Proper Risk Assessment and Management: The Key to Successful Banking

ORACLE WHITE PAPER | NOVEMBER 2017

Table of Contents

Author André Koch is managing director of Stachanov Solutions & Services BV in Amsterdam, an instructor at Oracle University, and a senior lecturer in bank financial management and risk management at Nyenrode University, The Netherlands Business School (andre@stachanov.com).

Executive Overview

Proper risk assessment and management is essential to successful banking. Often, the risk- modeling process is farmed out to the bank's quantitative modeling departments, where well- trained Ph.D. teams of specialists take on the number crunching. This paper will make a case for bringing some risk-modeling tasks back to the decision-makers, demonstrating how Monte- Carlo simulations can be carried out on the desktop using Oracle Crystal Ball software.

Oracle Crystal Ball is a Microsoft Excel plug-in that allows users to apply sophisticated Monte- Carlo modeling techniques on the desktop platform bankers are most familiar with. Even complex concepts such as the calculation of a bank's economic capital can be modeled without much effort. In a Monte-Carlo simulation, a problem is addressed in a probabilistic way rather than through a straightforward deterministic calculation. In this probabilistic model, calculations are repeated many times with different, randomized values used for the input variables. By analyzing the results of such simulations, bankers can glean useful insights into uncertainties and risks.

Introduction: Capital, and More Capital

In banking, the term *capital* is used in numerous ways, many of which are difficult for nonbankers to distinguish. First, there is the accounting concept of capital, which refers to shareholders' equity on the balance sheet. Then, there is the regulator's perspective on capital—with *regulatory capital* referring to the buffer the bank should maintain to protect itself against adverse conditions.

These regulatory rules stem from the international Basel Accords, which set forward the capital adequacy conventions in banking. National governments have adopted this Basel Framework and charged national banks seeing that the banking sector is in compliance with the Basel rules. Although the Basel II and Basel III Accords can hardly be classified as light reading material, the essence of the rules is straightforward.

This is not true for the third capital concept: *economic capital.* The economic capital buffer refers to the financial shock absorbers maintained by banks to prepare for unexpected risks. In this paper we will zoom in on the essence of economic capital and describe how it can be modeled and calculated.

Probability of Default

Whenever a bank issues a loan, it assesses the risk of the obligor being unable to meet his or her obligations to the bank. Known as the *probability of default* (PD), this risk is expressed as a percentage indicating the odds that the client will default on his or her obligations. This chance multiplied by the loan sum is the amount the bank will set aside to cover the risk of default.

With many outstanding loans, the bank can pool these risks. On average, the pool of provisions for default will cover the default losses that occur—referred to as the *expected losses* (ELs). However, averages can be deceiving: putting one foot in a bucket of ice cubes and the other foot in a pot of boiling water does not necessarily make one feel comfortable.

The same caution applies to the pool of default provisions. In a way, default risk itself does not pose much of a risk for the bank. It is like a traffic jam that forces you to add ten minutes to your daily travel routine to reach office on time. As long as the traffic is jammed every morning and the extra commuting time always adds ten minutes, this can hardly be classified as a risk. You know it will be always there, and you have factored it in to your estimated your travel time. The same goes for probability of default. Loan defaults can hardly come as a surprises to a bank, since the risk of just such events was considered in setting the interest rate and creating the provisions capital buffer.

Referring back to our traffic jam scenario, the real risk for the commuter comes when the delay from the traffic jam varies wildly, forcing her to start her drive much earlier. Similarly, in the case of the bank, the PD percentage itself does not pose a risk, but the variance in the probability of default does. If the real defaults exceed the expected defaults this may cause the bank to end up in dire straits. The variance in PD forces the bank to keep an additional financial buffer to absorb these unexpected losses. The buffer to protect the financial institute against expected losses is called the economic capital. Before we can proceed with calculating the economic capital, we have to refine our model by introducing the concepts of *loss given default* and *correlations.*

