RiskMetrics Group™ Research Monthly

# Forecasting for solvency risk

Christopher C. Finger
chris.finger@riskmetrics.com

Gilles Zumbach
gilles.zumbach@riskmetrics.com

December 2005

In 1714, the English Parliament announced a great prize for the inventor of an accurate method to determine longitude. The expectation was that the winning method would be an elegant set of astronomical observations. In the end, the prize went to a John Harrison, who was able to build a chronometer that, even after weeks at sea, could keep time at the ship's origin accurately enough to determine longitude.

Forecasting financial time series at a one-year horizon begins with a dilemma. We may choose methods typically used to forecast at ten- to twenty-year horizons, examining extremely long histories of infrequent data, and making assertions at the comparatively short annual horizon. Alternately, we may look to daily data, and seek to scale our forecasts to the seemingly long one-year horizon. Daily data seeming somewhat more terrestrial (and our new location being associated somewhat with timepieces), we take inspiration from Mr. Harrison, and choose the second route.

## The business problem

We have tackled the issue of one-year risk analysis before, with a focus on the treasury operations of non-financial corporations. There are two key characteristics of this business problem. First, the risk that concerns us is volatility in financial statements, and so the exposures we track are cashflows, accruals, and mark-to-market changes, and the timing of these within a one-year planning horizon. Second, because we are examining corporate hedging programs, the risk factors we consider (commodity prices, exchange rates, short-term interest rates) for the most part have traded derivatives; thus, we can use option prices to estimate the volatility of the risk factors at the one-year risk horizon.

We take up the issue again, though this time in the context of pension and insurance firms. Though the horizon is the same as before, the risk is now solvency—that in the future, the value of the firm's assets may be insufficient to meet its liabilities. While the risk of explicit insolvency is likely remote, the event that insolvency becomes "closer" is in fact relevant. Thus, we examine notions of distances to insolvency, such as the surplus of the market value of the assets over the discounted value of the liabilities. We concern ourselves with the risk that in the next year, the distance to insolvency drops below some barrier set, likely by regulators, as warranting intervention or restrictions on the firm's activities. Importantly, we focus not on the risk of the firm meeting obligations within the current year, but rather on the risk of it covering the present value of liabilities beyond one year.

The financial aspect of the problem is now not a hedging program, but an investment portfolio and a future stream of liabilities. Consequently, the financial variables that concern us are those factors that drive the asset portfolio performance, those that influence the magnitude or timing of the liabilities, and those that are used to discount the future liabilities to the present. As a contrast to the corporate problem, there are not option markets for these financial variables from which we can infer volatility estimates; here, we must seek a direct approach.

Importantly, our treatment excludes risks that derive from business uncertainties. For corporates, this means that we do not treat uncertainties in sales volumes; for pensions and insurance, this exclusion applies to underwriting and actuarial risks. Thus, for the purpose of our analysis, we will assume that actuarial and underwriting factors are static. We will not assess the risk, for instance, that casualty claims are unexpectedly high, nor that members of a pension fund have surprisingly high rates of early retirement.

These assumptions primarily affect our treatment of liabilities. For casualty insurance, the magnitude and timing of claims are unrelated to financial variables, and the only dependence of the liabilities on the factors we treat is through the interest rate used for discounting. Here, we simply consider the expected magnitude and timing of claims, and model the financial risk that derives from the discount curve.

Pension liabilities are somewhat trickier, in that pension benefits are often linked to inflation, which in turn is correlated with financial variables. For practical purposes, we may still represent the liabilities as a known set of future benefits, given a specific set of actuarial and inflation projections, and analyze the risk due to discounting. To the extent this is insufficient, a somewhat more sophisticated approach is to assess the sensitivity of the present value of the liabilities to inflation changes, and model this through a link to real interest rates.

Liabilities in life insurance provide the greatest challenge, in that the payouts themselves may depend on the same financial variables, and in a non-trivial way. Consider, for instance, a life insurance product with a guaranteed minimum return, but with upside potential linked to the performance of an asset portfolio. For these products, in order to truly capture the financial risk, it is not sufficient to set the liabilities to their mean levels and discount; here, we need either to model the liabilities explicitly as financial instruments (holding the actuarial parameters fixed), or to create a replicating portfolio of (possibly non-linear) financial instruments.

