Analyzing Concentration Risk

Diane Reynolds

Abstract

Concentration risks, particularly concentrations in credit risk, have played a key role in the financial instability of the banking sector in 2008. This paper examines different objectives in managing credit concentration risk. The suitability of different measures and presentation techniques is discussed and illustrated in the context of a case study. We also look deeper into "measures," examining the contributions and interactions of migration, name and sector concentration risk on the portfolio.

Introduction

Credit risk is an important risk faced by banks. The Basel Committee devoted a significant amount of attention to this type of risk when creating Basel II. The models used in Basel II calculations under Pillar 1 are formulaic in nature, facilitating standardization and implementation by making assumptions. One of the key assumptions in the formulae for capital calculation is full diversification—the assumption that no concentrations are present in the portfolio. For this reason, credit concentration risks are a focus of supervisory review under Pillar 2 of Basel II.

More business-focused models of credit risk tend to include concentration risk. Without considering existing portfolio concentrations, the signals communicated by the model to business decisions become corrupted. Consider a case where two investments are possible. Both have exactly the same profit profile and characteristics when input to Basel II, Pillar 1 formulae: product type, default probability, maturity horizon and loss given default (LGD). However, one investment lies in the sector in which the bank is highly concentrated while the other is in an emerging business area for the bank. The second investment clearly adds more diversification, creating an overall more optimal portfolio.

Given the importance of credit concentrations to both capital reviews by regulators and business decisions, this paper attempts to consolidate into a single location a variety of measurement and management techniques. Herein we provide an overview of the measurement, attribution and representation of credit risk concentrations, illustrating our discussions with a case study.

Measuring Credit Risk

In order to assess concentrations in credit risk, it is important to first measure credit risk. Many measures have been used and developed to communicate credit risk levels. For simple loans, it is often enough to state the amount of the loan—either using face value or amortized value. However, as financial instruments grew in complexity over the course of the last three decades, these concepts lost their power to accurately represent credit risk across all facets of typical portfolios today.

The measure most often used to express risk levels in general is exposure. Credit risk is no exception. Exposure indicates the level of loss should a credit event occur, in the absence of mitigating factors. Typically, it is calculated as the total amount at risk, adjusted for mitigating factors such as netting and collateral. In the case of derivatives transactions, both the total risk and the mitigating factors may be estimated using simulation techniques.

But, is exposure really a suitable measure of risk? If the bank is equally exposed to two firms, A and B, are A and B equally risky? Of course, the answer is no. Other factors determine the credit quality of the name and the nature of the exposure. These factors are typically incorporated into measures of probability of default (PD), or rating, and LGD.

To get a better picture of risk we might use a measure such as expected loss so as to include all three factors: PD, LGD and exposure. Expected loss is commonly used to set spreads for loans, quantify provisions or estimate profitability.

Using expected loss as a risk measure in these applications, however, is not without its own issues. Specifically, it is an average value, and so underestimates actual potential losses about half the time.[1] To more conservatively assess risk, we must look beyond the average.

TABLE 1
Example Measures for One Name

Measure	HSBC Holdings PLC
Exposure	3,462,675,500
Expected Loss	6,942,475
CVaR (99.9%)	436,616,264
Relative CVaR (99.9%)	12.61%

This means looking at cases where the overall portfolio experiences problems; specifically at groupings of defaults and downgrades. The tendency for losses to bunch together (and create problems) is quantified through correlation. Capturing correlation, along with the other key factors of exposure, PD and LGD, leads to a measure such as credit value-at-risk (CVaR[2]) defined at a particular confidence level.

1 This assumes that the mean and the median of the actual loss distribution are approximately equal. While this may not be the case, the general principle (that expected loss underestimates losses) is generally true.

2 In this context we have defined CVaR(quantile) as the unexpected loss at the specified quantile. The mean (EL) is subtracted from the total loss at the quantile to attain unexpected loss.

