NORTHFIELD
modeling
financial markets
worldwide

Risk Model Testing and Regulatory Reporting

Dan diBartolomeo

Northfield Webinar

August 2015

Motivation

- Over the past couple years we have received an increasing number of requests from clients for some form of test data validating our risk models
  - Most of these requests have been related to either the UCITS regulations in Europe, or US Federal Reserve Supervisory Release 11-7, which covers the issue of "model risk".
- Increasingly, we see asset managers and some asset owners having two risk systems
  - One focused on measuring risk in order to make better investment decisions, and a second system focused on regulatory risk reporting, which typically has a much shorter time horizon. Often the two approaches to risk assessment produce conflicting results

Presentation Outline

- The semantics of "model risk"
  - Differences in understanding across investors, managers and regulators are driven by the context of financial objectives
- Three different ways of testing risk models "out of sample"
  - Cross-sectional discrimination
  - Normality analysis
  - Exceedance tests
- How to avoid the need for two risks systems by generating good short term VaR estimates (the usual regulatory requirement) from regular Northfield systems and risk models

Key Concept

- Most regulatory risk requirements have been adapted (or in some cases adopted directly) from regulation of commercial banks
  - Banks are highly levered entities (often 40 to 1 or higher)
  - Banks have most of the liabilities at call (depositors can withdraw, commercial paper matures regularly)
  - The regulatory focus is balance sheet ***solvency at each moment in time***
- Bank objective function
  - Make as much money as you can every day, subject to an upper bound on the daily probability of financial default
  - If asset returns are normally distributed this is maximizing the Sharpe ratio

Model Risk

- Both the UCITS rules and SR 11-7 discuss "model risk"
  - Some analytical models being used in any part of the investment process can be wrong. Models should be validated somehow
  - Since the regulatory focus is on "everyday" balance sheet solvency, the practical emphasis *is computational accuracy and data integrity in the pricing of financial assets at a time point*

- In many countries the risk reporting requirements are for VaR values based on daily historic repricing of the current portfolio over some historical sample period
  - Simulate T day daily returns for N days. Take the 95th percentile worst return observation and multiply it by portfolio value.
  - Nothing is cumulative and the *fact that the future may not be like the past is not part of the conception of model risk*

Asset Owners are Not Banks

- Long term asset owners do not have any liabilities “at call”. Solvency is NOT the issue
  - Even pension “liabilities” are really the present value of liabilities often decades in the future. This is why underfunded pensions can continue to operate
  - Many long term asset owners such as sovereign wealth funds have no legal liabilities at all. It’s hard to go broke if you have say $500 Billion and don’t owe anyone anything
- Asset owner objective function is to maximize risk adjusted return over many periods
  - Markowitz and Levy (1979)
  - Short term downside risk as measured by VaR is irrelevant, except for some rare cases

Asset Managers are Not Banks Either

- Traditional asset managers are pure agents. There is no solvency issue
  - Good managers conform their activities to the objectives of their asset owner clients, meaning long term risk adjusted return.
  - You might argue that retail mutual investors need to be protected from themselves (panic selling during a downturn)?

- Hedge funds often use leverage but unlike a bank depositor, hedge fund shareholders are not promised a fixed redemption value.
  - The only solvency risk would be to the prime broker who already has custody of the assets and can legally liquidate assets at will to meet margin calls.

Semantics of Model Risk in Asset Management

- Analytical models in asset management *are predictive of the future* in an effort to enhance investor returns
- The need to distinguish between *risk* and *uncertainty*.
  - Risk is what can happen because returns arise from a known distribution. Think roulette.
  - Uncertainty is what can happen because the true return distribution is not known. Think playing poker.
  - We sub-divide uncertainty into estimation error and model risk
  - Estimation error arises because the world changes from the conditions in the data sample from which the model was built
  - Model risk arises because we did something wrong like using flawed data, improper statistical techniques or failed to incorporate an important feature (e.g. higher moments)

The Dumbest Question About Risk Models

- "What is the R-Squared of your model?"
  - In sample or out of sample?
  - Individual securities or portfolios?
  - Equal weighted statistics, cap weighted, something else?
  - What universe of securities?
  - What time period is covered?
  - If in sample, what observation frequency (daily, weekly, monthly)?
  - If out of sample what is the time horizon of the risk forecast? (a day, a month, a year)
  - Adjusted for higher moments or not?
- By framing the problem correctly I can get any answer I want to give out