The one-year risk horizon turns out to be an interesting one, in that it has both short- and long-term characteristics; perhaps for this reason, the amount of literature discussing this horizon is relatively meager. From an actuarial point of view, the one-year horizon is decidedly short-term; it is defensible to say that projections of mortality, property claims, or even inflation will not change over this horizon. From a financial point of view, the picture is more complex.

One year is long relative to the trading horizons at which we typically forecast volatility and risk, and certainly long relative to the amount of data we typically have at our disposal to study. On the other hand, one year is short enough such that recent information still matters: the fact that a market is volatile today may or may not affect our projections of returns over twenty years, but this recent information certainly influences our projections at one year.

Figure 1: Properties of financial time series. S&P500 returns (left) and lagged correlation of return magnitudes (right). Time series represented are equity indices (green), European interest rates (red), US interest rates (light blue), and foreign exchange (blue). The vertical bar indicates a lag of one month.

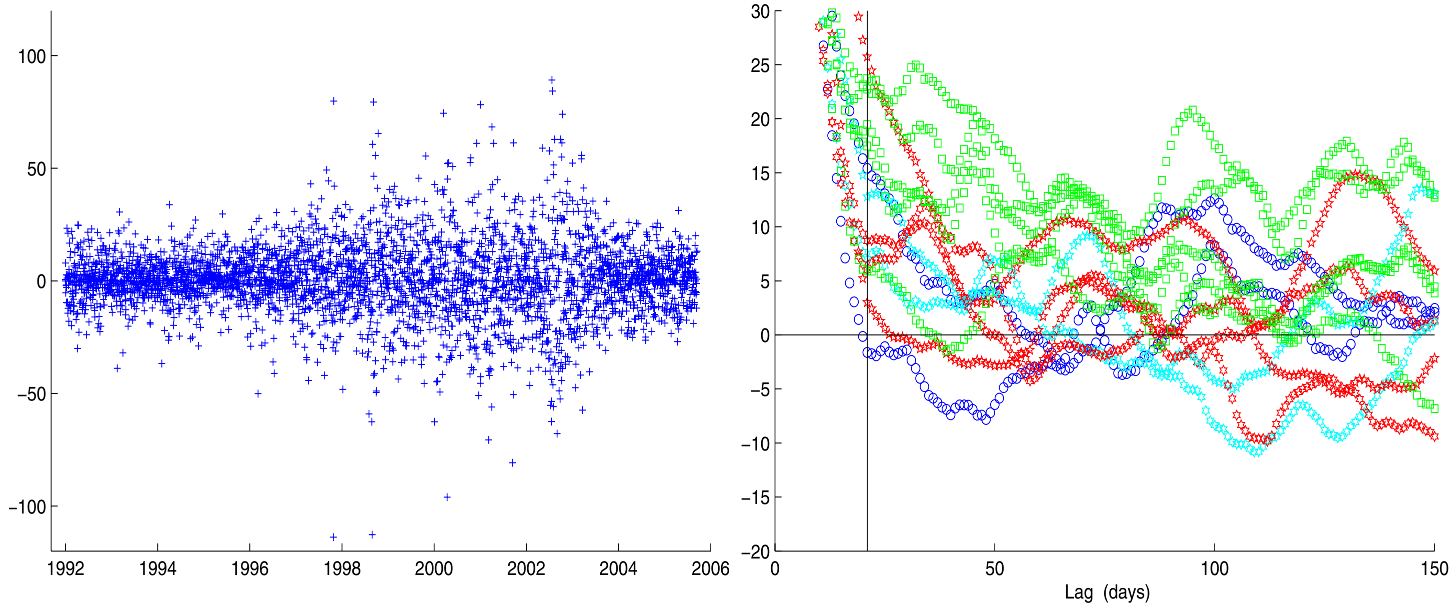


## The data

Before discussing our modeling approach, we review a few key features of the data. First, financial returns are heteroskedastic, meaning their volatility changes through time. This is clear for the S&P 500 in the left graph of Figure 1. Moreover, the returns displayed there exhibit volatility clustering: large (and small) returns tend to appear together. That this clustering exists, and large returns do not in general appear out of nowhere, means there is hope of forecasting the volatility of future returns given today's information.