Depending upon the objective, the confidence level may vary. For example, management decisions are typically made on a fairly short-term basis, for example 10 years, translating to a confidence level of 90 percent. However, capitalization levels must account for much less common events. Hence, the most commonly used level is 99.9 percent; it matches the Pillar 1 standards of Basel II for minimum capital calculations. Many financial institutions maintain a higher rating than the BBB average of Basel II, which drove the choice of the 99.9% confidence level. Accordingly, they seek a higher confidence interval for CVaR; for example, an AA rating would indicate a 99.98 percent level.

Relative measures can also be informative. Considering any of the above values not in absolute terms but as a percentage would highlight disproportionate risk allocations. In some cases, comparing the relative shares of CVaR (99.9%) against the relative shares of exposure along the various dimensions can yield additional insights by identifying the most risk-expensive names per dollar of exposure.

Measuring Credit Concentrations

When searching specifically for concentrations, combinations of measures can uncover hidden information. One such measure is the name concentration ratio. The name concentration ratio is defined as the name level CVaR divided by the absolute loss.[3] Representing the concentration ratio in a scatter plot comparing it to relative CVaR provides particular insights. Names that have the lowest concentration ratios provide natural hedging in the portfolio, while those with the highest concentration ratios are good targets for increased monitoring or hedging. Of course, this analysis is most useful in the context of material exposures (hence the second dimension of the plot). Such information allows risk managers not only to select names that would be good hedges, but also to select the ones that have a low cost of capital per dollar of exposure.

It is possible to investigate further which names add concentration and diversification to our portfolio. To do so, we look at the concentration index[4] (CI) of each name in our portfolio. Loosely speaking, the CI is a benchmarked version of the name concentration ratio. It converts the calculated name concentration ratio into a standardized measure by comparing it to the portfolio average. A name with a CI >1 adds concentration risk to the portfolio while those with a CI < 1 add diversification to the portfolio. This standardization increases comparability across portfolios and across time.

Figure 1
Exposure Versus CI 99.9%

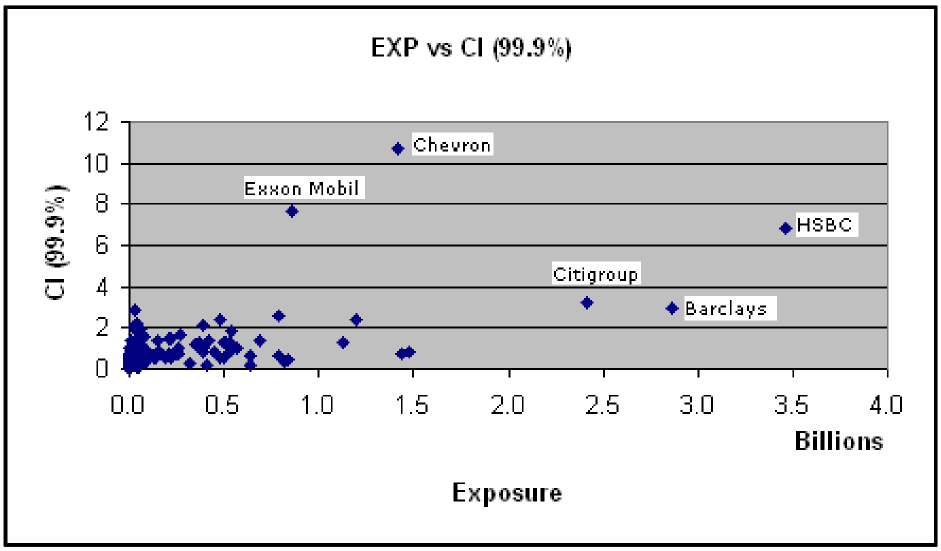

[3] The absolute loss for name N at quantile Q is defined as the CVaR(@Q) for a portfolio containing only the exposures to N that appear in the overall portfolio. This measure is also known as “stand alone CVaR.”