Loss Given Default

A bank does not lose the complete loan amount each time a default occurs. Usually, it can recover at least part of the claim—for instance, by selling off the collateral pledged as part of the agreement. The average percentage actually lost is called the *loss given default* (LGD). The amount set aside for expected losses is reduced by the LGD—which results in the following formula for expected losses:

EL = loan amount * PD * LGD

Correlations

A second correction concerns the correlation between losses. Frequently, defaults are not unrelated, independent events. The banking crises show that trouble rarely comes alone—with shifts in the housing market, for example, leading to a deluge of mortgage defaults. This correlation increases uncertainty and diminishes the portfolio benefits of spreading the default risk over multiple loans. Clearly, these correlations are major risk drivers and should be included in the model.

Economic Capital

Now, it's time to take a closer look at the economic capital model for our bank, which has a U.S. $6 billion loan portfolio spread over three economic sectors: IT, telecoms, and retail. For each sector, we know the PD and LGD—numbers that can be obtained by analyzing historical data (which the bank has access to), seeking expert opinions, or contacting rating agencies.

	A	B	C	D	E	F	G
1	**Sector**	**Exposure**	**Probability of Default (PD)**	**Stand. Dev. PD**	**Loss Given Default (LGD)**	**Stand. Dev. LGD**	**Provisions**
2	IT	1,500,000,000.00	1.50%	0.25%	60.0%	2.00%	13,500,000.00
3	Telecoms	2,000,000,000.00	1.00%	0.15%	50.0%	2.00%	10,000,000.00
4	Retail	2,500,000,000.00	0.50%	0.12%	40.0%	1.00%	5,000,000.00
5	**Totals**	**6,000,000,000.00**					**28,500,000.00**

Figure 1. The bank's loan portfolio, with probability of default (PD) and loss given default (LGD) showing for each sector.

To calculate the provisions for expected losses, we simply multiply the loan amount (exposure) by the PD and LGD.

How Oracle Crystal Ball Can Help

As indicated previously, the bank's real risk comes from deviations from the average. In other words, the higher the variance in the possibility of default, the higher the risk. Oracle Crystal Ball software allows users to model the behavior of the PD IT-sector input variable by associating this parameter with a probability distribution. The behavior of the PD is also skewed to the right, because on the left side zero provides a boundary, while on the right side, the maxim can go up to 100 percent.

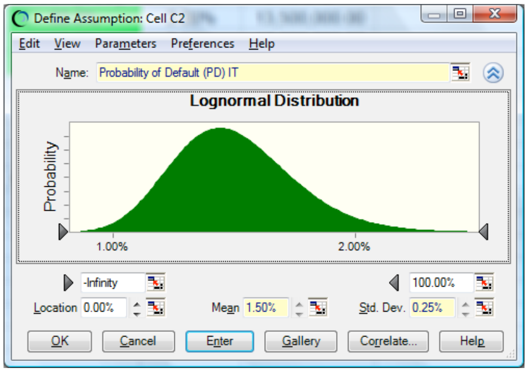

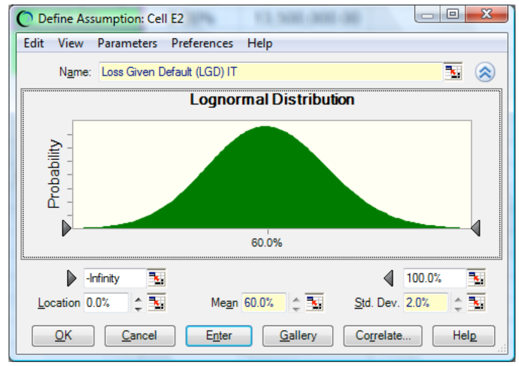

Figures 2 and 3. On the left Oracle Crystal Ball uses lognormal distribution to model the PD for the ID sector; on the right it uses lognormal distribution to model the LGD for the IT sector.

Similarly, PDs can be modeled for other industries, and the same approach can be used for LGDs.