Volatility forecasting is well established in risk management practice. However, applications are typically for short (one- to ten-day) horizons. A relevant question to ask is whether this type of analysis matters for longer horizons. That is, is heteroskedasticity something we need concern ourselves with when forecasting the volatility of annual returns?

If it were the case that the data exhibited very short bursts of volatility—one or two weeks—followed by longer periods of "normal" market conditions, then we could perhaps argue that at a one-year horizon, these small bursts of volatility should average out, and our forecasts should not be influenced by whether we are currently in one of these volatile periods. This is not the case, however. Returning to our plot of S&P 500 returns, we see that volatile (and quiet) periods sometimes last for a month, but more often persist for a year or more.

A more precise description of the persistence of volatility is through the lagged correlation of the sizes of returns. For example, we might examine at the correlation between the magnitude of a monthly return at one point in time with the magnitude of the monthly return from two months before; here, we would describe, for instance, the effect of the size of the return in January on the size of the return the following March. In this example, one month is our forecasting horizon, and two months is our lag.

In the right plot in Figure 1, we present the lagged correlation of the sizes of monthly returns for a selection of time series. Note that there is a reasonable level of correlation (fifteen to twenty percent) even at lags as long as six months. We refer to this feature of the data as a long memory.[1] The implication of the long memory is that volatility in January still has some influence on the size of the return in July. As we consider the return for the entire year, it is clear that knowing where we stand in January is a crucial element of forecasting.

For interest rates, one typically expects to observe some indication of reversion to a long-run mean level: a specific interest rate should not necessarily grow without bound, but rather tend to revert, over sufficiently long horizons, to some equilibrium rate, such as six percent. This type of mean reversion[2] should show itself in the lagged correlation of the returns themselves (as opposed to the lagged correlation in magnitudes we examined earlier). If mean reversion is displayed, then at some lag, we should see negative correlation, since positive returns (which take the rate above its long-run mean level) should eventually be compensated by negative returns. This is not actually something we observe, at least within the one-year horizon.[3]

## The tried and true

From a modeling perspective, our first goal is to forecast volatility over the relevant risk horizon. For short horizons, particularly one day, our current model is well understood: mean returns are zero, the volatility of returns is estimated using an exponentially weighted moving average (EWMA), and the distribution of returns is normal. The EWMA estimator, with the parameter (decay factor) we typically use for one-day forecasts, can be seen as a balance between the need for enough historical data to produce stable estimates and the desire to base the estimate on only the very most recent data.

Under the current model, to forecast risk at longer horizons, we make the additional assumptions that volatility is constant over our forecasting horizon and that returns are independent from one day to the next. These two assumptions then imply that returns over multiple days are still normally distributed, and that their volatility simply scales with the square root of the risk horizon. An important, but somewhat subtle, point about our model is that we use the EWMA framework purely as an estimator. From one day to the next, our estimate of volatility will change; but when we forecast risk on a given day, we assume that returns follow a process wherein the volatility over the entire horizon is constant.

A slight variation on the constant volatility process is the I-GARCH process. Under I-GARCH, the (squared) volatility of returns tomorrow is equal to a weighted average of the effective (squared) volatility

[1] Technically, long memory means that the influence of past returns decays as a power law, not as an exponential, in the lag.

[2] It is important to note here that we are speaking of mean reversion of the rates themselves, not of their volatility.

[3] In fact, the strongest lagged correlations we see are for short maturity interest rates. These correlations are positive, indicating that there are persistent trends in these rates—the opposite of mean reversion. This is a sensible observation, as these rates are the most sensitive to central bank policy. It is also less concerning for our purposes here, since the liabilities we treat are typically much longer than these rates; moreover, sufficient options markets exist for these rates to produce market-implied forecasts of volatility.

today and the square of today's return; and given the volatility, subsequent days' returns are independent. By construction, the size of one day's return impacts the size of future returns, but the directions of returns from one day to the next are uncorrelated. As a process, I-GARCH is interesting in that it shares the square root scaling law with the constant volatility process. For example, if we assume that returns follow an I-GARCH process, then the volatility of the cumulative return over the next 250 days is equal to $\sqrt{250}$ times the current daily volatility.