[4] The concentration index was developed by Nasakkala and Agustsson at Nordic Investment Bank circa 2004. To calculate the concentration index, define $CVaR_i$ as the credit value at risk contribution of name ‘i’ within the portfolio, and $ABSL_i$ as the standalone credit value at risk when name ‘i’ is considered in isolation. In the case study portfolio of 500 distinct names, the concentration index for name ‘i’ is given mathematically by

$$(CVaR_i / ABSL_i) / \left(\sum_{j=1}^{500} CVaR_j / \sum_{j=1}^{500} ABSL_j \right)$$

The CI can be plotted against exposure to identify which names are good candidates for investment. Risk managers can also use the CI to help identify new business partners. For example, if a bank would like to do new business with a new partner who has low relative CVaR, this could be a good investment for the bank, but it may not be. If this new partner has a very high CI and adds a substantial amount of concentration risk to the portfolio, the bank may want to reconsider.

The choice of measure can materially impact the concentration risk reported for a portfolio. From loan amount outstanding to exposure, expected loss to CVaR and on to relative measures, the number of factors considered increases. However, this leads to a larger number of input parameters to be calibrated, increasing the inherent uncertainties in the numbers themselves. It is thus a trade-off, and using a variety of measures is likely to lead to the most complete information.

While important, the actual numerical results are not the only factor in creating and using information on credit concentrations. Relying on humans to pick out numbers from vast quantities of data is ineffective—no matter how good the numbers. For this reason, it is worthwhile to consider the possible representations of concentration risks.

Representing Concentrations

In a search for simplicity, it is tempting to reduce concentration risk to a single number, or index. Various indexing techniques have been examined in the credit risk literature. All have a common approach: to identify the extent of the concentration in a portfolio through a single measure. One of the most straightforward indexes is the Herfindahl-Hirschman Index (HHI).

Originally used in the context of quantifying diversification within an industry to assess the level of competition in the marketplace, the HHI can also be used to calculate portfolio concentration risk. The HHI is calculated by summing the squares of the portfolio share of each contributor. Two different decisions must be made: what is a contributor and from which measure should the shares be calculated.

For example, a contributor to concentration risk might be an individual name in the portfolio, a sector, rating grade class, geographical area or product type. Assessing the HHI in varying dimensions can yield very different results for the same portfolio. For example, an equal investment in each of 100 companies would lead to a diversified HHI of 0.01. In contrast, if the firms are divided amongst five sectors in the ratio 5:2:1:1:1, then the implied HHI by sector is 0.32, indicating a significant concentration.

Like other indexes, the HHI suffers from two key issues: it provides different results for the same portfolio depending upon the dimension measured and it fails to provide any actionable information. Specifically, the HHI of 0.32 does not hint at ways to lower the concentration. So, the temptation of simplicity may be outweighed by the need for directed action.

TABLE 2
HHI & HHI Contributions by Sector

Industry Sector	HHI	Contributions to Portfolio HHI
Telecommunications	0.373	0.010
Health care	0.301	0.001
Utilities	0.168	0.012
Materials	0.134	0.002
Consumer staples	0.096	<0.001
Financials	0.080	0.641
Energy	0.042	0.330
Consumer discretionary	0.041	<0.001
Information technology	0.035	0.001
Industrials	0.022	0.002

Moving from a single number to a ranked list of "high risk" items is a logical step towards more actionable, granular information. The concept of a "Top 10 List" is simple: pick the 10 names with the largest (or smallest) values and manage them more closely than the others. Most firms do this with their largest exposures: extra care is taken in managing the relationship with these large clients or counterparties, more diligence is exerted in operations involving large accounts, and senior management is more engaged in the details of the exposures and relationships. When asked, almost any bank can produce its list of top exposures.

Armed with a Top 10 List managers can make better-informed lending, trading and relationship management decisions regarding key names in the portfolio. Alternatively, exposures beyond management comfort levels might be mitigated in the secondary markets, for example, individually through a credit default swap or collectively through securitization.

The saying “a picture is worth a thousand words” leads us from tables to graphs. Simple pie charts mimic the Top 10 list idea, showing contributions to the whole from key elements. Combining relative measures with graphical representations informs immediate impressions of concentrations.