Exceedance Tests of Risk Models

- This is what most regulators look for
  - Take your portfolio, and estimate what a one period 95% VaR (one day, one week, one month) would have been at various moments in history
  - Calculate the return for the period subsequent to the VaR estimate and convert to dollar loss
  - Count the percentage of events when your loss exceeds estimated VaR. If the percentage of exceedance events is bigger than 5%, the model is underestimating risk.
  - If you are assuming a normal distribution, this is just a return that is worse than a negative 1.645 standard deviation event
- Consider what happens if volatility on average is *known to be 9%* but varies randomly each day from 7% to 11%.
  - You can be exactly right on average but have 50% exceedance

Ratio Normality Tests

- Take your portfolio, and estimate what the tracking error or absolute volatility (monthly units) would have been at various moments in history. **Normality assumptions are not proper for shorter periods**
- Calculate the one month return for the period subsequent to the estimate. Divide each observed return by the related standard deviation (TE or absolute) to get the ratio
- Repeat the process for other portfolios and other time periods
- Look at the distribution of the ratio values. If the model is doing a good job, the distribution of the ratios will be approximately unit normal with mean close to zero and standard deviation one.
- *If the standard deviation of the ratio is less than one, you are overestimating risk. If the standard deviation of the ratio is greater than one you are underestimating risk.*

Cross Sectional Discrimination

- The most important test is whether we can discriminate between a low risk portfolio and a high risk portfolio.
  - Most investment management mandates require the funds be fully invested in a specific universe of securities
  - Simulate several hundred portfolios and benchmarks at a moment at least a year ago.
  - Randomly select securities and randomly switch weighting schemes between equal, cap and square root of cap weight
  - Calculate portfolio absolute and active performance for 12 months, rebalancing portfolios back to original weights monthly
  - At each of the rebalancing points observe the risk forecasts for tracking error and absolute volatility.
  - For major model overhauls we do thousands of years of portfolio simulations

Cross Sectional Discrimination Diagnostics

- Test TE, Absolute Portfolio Risk and Absolute Benchmark Risk across all portfolio/benchmark cases
  - What is the cross-sectional correlation between the TE forecast at the start of the period, and the realized TE over the subsequent 12 months. If this is high, the model does a good job separating high risk and low risk portfolios
  - You can also look at bias. Was the average risk forecast too high or too low? Check both absolute differences and differences as a percentage of forecast. *We prefer to be high*
  - What is the standard error of the forecast? (standard deviation of the difference between the forecast and realizations)
- Repeat the whole process using the time series average risk forecast rather than the starting point forecast

A Sample Output

-0.69408	-4.76598	1.090423	-2.08156			6.240271	6.450881	-0.21061		13.09072	11.78212	1.308602		13.78317	12.03092	1.752243
1.03433	1.82515	-0.6531	1.19695			4.226763	4.186008	0.040755		13.31263	11.42908	1.883549		15.56081	13.12767	2.433139
1.204983	-1.11859	-0.19292	-1.42571			4.47047	4.550398	-0.07993		16.30259	12.74244	3.560152		14.61033	10.99533	3.614992
-0.31854	0.69255	-0.24006	0.26044			1.853769	3.621345	-1.76758		13.59244	11.19288	2.399558		14.41371	12.86559	1.548126
0.03854	-2.08754	0.739227	-0.66974			3.950481	5.397343	-1.44686		15.59543	12.28511	3.310319		13.15718	11.0279	2.12928
0.11905	0.44503	-0.7486	0.09463			1.288657	1.183335	0.105322		13.77238	10.97266	2.799718		13.38144	10.87177	2.509672
-0.36518	-0.29254	-0.07238	-0.68597			2.558233	3.113148	-0.55491		14.28915	12.64795	1.641209		13.8734	11.27126	2.602144
-0.38983	0.39104	-0.59581	0.71605			2.097175	2.777655	-0.68048		14.99062	12.73091	2.259712		16.12416	12.7286	3.395558
-0.64956	-0.47019	1.09406	-0.168			3.865755	4.953447	-1.08769		16.13228	12.90876	3.22352		14.15875	11.12841	3.030335
0.22648	-0.2461	-0.40465	-0.54631			1.912885	1.810157	0.102729		14.32602	11.02714	3.298877		13.18207	11.14047	2.041606
0.00284	1.527	0.372785	-0.31116			2.126351	2.392305	-0.26595		13.58506	10.9106	2.674461		14.72546	11.27177	3.453693
0.94638	0.70401	-0.07016	0.93851			2.856656	3.619906	-0.76325		13.8695	10.94798	2.921523		14.67349	12.85641	1.817079
-1.35993	0.41988	1.013425	0.32676			3.934226	5.77113	-1.8369		14.52554	12.71645	1.809093		14.40937	12.08792	2.321453
1.454595	0.07573	-0.35159	1.36731			2.751822	4.110385	-1.35856		13.91536	11.08522	2.830134		14.43398	12.42103	2.01295
					Average	3.082947	3.791251			14.22534	11.83587			14.4168	12.04448	
					Correl(Init)	0.817071				0.638245				0.594795		
					Correl (Average)	0.835842				0.672253				0.638854		
					StanErr Corr	0.030443				0.055364				0.059787		
					T	27.4561				12.14241				10.68542		
					% Bias	22.97491				-16.7973				-16.4552		
					Confidence			0.790713				0.721117				0.889841
					% Confidence			20.85626				6.092636				7.387956