Because we are accustomed to the constant volatility process, this does not seem surprising, but it is in fact not trivial. The volatility of the 250-day return is (loosely) the average of the volatilities of the individual daily returns in the period. But a large return on one day leads, through the I-GARCH process, to a higher volatility the next day, and a small return to a smaller volatility; it is not obvious that averaging over such a process yields such a simple scaling rule.

But does I-GARCH work for longer horizons? Unfortunately, no. We have seen before that forecasts at a one-month horizon require a different decay factor from forecasts at a one-day horizon. This in itself is an indication that the model does not fully describe the dynamics of the time series over even a one-month horizon. This leaves us with the prospect of searching for a new decay factor for each horizon, yet once we move beyond one month, our data will not support this. We need a model which better describes the features of the data, and from which we can extrapolate forecasts at the one-year horizon. The processes we consider as alternatives still lend themselves to forecasting. At their core, the processes are corrections to the random walk, and so their scaling rules (that is, how to go from a one day to a longer horizon volatility forecast) are still primarily driven by the square root of the horizon; their actual scaling rules, though, are more complex.

## A new approach

Forecasting at a one-year horizon is a daunting challenge, not least because we typically can count on only twenty years of data with which to calibrate our model. It is a mistake, though, to think of this data as only twenty points when we in fact have daily observations. (We should not complain about lack of data, and then proceed to throw away 249 of every 250 observations.) Similarly, it is a mistake to endeavor to develop a model exclusively for the one-year horizon: on the one hand, such a model would have no connection to forecasts at shorter horizons; on the other, modeling only at the long horizon would force us to make empirical assertions at the point where our data is most sparse and our statistical power weakest.

Our strategy, then, is to make the most of the data we do have, and to build a model from which we stand on firm theoretical ground when we make assertions for longer horizons than we have data to support. We would like the features of our process to tell us how to scale to longer horizons. Contrast this with the approach where we suggest a new decay factor for longer horizons, but without a model that in any rigorous way links the decay factor to the horizon.

We will validate the model empirically for as long a horizon as the data will permit, and then rely on the theoretical properties of our model process to take us the rest of the way. In the end, our forecasts for one year will be supported by solid empirical evidence at slightly shorter horizons, plus a theoretical model to bridge to one year, in addition to the (perhaps) shaky

Figure 2: Effective weights for volatility forecasts under EWMA and LM-ARCH models

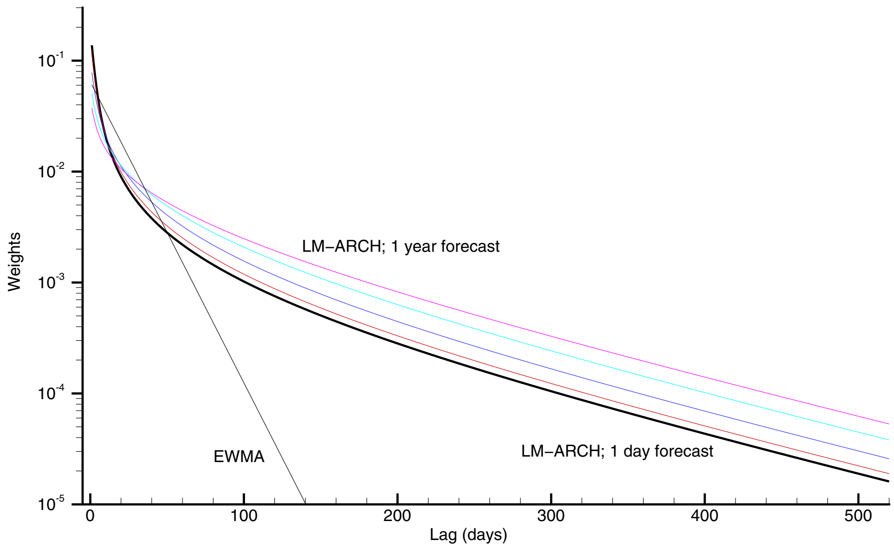


direct empirical evidence from the annual observations. Contrast this with the approach of modeling with annual data only, wherein the shaky empirical evidence is all we have.

