in the case of indexes, funds and futures, the number of exceptions is about one half of what is expected. In reality, since futures are written on indexes and funds are most of the time managed against a benchmark, the exceptions are not independent and are explained by the indexes. In figure 4, these asset classes can be regarded as a single one. The observed overvaluation of risk comes from the trend of decreasing volatility of the past few years, such as stocks, and can also be partly explained by the rebalancing of the equity indexes. Since equity volatility is normally half-smiled and increases as prices come down and decreases as prices rise, in a general market environment of declining volatility, the index rebalancing process can accelerate the reduction of volatility by replacing losers, (with declining prices and increasing volatility) with winners (by increasing or stable prices and declining or stable volatility). This is one of the reasons why the back testing of indexes shows a lesser number of exceptions than stocks.

Conclusion

In this article we have explained how StatPro has implemented the historical-simulation model. We have shown a wide set of back-testing results, while considering the exceptions for the past two years, which is equivalent to four years of historical data, and a large set of instruments totalling millions of observations. This set shows a clear evidence of the reliability of the model on many asset classes – on both linear and non-linear instruments – as shown in figure 4. It is also remarkable to notice that in some assets, such as stocks, the strong declining volatility of the past four years had only a minor impact on the back-testing results, especially in the 99% tail: the violation of the stationary assumption has only been marginally transferred into the effective number of exceptions.

At StatPro, we have improved the model and extended it to the needs of the asset management industry, because we believe that the recent trends in the buy-side will force the industry to gradually replace (or complement) traditional linear econometric models with option-based models. In our opinion, the historical-simulation model is currently the most appropriate model for asset managers for the following reasons:

Wide asset coverage

It is very easy to cover new types of assets as you only need the pricing function for the asset and the history of the underlying risk factors. In this way, it is possible to cover any type of asset: liquid or illiquid, simple or exotic.

4. Back testing for each asset class

Asset type	99%	95%	Samples
Stocks	0.76%	3.41%	4,564,381
Funds	0.53%	2.54%	2,908,593
Warrants	1.01%	2.98%	204,193
Indexes	0.50%	2.22%	346,631
Deposit/swaps	1.08%	3.78%	140,973
Currency exchange	1.00%	3.85%	23,786
Bonds	0.84%	3.06%	5,220,004
Futures	0.51%	2.20%	117,087
Plain options	0.85%	2.89%	3,743,164
Exotic options	0.81%	2.80%	95,021

Integrated framework

The simplicity of the historical-simulation framework can be applied to any asset class that you need to manage. Indeed, it is not necessary to write a new model, and therefore add new assumptions, every time a new asset class requires a new type of risk factor. However, this is not true for the Monte Carlo method, for example, or for some other approaches. In the historical-simulation model every risk factor is automatically and consistently integrated with each other.

No calibration parameters

Another feature of the historical-simulation framework is that it does not need the calibration of any parameter to obtain a distribution that matches the market. In particular, no assumptions are made on the shape of the distribution. It captures the distributions of risk factors as it is. The relevance of this property is particularly evident in non-linear products, such as the double-performance bond considered earlier.

Speed and flexibility

One of the traditional shortcomings of the historical-simulation model is the speed of computation. StatPro has solved this problem, thanks to a special proprietary architecture, which allows for speed of execution that compares to analytical models. While it is obvious that the historical simulation is faster than the Monte Carlo model, we found that for large portfolios, the actual speed is even greater than that of many analytical models.

Risk attribution and decomposition

StatPro has developed a method that reliably computes risk decomposition and attribution in the framework of historical simulation. The description of this method is beyond the scope of this article and will be given in a future paper. With this development, StatPro has removed the final obstacle that was present between asset managers and the historical-simulation framework. ■

References

[1] John C. Hull Options, Futures, and other Derivatives Fifth edition, Prentice Hall Finance Series

[2] QuantLib financial library, http://www.quantlib.org

[3] J. Hull and A. White Pricing Interest Rate Derivative Securities Review of Financial Studies, 3, 4 (1990), 573–592

[4] A. Brace, D. Gatarek, and M. Musiela The Market Model of Interest Rate Dynamics Mathematical Finance, 7, 2 (1997), pp. 127–155

[5] D. Tavella and C. Randall Pricing Financial Instruments: The Finite Difference Method Wiley series in financial engineering

[6] A. Kalotay, D.Yang, and F.J. Fabozzi An Option-Theoretic Prepayment Model for Mortgages and Mortgages-Backed Securities Int. Journal of Theoretical and Applied Finance, Vol. 7, No. 8 (December 2004)

Contacts

Visit: www.statpro.com

Dario Cintioli, Managing Director, StatPro Italia

Marco Marchioro, Head of Research & Development, StatPro Italia

T: +39 02 4331 7503

Via G.B. Vico 4, 20123, Milano

E: dario.cintioli@statpro.com

marco.marchioro@statpro.com

overmatter for ‘statpro’