Client Testing

- SR 11- 7 has some interesting requirements
  - The user organization must show that the risk model fits their particular usage, not just that the vendor tests generic cases
  - The user organization must also be able to show a thorough understanding of vendor models. No black boxes allowed.

- Clients often do some sort of test comparing forecasts and realizations on a single portfolio (theirs)
  - This *might* satisfy SR 11-7 but it's hopeless for comparing models from different vendors. The sample is just to small to be statistically meaningful.
  - For testing different vendor models, look for outliers in the risk estimates *on the same portfolio for the same time horizon.* If Model A says 8%, Model B says 8.1%, Model C says 7.8%, and model D says 3% you might get an idea of what not to buy.

Avoiding Two Risk Systems

- Many regulatory organizations require short horizon VaR values as compliance reporting.
  - Some of these have time horizons as short as one day (Mexico)
  - Some of these require that the process arises from some form of historical or Monte Carlo simulation
- There are several different ways Northfield can fulfill these requirements without clients having to deal with another vendor
  - We have "near horizon" versions of all models that have recently enhanced with our Risk Systems That Read[sm] technique.
  - There are multiple computational modules available to provide the required analysis depending on applicable regulations
  - Avoids conflicts and confusion over two disparate systems

Short Horizon Parametric VaR

- The simplest formulation for T day VaR(p) from any parametric distribution can be framed as:

$$VaR(T,p) = S * Z(p) * (T/252)^{.5} * V$$

S = annualized forecast for portfolio volatility

Z(p) = the number of standard deviations for cumulative probability P

T = the number of days in the VaR horizon

V = dollar value of the portfolio

The helpful bit here is the S can be estimated from longer observation data builds in the serial properties of daily returns, so the square root of time assumption suffices

One Change Handles Most of the Problem

- The most obvious mistake is to assume that returns are normally distributed.
  - Demonstrably untrue for daily financial market returns
  - A popular and well researched alternative is the T-5 distribution.
  - See diBartolomeo, Dan. "Fat Tails, Tall Tales and Puppy Dog Tails", *Professional Investor*, Autumn, 2007.
  - For the normal, Z(95%) = -1.645 Z(99%) = -2.33
  - For the T-5 distribution, Z(95%) = -2.02, Z(99%) = -3.38

- If we make no assumptions on the shape of the distribution we can invoke Chebyshev's inequality
  - Z(95%) = -3.165, Z(99%) = -7.07

Incorporating Historical Calibration

- Many regulatory schemes require that VaR be derived from a historical simulation of the current portfolio
  - To reconcile the forward looking forecast from the risk model to the historical simulation, we need to calibrate our estimate of portfolio volatility S against historical experience
  - Take your existing portfolio and calculate it's annualized historic return volatility over some past period (e.g. 24 months).
  - Compare this to the average of the volatility forecasts that you would have gotten during the sample period.
  - Let B = historic realization / average estimate
  - This is really easy to implement in our Performance Attribution module.

$$VaR(T,p) = S * B * Z(p) * (T/252)^{.5} * V$$

Conclusions

- Regulatory schemes impacting the management of institutional assets have been largely translated from banking and focus on short term downside risk, i.e. VaR
- As long term investors are appropriately less interested in short horizon downside risk than banks, a divergence has developed between forecasting risk for the purpose of investment decisions, and forecasting risk for regulatory reporting.
- Some regulations explicitly address "model risk". Compliance requires demonstrations of model validation and testing. Northfield regularly conducts a variety of rigorous tests on our models, and can do more specific analysis on client request
- We have suggested an approach to estimation of short horizon VaR which we believe is robust, and easily implemented. It will satisfy most regulatory requirements, while avoiding the cost and confusion of disparate but parallel risk systems.