To build our model, we assume again that returns are uncorrelated in time. As to the volatility, we have two requirements of the model. First, it should capture the long memory behavior of the data that I-GARCH misses. Second, in order to limit the number of parameters we need to estimate, the process should have no time series-specific parameters. The second requirement admits the I-GARCH model, for instance, but precludes most standard GARCH models. Standard GARCH models effectively have a long-term volatility to which the process reverts; this volatility clearly cannot be the same for all time series. On the other hand, the only parameter in the I-GARCH model is the decay factor, which can be seen as universal. A process satisfying these two requirements is the Long Memory ARCH, or LM-ARCH.

The LM-ARCH model can be seen as a mixture of I-GARCH processes with different decay factors, or equivalently, with different characteristic time scales. The processes are parameterized effectively with the shortest and longest time scales (equivalently, greatest and smallest decay factors) employed. An attractive consequence of these processes is that our volatility estimates are no more than a linear combination of past (squared) returns.[4]

In Figure 2, we display the effective weights as a function of lag for our standard I-GARCH model, and for the LM-ARCH model with a variety of forecast horizons. For I-GARCH, since the volatility forecast scales simply with the square root of the horizon, the shape of the effective weights is the same for all horizons. For LM-ARCH, we observe that the shape of the weights is convex, an indication of its long memory. In addition, we observe that the shape of the weights for LM-ARCH changes with the forecast horizon; as we move from one day to longer horizons, the weights for short lags decrease and those for longer lags increase. Importantly, this

[4] This is not true of all processes: the GARCH family implies estimators that include the long term mean volatility.

Figure 3: Properties of time series of residuals. S&P500 residuals (left) and lagged correlation of residual magnitudes (right). Time series represented are equity indices (green), European interest rates (red), US interest rates (light blue), and foreign exchange (blue). The vertical bar indicates a lag of one month.