Decompositions along potential concentration dimensions aren’t the only graphing tool available. Often it is a combination of factors that helps to identify hedges that would improve the portfolio risk profile or identify underlying risks. Such comparisons are typically represented in scatter plots.

Attributing Risks

While name and sector concentrations are the most commonly discussed concentrations, manipulating the risk profile to best serve the needs and preferences of an organization is a multidimensional task. Because CVaR incorporates many facets of credit risk into a single number, it is very useful in the ranking and relative analysis—as we have already seen. However, CVaR is a complex calculation. Typically, models from which it is estimated acknowledge and capitalize various types of risk and the interactions between key risk types.

Thus, it is also important to understand the sources of risk in the portfolio at a more granular level. This can be done through the process of capital attribution. Estimates of CVaR at the portfolio level can be attributed to different risk types (or sources) by recalculating using different combinations of assumptions to isolate particular risks. Usually, the most influential risks are: default risk, migration risk, name risk and sector risk. It is possible to isolate each of these risks in turn by varying the model assumptions.

The information gleaned from such analysis, and the details of its implementation are discussed in the context of the case study to provide additional clarity.

Case Study

To illustrate the concepts, measures and applications we make extensive use of a case study based on a portfolio of 500 exposures. The remainder of the paper presents the case study data set and then uses it to illustrate a series of concentration risk assessments.

The case study is based on an international portfolio of 500 publicly traded and rated names. The overall exposure of the portfolio is approximately U.S. $44.5B, with individual exposures ranging from just over $1M to almost $3.5B. The average exposure is approximately $88M but the median exposure is only $10M. Other studies[5] have shown that such skewness is typical of trading portfolios, where only about 10 percent of names carry 90 percent or more of the exposure. In this case study, 10 percent of names represent a more modest 77 percent of total exposure.

Exposure breaks down into six major ratings grades[6] and 10 major industries (Dow Jones) across 10 countries.

[5] Higo (2006).
[6] Default probabilities are based on ratings by Fitch Ratings in June 2007.