With the input parameters modeled, we only need to pinpoint the output variable that we would like to calculate. In our case, this is the total of provisions. Oracle Crystal Ball colors the input variables green and the output variable sky-blue.

Sector	Exposure	Probability of Default (PD)	Stand. Dev. PD	Loss Given Default (LGD)	Stand. Dev. LGD	Provisions
IT	1,500,000,000.00	1.50%	0.25%	60.0%	2.00%	13,500,000.00
Telecoms	2,000,000,000.00	1.00%	0.15%	50.0%	2.00%	10,000,000.00
Retail	2,500,000,000.00	0.50%	0.12%	40.0%	1.00%	5,000,000.00
Totals	6,000,000,000.00					28,500,000.00

Figure 4. Oracle Crystal Ball model with input variables in green and output variable in sky-blue.

In the case of the IT sector the provisions amount to 1.5 billion *times* 1.5 percent *times* 60 percent, or

13.5 million. Repeating this calculation for each of the economic sector results in a total of 28.5 million in provisions for our bank.

	Sector	Exposure	Probability of Default (PD)	Stand. Dev. PD	Loss Given Default (LGD)	Stand. Dev. LGD	Provisions
1							
2	IT	1,500,000,000.00	1.50%	0.25%	60.0%	2.00%	=B2*C2*E2
3	Telecoms	2,000,000,000.00	1.00%	0.15%	50.0%	2.00%	10,000,000.00

Figure 5. Calculating provisions for the IT sector.

Running Simulations

Now, it's time to start spinning the roulette wheel: Oracle Crystal Ball will generate random numbers for each of the input variables. This random process is nonetheless governed by the distributions associated with these input parameters. As the input variables take on different values, the output parameter and the sum of the provisions will change as well. Thus, if we were to run 10,000 trials in a simulation, we would end up with close to 10,000 different results, while our original Excel model would produce just one outcome. We can then display those 10,000 results in a histogram with the results grouped in bins. This histogram is basically a frequency diagram with the various amounts for provisions on the horizontal axis and the rate of occurrence on the right-hand vertical axis. It is important to realize that a higher frequency translates to higher bars in the graph and an increased relative probability.

Looking at the histogram, you will notice that the majority of the observations stay close to the original Excel amount of 28.5 million. Still, the results are scattered over a wide range, indicating a fair amount of risk. Ideally, the results are clustered in the middle: the more spread out the graph is, the higher the risk. Often, the standard deviation is used as a proxy for risk. Oracle Crystal Ball tells us that the standard deviation is close to 3 million (see Figure 7).