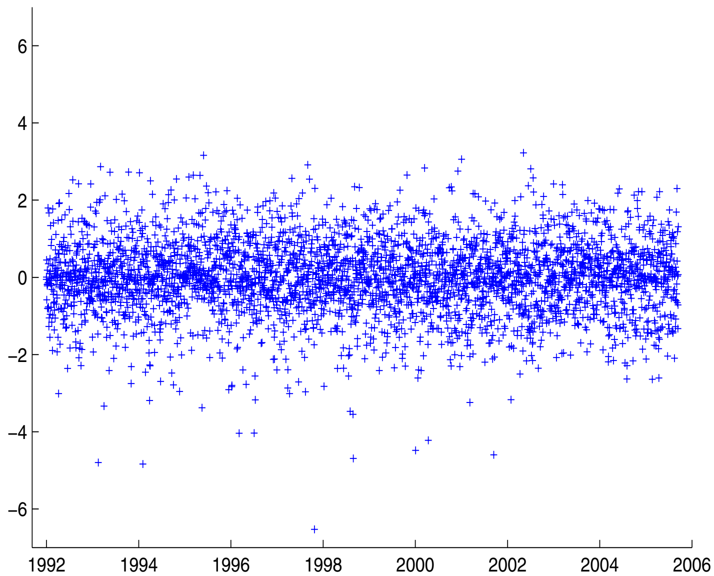


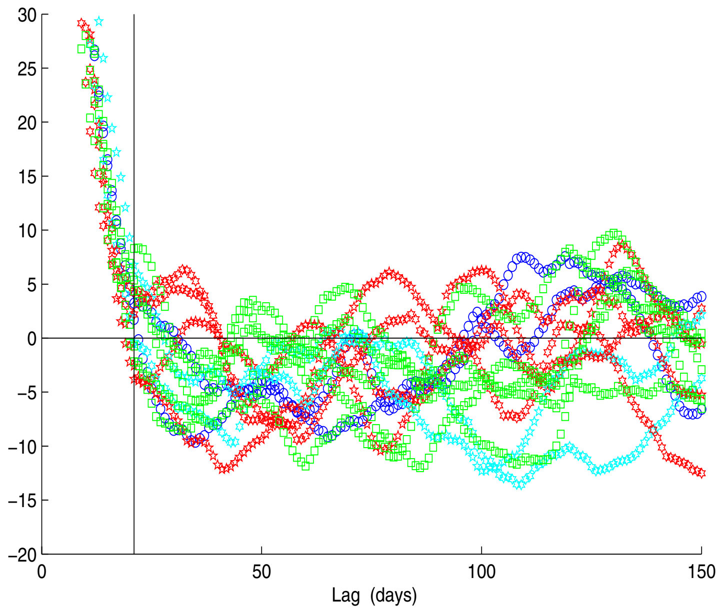


is not a result of changing model parameters from one horizon to the next; it is a result of the model itself, with a fixed set of parameters.

## Evaluating

To evaluate the model's performance, we examine the residuals—that is, what is left over once we have let our model do its work. The residual for a particular time period is defined as the actual return for that period divided by the volatility forecast we would have made at the beginning of that period. For example, the residual for January is the actual January return divided by the monthly volatility we forecasted on January 1. Note here that we still forecast using the daily data, but that we examine the model performance (through the residuals) at the horizon of interest.

If our model performs well, then the volatility forecasts at various horizons should account for the two key aspects of the data mentioned previously (volatility clustering and long memory), and the time series of residuals should not display either of these properties.

In the left plot of Figure 3, we plot the daily residuals for the S&P 500. In sharp contrast with Figure 1, we see that the volatility does not appear to vary through time. This is an indication that for a daily forecast horizon, the model has well described the volatility clustering.

In the right plot of Figure 3, we examine the lagged correlation of the sizes of monthly residuals. Recall that in Figure 1, we observed correlations of fifteen percent, even at lags as long as six months. Here, we see the correlation drop to below five percent at only a one-month lag; this holds for all of the time series represented. The model thus describes well the long memory in monthly returns, and is robust across time series. We emphasize that all of these results derive from a fixed set of model parameters; there is no need to tweak the model for the longer horizon, nor to provide custom parameters for specific time series.

## What next?

As we lengthen the forecast horizon, the amount of data (that is, the number of non-overlapping residuals) available to validate our forecasts becomes sparse. So it is at a three- to six-month horizon where we abandon the approach of validating based purely on the data, and start to assert that the model should carry us the rest of the way. The point here is not that the model performs badly at a nine- or twelve-month horizon, but that we cannot make strong conclusions about the model based only on empirical evidence. One line of ongoing work is to continue to refine the volatility forecast, in order to provide better fits at one month, and with the hope of strengthening our assertions at one year.

A second line of ongoing research is to describe the distribution of the residuals, and in particular to determine if the fat tails (and asymmetry) we observe in the daily residuals[5] are still present at longer horizons. We do this by fitting distributions to the longer horizon residuals directly. Under this approach, we can once again hope to stand on solid empirical ground up to one month or more, and to assert that the residual distribution is reasonably stable from this point to one year. Using Monte Carlo, we may validate this assertion by examining the stability of the residual distribution, for instance, from one month to one year, under our assumed process.

In the end, we have addressed the one year horizon with a process that fits our data at short and medium horizons. Of immediate interest here is that we can scale this process to one year. Of perhaps even broader interest is that the process improves our forecasts at all horizons, making it possible to examine the risk from strategic and tactical points of view under a single framework. There is work to be done still, but we remind ourselves that it took John Harrison thirteen years to construct his chronometer. By that standard, we are well ahead of schedule.

## Further reading


- Stern, E. (2004). Risk management for non-financial corporations. *RiskMetrics Journal*, **5**(1): 41–69.
- Mina, J. and Xiao J. (2001). *Return to RiskMetrics: The Evolution of a Standard*, RiskMetrics Group.
- Zumbach, G. (2004). Volatility processes and volatility forecast with long memory. *Quantitative Finance*, **4**(1): 70–86.
- Zumbach, G. (2005). Long term risk evaluation for pension funds, Presentation, RiskMetrics Group European Client Conference.


[5] See Figure 3.