Figure 2
Exposure and CVaR (99.9%) by Industry

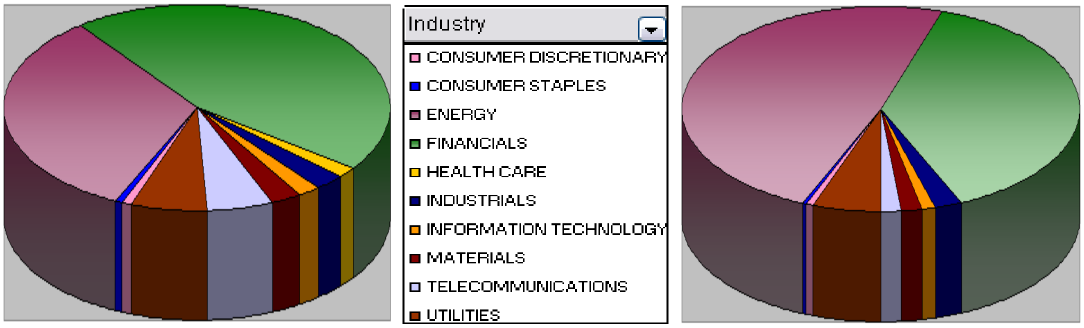

Figure 3
Exposure and CVaR (99.9%) by Rating Grade

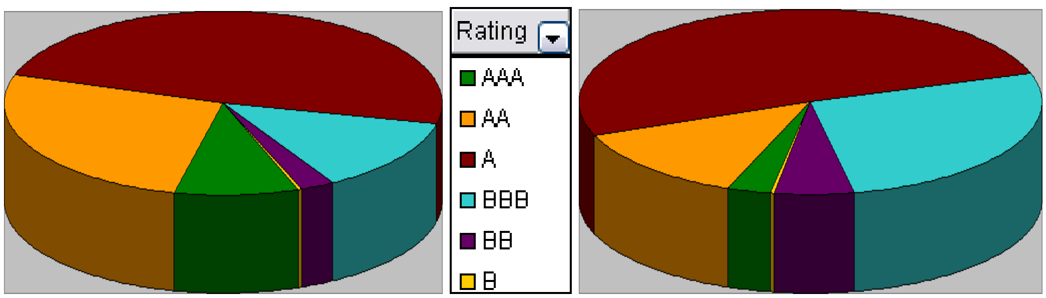

Figure 4
Exposure and CVaR (99.9%) by Country

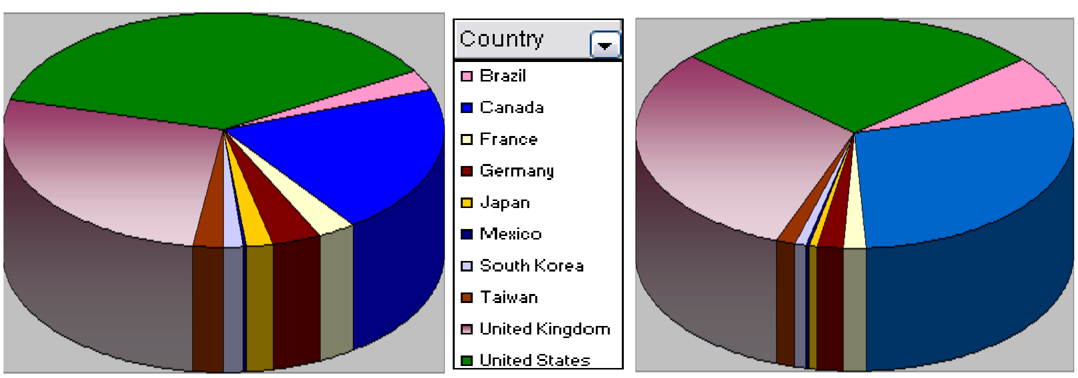

Typical of a trading book, clear concentrations are apparent in the financial and energy sectors, in single-A-rated firms and in the United Kingdom, United States and Canada. A few counterparties are its main trading partners: other banks, while the majority of its clients (for derivatives) may lie in a single sector or geography where the bank is dominant or particularly focused.

We begin by examining common portfolio-level index measures of concentration risk before moving on to the ever-popular “Top 10 list” where we compare the usefulness of different measures in constructing such a list. Measures are then compared further through the use of scatter plots and bar charts.

Next, capital attribution to various model assumptions such as the inclusion or exclusion of migration risks and the benchmark of full diversification permits more detailed analysis of the source of concentrations. Finally, the concept of concentration index is revisited, but with a twist based on capital allocation techniques.

Looking at shares of exposure for each name in the portfolio using the HHI yields **HHI = 0.023.** By comparison, a perfectly diversified portfolio of 500 names would have HHI = 0.002, an order of magnitude smaller than the case study example. By using industry shares (instead of name shares) of the portfolio across 10 industries, HHI = **0.321** indicating a significant concentration compared to the fully diversified results of HHI = 0.100.

To assess which industries might be contributing concentration risk to the portfolio, we can examine the industry contribution to HHI based only on exposures. We might also calculate the name-level HHI for each industry independently to assess intra-industry concentrations.

Table 2 shows the individual HHI for each of the 10 industry sectors, sorted from highest to lowest, and the contribution of the sector to overall portfolio HHI.

The contrasting information here is interesting: telecommunications are highly concentrated, but add little at the portfolio level. In contrast, financials are well-diversified, but add greatly to the overall portfolio concentration. To act upon such results, a portfolio manager might look into individual name credit default swaps within telecommunications and health care, but take a basket approach for financials and the energy sector. If forced to choose, however, the decision as to which of these four possible actions should take priority remains difficult to answer.

The

Top 10 List by Exposure is shown in Table 3. The industrial classification of the names in this list helps to explain the industry rankings of concentration by the HHI. While most of the Top 10 names are in the financials sector, we do not see concentration within this sector. The large number of large exposures in the sector is creating diversification within the sector, even though the sector is a source of concentration within the overall portfolio.