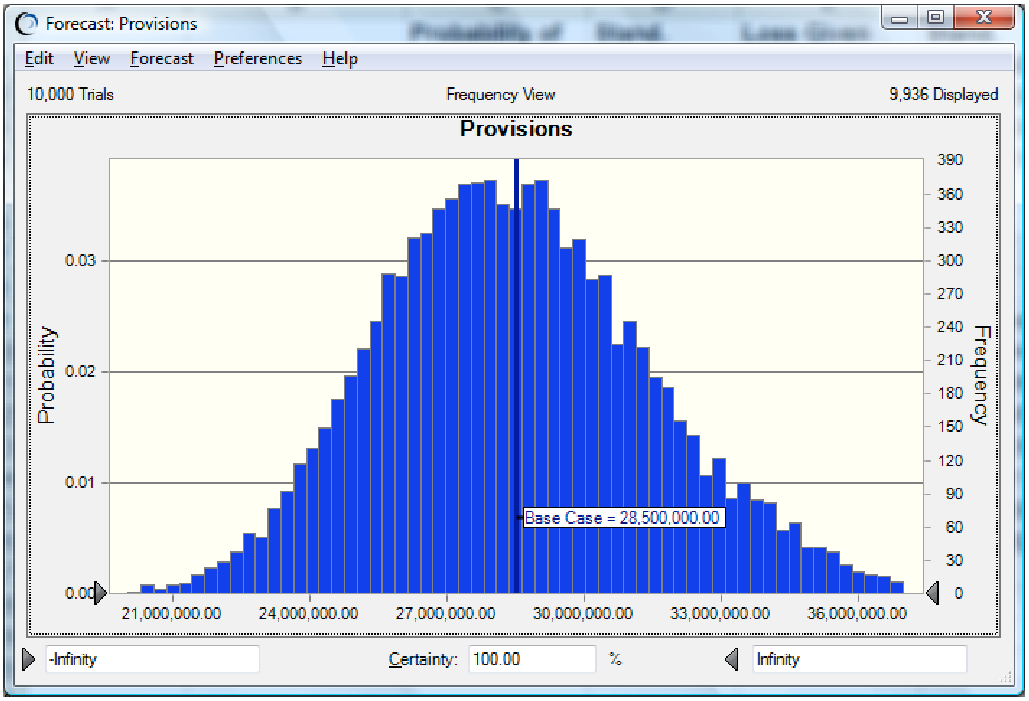

Figure 6. Ten thousand simulation results grouped in bins with the original Excel result as the base case.

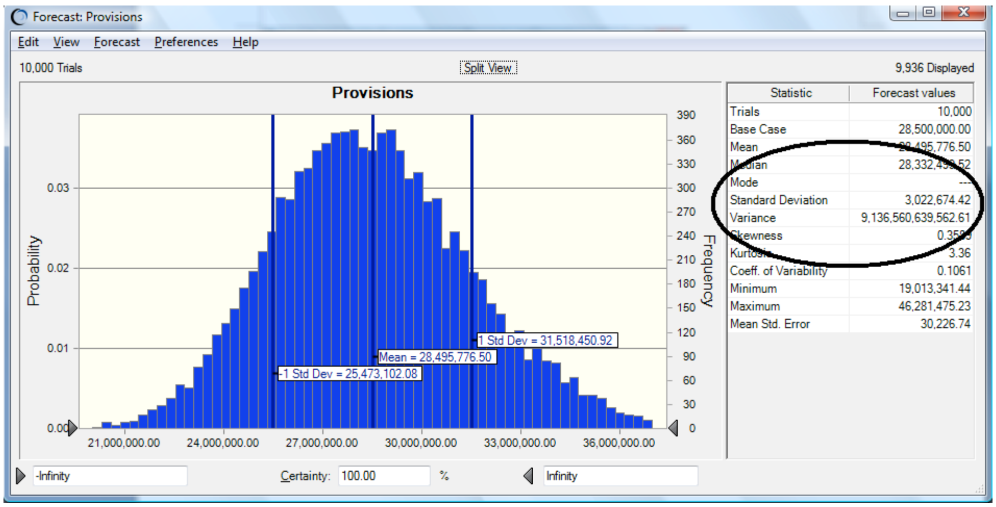

Figure 7. Standard deviation of 3 million below and above the mean.

Required Economic Capital

Now, it's time to return to our original question—that is, how much economic capital should this bank maintain to fend off unexpected losses (ULs) *in addition to* the provisions it has already set aside to protect against *expected* losses (ELs)? Basically, the unexpected losses can be seen to the right of the mean in the forecast histogram. Obviously, if we were to increase the amount of provisions, more of the possible variance in outcomes would be covered. For example, if we were to increase the provisions from 28.5 million to 30 million, we would cover more than 70 percent of the surface of the graph, and hence of the risk.