TABLE 3
Top 10 List by Exposure

Name	Sector	Exposure
HSBC Holdings PLC	Financials	3,462,675,500
Barclays PLC	Financials	2,862,294,893
Citigroup Inc.	Financials	2,409,106,024
Goldman Sachs Group Inc.	Financials	1,476,020,570
BCE Inc.	Telecommunications	1,436,582,746
Chevron Corp.	Energy	1,412,888,499
Royal Bank of Scotland Group	Financials	1,197,653,154
JP Morgan Chase Co.	Financials	1,126,490,957
Exxon Mobil Corp.	Energy	856,848,012
Exelon Corp.	Utilities	835,715,965

Table 4 shows the Top 10 List by Expected Loss. While there is some overlap, the differences between the Top 10 List by Exposure and the Top 10 List by Expected Loss are considerable.

TABLE 4
Top 10 List by Expected Loss

Name	Sector	Expected Loss
Ensign Energy Services Inc.	Energy	7,723,727
Companhia Energetica de Minas Gerais (CEMIG)	Utilities	7,344,251
HSBC Holdings PLC	Financials	6,942,475
Petroleo Brasileiro S.A. Ord	Energy	6,249,301
Husky Energy Inc.	Energy	4,751,980
Shell Canada Ltd.	Energy	4,327,390
Citigroup Inc.	Financials	4,293,962
Canadian Natural Resources Ltd.	Energy	4,262,384
Barclays PLC	Financials	3,566,929
Xstrata PLC	Energy	3,509,872

Table 5 shows the Top 10 List by CVaR (99.9%), illustrating the importance of correlation in quantifying risk. Table 5 also illustrates the benefit of considering more factors in the ultimate measure of credit risk: many names not apparent in the previous lists suddenly appear in this assessment. The choice of 99.9 percent matches the Pillar 1 standards of Basel II for capital calculation.

TABLE 5
Top 10 List by CVaR (99.9%)

Name	Sector	CVaR (99.9%)
HSBC Holdings PLC	Financials	436,616,264
Barclays PLC	Financials	167,355,654
Citigroup Inc.	Financials	129,752,971
Petroleo Brasileiro S.A. Ord	Energy	121,368,636
Husky Energy Inc.	Energy	88,305,663
Shell Canada Ltd.	Energy	83,434,673
Canadian Natural Resources Ltd.	Energy	81,056,541
Talisman Energy Inc.	Energy	71,528,055
Ensign Energy Services Inc.	Energy	69,574,640
Suncor Energy Inc.	Energy	53,734,493

Many financial institutions maintain a higher rating than the BBB average of Basel II, which drove the choice of the 99.9 percent confidence level. Accordingly, they seek a higher confidence interval for CVaR, for example an AA rating would indicate a 99.98 percent level. In the case study, the choice between 99.9 percent and 99.98 percent has little effect on the Top 10 List, with a slight re-ordering amongst the top 12 names in the overall list. This may not be the case for all portfolios.

Relative measures are also possible. For example, the Top 10 List by Relative CVaR(99.9%) would list the most capital-expensive names per dollar of exposure. It can be created using the ratio of CVaR (99.9%) to exposure**.** Table 6 contains the resulting Top 10 List.

TABLE 6
Top 10 List by Relative CVaR (99.9%)

Name	**Sector**	**Relative CVaR (99.9%)**
Transocean Inc.	Energy	18.59%
Petroleo Brasileiro S.A. Ord	Energy	17.72%
Husky Energy Inc.	Energy	16.85%
Ensign Energy Services Inc.	Energy	16.74%
Shell Canada Ltd.	Energy	16.55%
Canadian Natural Resources Ltd.	Energy	16.31%
Suncor Energy Inc.	Energy	15.34%
Imperial Oil Ltd.	Energy	15.21%
HSBC Holdings PLC	Financials	12.61%
Adobe Systems Inc.	Info Tech.	12.36%

The Top 10 List by Relative CVaR(99.9%) resembles the Top 10 List by Expected Loss with one noticeable difference: As expected, many of the highly rated names with large exposure that appear in the Top 10 List by Expected Loss do not appear in the Top 10 List by Relative CVaR(99.9%). They are replaced by lower-rated names that are more likely to default, and thus require more capital per dollar of exposure to protect against potential losses. This gives a better indication of which names are more expensive with which to do new business.