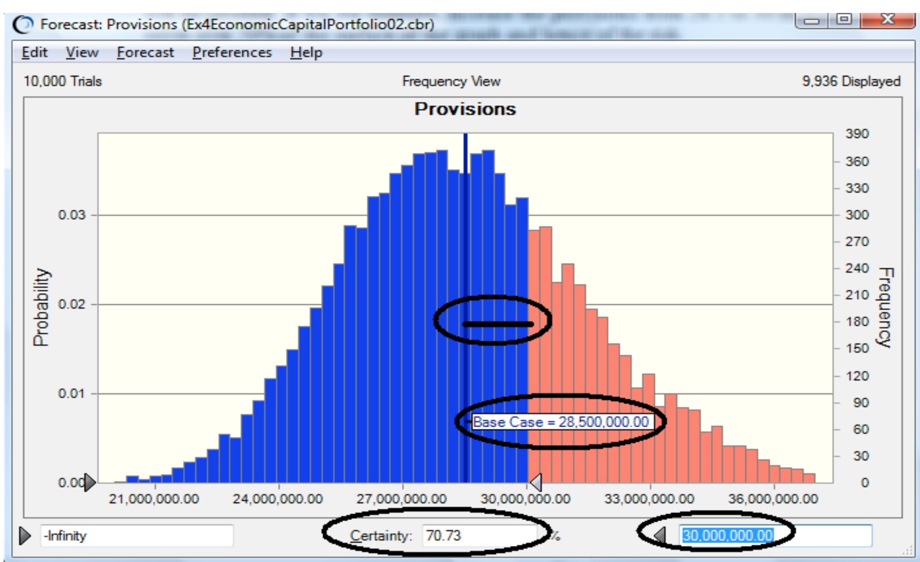

Figure 82. With provisions increased to 30 million, 70 percent of the bank's risk would be covered.

In the Oracle Crystal Ball forecast chart shown in Figure 8, the area in blue represents the risk covered and the area in red represents the uncovered risk.

Risk Appetite

Now we get to the crux of the model. The economic capital is the difference between the provisions depicted as the base case in the graph and the cut-off point on the horizontal axis for a certain risk level. In our example, the risk covered is 70.73 percent, which corresponds with a cut-off point of 30 million. Since, 28.5 million of these 30 million were already accounted for in the provisions, this will leave 1.5 million for the economic capital to serve as a buffer for unexpected losses. If we demand more certainty and less risk, we should allocate more economic capital—making more of the graph's surface blue. As anticipated in our traffic jam analogy, if the traffic jam duration varies wildly, the commuter needs to leave home earlier—how much, depends on how much the traffic jam duration varies. Since one runs the risk of job loss for reporting late to work, wider buffers are called for in the traffic jam analogy. The amount of risk one is willing to stomach in a given scenario is known as *risk appetite.*

Credit Ratings

Returning to the bank scenario, risk appetite is determined by a bank's credit rating. A triple-A bank will maintain higher capital buffers than a single-A bank. These credit ratings correspond to certainty levels in the Oracle Crystal Ball forecast chart as shown in Figure 9 (a higher certainty level calls for more capital). In reality, most banks do not allow for more than a couple of basis points of uncovered risk, which takes the certainty levels for the most common credit ratings to more than 99 percent.

Rating Bank	Certainty level %	Uncov'd Risk BP	Total Capital	Provisions	Economic Capital
AAA	99.99	1	42,112,650.19	28,500,000.00	13,612,650.19
AA	99.97	3	40,741,947.25	28,500,000.00	12,241,947.25
A	99.93	7	39,466,961.56	28,500,000.00	10,966,961.56
BBB	99.78	22	38,353,746.36	28,500,000.00	9,853,746.36
BB	98.85	115	36,005,875.74	28,500,000.00	7,505,875.74
B	93.90	610	33,470,215.26	28,500,000.00	4,970,215.26
CCC	81.73	1,827	31,171,885.39	28,500,000.00	2,671,885.39

Figure 9. Economic capital for banks with ratings from AAA through CCC.

Thus, a bank's economic capital need depends on its risk appetite. In our example, a triple-A bank would need about 13.6 million, while a triple-B bank could get away with just 9.9 million.

Besides the size of the loan portfolio and the credit rating of the bank itself, one of the main motors of economic capital is the uncertainty surrounding unexpected losses. But what is the source of this uncertainty? Oracle Crystal Ball's sensitivity analysis identifies the main risk drivers and ranks them by their contribution to the variance of the provisions. Sensitivity analysis is a useful tool in risk mitigation, providing a handy to-do-list for further action.

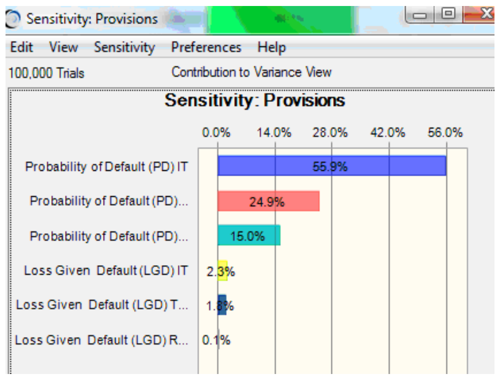

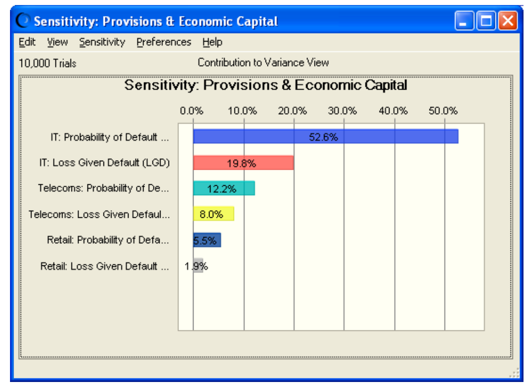

Figures 10 and 11. From the sensitivity analyses shown in these Oracle Crystal Ball graphs, it's clear that the sector needing the most attention is IT—particularly as it relates to probability of default.

In this case we can see it is the IT sector we would be best focusing our attention and in particular on the uncertainty in PD.

There is one more thing to add to the model: correlation. Oracle Crystal Ball allows for the input of correlations on a variable-per-variable level; however, it can also read a correlation matrix.

	A	B	C	D	E	F	G
1		PD IT	PD Telecoms	PDfRetail	LGD IT	LGD Telecoms	LGD Retail
2	PD IT	1.000	0.900	0.100	0.500	0.500	0.100
3	PD Telecoms		1.000	0.100	0.500	0.500	0.100
4	PD Retail			1.000	0.100	0.100	0.100
5	LGD IT				1.000	0.500	0.100
6	LGD Telecoms					1.000	0.100
7	LGD Retail						1.000
8							

Figure 12. Correlation matrix.

In the economic capital calculation, correlation has an impact on the risk level of the bank and leads to considerably higher levels of economic capital.

Rating Bank	Certainty level %	Uncov'd Risk BP	Total Capital	Provisions	Economic Capital Correlation
AAA	99.99	1	49,371,718.25	28,500,000.00	20,871,718.25
AA	99.97	3	47,915,016.07	28,500,000.00	19,415,016.07
A	99.93	7	46,555,625.02	28,500,000.00	18,055,625.02
BBB	99.78	22	44,013,285.45	28,500,000.00	15,513,285.45
BB	98.85	115	40,271,603.73	28,500,000.00	11,771,603.73
B	93.90	610	35,970,279.25	28,500,000.00	7,470,279.25
CCC	81.73	1,827	32,512,667.10	28,500,000.00	4,012,667.10

Figure 13. Economic capital under correlation.

In the case of a triple-A bank, we notice that the economic capital jumps from 13.6 million to almost 21 million due to correlation effects.

Conclusion

Oracle Crystal Ball is a useful tool for modeling capital adequacy. Ease of use and the interaction with Excel indicate that the bank's decision makers can take on more of the model creation themselves and break the artificial barriers between the quants and management.

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.303.633.0947
www.oracle.com/crystalball

CONNECT WITH US

- blogs.oracle.com/oracle
- facebook.com/oracle
- twitter.com/oracle
- oracle.com

Integrated Cloud Applications & Platform Services

Proper Risk Assessment and Management: The Key to Successful Banking
November 2017
Author: Andre Koch