Comparing the four Top 10 lists we see four distinct groups of firms. From the

Top 10 List by Exposure, the first three entries appear in most of the lists, while the remaining seven entries are unique to this list. Those names from the Top 10 List by Expected Loss generally continue to appear in lists created using more risk-sensitive metrics. Finally, the fourth group of names includes five firms present in one of the last two lists, but in neither of the first two. Without advanced, risk-sensitive CVaR measures, the risks associated with Talisman Energy, Suncor Energy, Transocean, Imperial Oil and Adobe Systems would remain hidden within the portfolio.

In some cases, comparing the relative shares of capital (or CVaR (99.9%)) against the relative shares of exposure along various dimensions such as rating, sector or country can yield portfolio management insights. Figures 1, 2 and 3 present CVaR (99.9%) decompositions by sector, rating and country alongside the corresponding exposure decompositions. While the key sectors and countries are clearly the same, it is noteworthy that the CVaR (99.9%) breakdowns highlight concentrations in the utilities sector and Brazil that remain hidden in the exposure-only analysis.

We now turn to the question of which specific names add concentration and diversification to our portfolio. We gain such insight by looking at the concentration index (CI) of each name in our portfolio. Recall that the CI is the standardized name concentration ratio relative to the portfolio. A counterparty with a CI >1 adds concentration risk to the portfolio while those with a CI < 1 add diversification to the portfolio.

According to the CI, we have 48 names that add concentration to the portfolio (using CVaR 99.9% to compute CI) leaving 452 to create the diversification. The CI can be plotted against exposure to identify which names are good candidates for investment. Figure 1 shows a scatter plot of the CI against exposure, by name.

Extended Case Study: Attributing Risks

In the base case study discussed above, all credit risk factors are modeled and the overall CVaR (99.9%) is estimated to be U.S. $2.58B. To attribute a portion of this risk to migration risk, we "turn off" migration risk and recalculate under the assumption of only two states: default and no-default. The resulting CVaR (99.9%) level is U.S. $1.94B. Thus, we can attribute the difference, i.e., U.S. $637M, to migration risk. TABLE 8 summarizes this calculation alongside the other first-order effects.

Our base case is a multifactor model which assumes that credit drivers are associated to names based on their country and industry. By changing to a single-factor model, similar to that used in Pillar 1 calculations under Basel II, we see a CVaR (99.9%) level of U.S. $3.53B. Thus, the multifactor environment (expressed partially through sector diversification within the portfolio) is providing U.S. $951M of diversification benefit.

If the increase in CVaR (99.9%) seems unintuitive initially, it can be explained by examining the correlation assumptions. Because the multiple factors are correlated, but not perfectly correlated, having only a single credit driver for all counterparties increases the average pair wise asset and default correlations significantly. Table 7 shows the average of the pair-wise asset and default correlations in the single factor and the multifactor models.

TABLE 7
Correlations under Various Model Assumptions

Model Assumption	Asset Correlation	Default Correlation
Single Factor	31.2%	2.1%
Multifactor	18.7%	0.9%

For the base case, we model idiosyncratic risk in the context of a multifactor model using Monte Carlo sampling techniques. To isolate the name risk we change this assumption to one of full diversification. As a result, CVaR (99.9%) drops to U.S. $2.13B. Clearly, the portfolio is not large enough to completely diversify away the idiosyncratic risk. Attributing the U.S. $456M difference to name concentration risk in the portfolio provides significant insight: no matter what diversification strategy we assume, we are unlikely to reduce capital requirements, as measured by CVaR (99.9%), beyond the U.S. $2.13B level.

For each of the three factors, or risk types, we have measured their isolated impact on the portfolio-level risk measure, CVaR (99.9%). Such first-order attributions allow ranking of the risks in order of importance: sector diversification, migration risk and name diversification. These are referred to as first-order effects because in each case, only one assumption was changed. TABLE 8 shows a summary of the results of the first-order attributions.

TABLE 8
First-Order Effects

Model	CVaR (99.9%)	Interpretation	Attribution (Base - Model)
Base Case	2,581,738,825		
Default / No Default (DND)	1,944,680,630	Migration Risk	637,058,195
No Name Concentration	2,125,727,123	Name Concentration	456,011,702
Single Factor	3,533,178,245	Sector Diversification	951,439,420
Total 1st Order Effects			141,630,477

Unfortunately, the CVaR (99.9%) result when all three assumptions are changed simultaneously differs significantly from the cumulative U.S. $141M shown above. It is U.S. $1.95B, creating the need to explain a total change of U.S. $626M. The reason for the discrepancy is that the model is not straightforward—or linear. Recall that it also includes interdependencies, correlations and interactions between these risks. Hence, higher-order effects create the U.S. $485M discrepancy. This calculation is summarized in Table 9 while Table 10 provides an overview of the second-order effects by assessing CVaR (99.9%) results using pairs of assumptions.

TABLE 9
Higher-Order Effects

Model	CVaR (99.9%)
Base Case	2,581,738,825
Single Factor, No Name Concentration, DND	1,954,955,399
Difference	626,783,426
Total 1st Order Effects	141,630,477
Total to be Explained by Higher Order Effects	485,152,949

Table 10. Second-Order Effects

Model	CVaR (99.9%)	Attribution	2nd Order Effect
Base Case	2,581,738,825		
No Name Concentration, DND	1,340,623,229	1,241,115,596	148,045,699
Single Factor, DND	2,404,043,868	177,694,957	492,076,182
Single Factor, No Name Concentration	3,241,306,962	-659,568,137	-164,140,419
Total 2nd order effects			475,921,462

From the results in Table 10, we see that the second-order effects account for almost all of the initial discrepancy. The largest of the second-order effects arises from the interaction between the multifactor models and migration risk. This might arise if migration risk has a regional or sector-specific component, indicating the need for further investigation before attempting to hedge either type of risk.

In contrast, the largest combined effect is from the removal of name concentration risk and migration risk. This second result is expected, since the single-factor model is based on an overall higher level of correlation. When we remove the idiosyncratic risk from the portfolio and no longer allow migration between credit states, the capital requirement decreases substantially. However, the decrease is only slightly more than the sum of the decreases from applying each of these assumptions individually. This implies that there is little relationship between name concentrations and migration risk, indicating independent hedging strategies are likely to be as effective as a coordinated effort.

Conclusions

We have presented various techniques for measuring, assessing and presenting concentration risk, observing that the use of any single measure or representation can be misleading when analyzing concentration. Accordingly, we have measured name level concentration through a variety of techniques, including the CI, which differentiates names adding concentration from those adding diversification to the portfolio. The CI can therefore help to identify good candidates for new business.

We have also studied various types of risk at different levels of granularity on our portfolio through the process of capital attribution. We have seen that it is not only important to assess the impact of migration risk, name risk and sector risk, but also to look at the interactions between these types of risk in creating a complete picture of concentrations.

Concentration risk is likely to remain an issue that requires a significant amount of time and effort to manage. However, regulators and other stakeholders are demanding more accurate and precise answers that can only be obtained by using more sophisticated models and providing more detailed analysis.

Acknowledgments

Many thanks to Patrick Tobin, financial engineering student at Algorithmics, for his work on the case study and preliminary (internal) white papers (August 2008) upon which this work relies heavily.

About Algorithmics

Algorithmics is the world's leading provider of enterprise risk solutions. Financial organizations from around the world use Algorithmics' software, analytics and advisory services to help them make risk-aware business decisions, maximize shareholder value, and meet regulatory requirements. Supported by a global team of risk experts based in all major financial centers, Algorithmics offers proven, award-winning solutions for market, credit and operational risk, as well as collateral and capital management. Algorithmics is a member of the Fitch Group.

Diane Reynolds is a senior director at Algorithmics Inc. She can be reached at *Diane.Reynolds@Algorithmics.com*.