

INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Case—Quantifying Operational Risk in Financial Institutions

Brian Keller, Güzin Bayraksan,

To cite this article:

Brian Keller, Güzin Bayraksan, (2012) Case—Quantifying Operational Risk in Financial Institutions. INFORMS Transactions on Education 12(2):106-113. <http://dx.doi.org/10.1287/ited.1110.0075cs>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

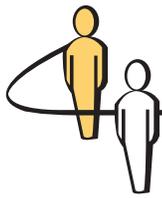
Copyright © 2012, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



Case

Quantifying Operational Risk in Financial Institutions

Brian Keller

Booz Allen Hamilton, Annapolis Junction, Maryland 20701,
keller_brian@bah.com

Güzin Bayraksan

Systems and Industrial Engineering, University of Arizona, Tucson, Arizona 85721,
guzinb@sie.arizona.edu

Risk management is essential in today's business environment for banks and other financial institutions to survive in highly competitive and volatile market environments. As the nexus of financial institutions grows ever more complex, sound risk management practices, or the lack thereof, affect more than just the individual institution. The subprime mortgage debacle in 2008 has clearly demonstrated this fact.

Key words: risk modeling; teaching statistics; maximum likelihood estimation; teaching simulation

History: Received: February 2009; accepted: February 2011.

Background

On June 26th, 1974, the troubled German bank Herstatt was ordered into liquidation by the German authorities. Earlier that day, several German banks had paid deutsche marks to Herstatt, believing they would receive U.S. dollars later in the same day. However, because of time zone differences, the German banks never received their dollars. It was only early morning in New York when Herstatt was liquidated and all outgoing U.S. dollar payments were suspended.

This messy incident brought about the interconnectedness of the banking system around the world. The national banking legislations were too limited in their geography to be able to handle such incidents. In response, the eleven G-10 nations (Belgium, Canada, France, Germany, Italy, Japan, The Netherlands, Sweden, Switzerland, the United Kingdom, and the United States) and Luxembourg decided to form a council to improve the quality of banking standards and supervision within the member states. Thus, a committee was formed in 1974, informally known as the *Basel Committee*, consisting of each country's central banker and representatives of her banking supervisory authorities. The committee meets four times a year in Basel, Switzerland.

Soon after its inception, the Basel Committee began efforts to harmonize the international banking system. In July of 1988, after six years of deliberation, the initial twelve nations plus Spain released a capital accord, commonly referred to as *Basel I* (Basel Committee 1998). Basel I focused on *credit risk*, that is, the risk that the borrower does not pay the agreed upon amount (principal and/or interest). Basel I proposed minimum capital requirements to be set aside for internationally active banks to protect against credit risk.

One of the criticisms of Basel I was its narrow scope. Credit risk, although being very important, is not the only risk banks face. Another major risk is the *market risk*, which is the risk that the portfolio of the bank's market investments (stocks, bonds, treasuries, etc.) will decrease because of moves in market factors such as changes in the interest rates. Basel I was expanded in January 1996 with the *Market Risk Amendment*.

This was just the beginning of an ever evolving supervisory risk mitigation in the international banking system.

Operational Risk

In the words of Roger Ferguson—the vice chairman of the board of governors of the Federal Reserve System

from 2001 to 2006—at a hearing before the Committee on Banking, Housing, and Urban Affairs of the United States Senate in 2003:

In an increasingly technologically driven banking system, operational risks have become an even larger share of total risk. Frankly, at some banks, they are probably the dominant risk.

But what is *operational risk*? Why and when did it become so important?

The currently accepted definition of operational risk is as follows (Basel Committee 2001)

Operational risk is the risk of loss resulting from inadequate or failed internal processes, people or systems or from external events.

This is quite a general definition and covers a multitude of risks such as employee fraud, credit card fraud, worker's compensation lawsuits, damage to physical assets due to weather, and data entry errors. It includes *legal risk* but excludes *strategic risk*. Legal risk is the risk that the bank will face lawsuits, adverse judgments, unenforceable contracts, and penalties and sanctions pronounced by a regulatory body. Strategic risk, on the other hand, is the risk associated with the bank's future business plans and strategies.

Operational risk can also be the cause of another type of risk referred to as the *reputation risk*. For instance, after a bank loses money because of internal fraud, this results in bad publicity, which could then result in negative consequences like loss of existing customers or fewer new customers. According to a study by Perry and de Fontnouvelle (2005) for internal fraud events, the actual total loss, taking into account losses due to reputational consequences, on average, is three times larger than the operational losses.

Most operational losses happen frequently and do not result in major damages. These include small data entry mistakes and minor fraud cases. However, banks (and other financial institutions) can suffer from operational risks that can cause major losses in excess of billions of dollars and even cause them to go bankrupt. Therefore, it is paramount for banks to protect themselves from losses due to operational risks. Below are some examples of operational risk that show the range and magnitude of this risk.

- *Rogue Trading*: A *rogue trader* is a professional trader who makes unapproved financial transactions to the detriment of his/her clients and institution. Some of the most famous rogue traders are Nick Leeson (Barings Bank 1995, losses of \$1.3 billion), John Rusnak (Allied Irish Banks 2002, losses of \$691 million), Toshihide Iguchi (Daiwa Bank Group 1995, losses more than \$1 billion), and Jérôme Kervie

(Société Générale 2008, losses of \$7.2 billion). Nick Leeson wrote a book *Rogue Trader*, which was published in 1997. It was later turned into a movie starring Ewan McGregor, released in 1999. Leeson had been making speculative unauthorized transactions in Asian markets. At some point he had over \$3 billion in index futures on the Nikkei. In 1995, Leeson incurred \$1.3 billion in losses after the Kobe earthquake in Japan resulted in broad-based sell-off. As a result, the 233-year-old Barings Bank went into bankruptcy. Leeson was charged with fraud and sentenced to six and a half years in jail. This is an example of "inadequate or failed internal processes" that should have detected such activity earlier and "people" (fraud) in the definition of operational risk.

- *Terrorist Attacks of September 11, 2001*: On the morning of September 11, 2001, the world financial system was shaken when two of the four hijacked planes crashed into the World Trade Center in New York City. Direct losses to the financial system included damages to physical property, disruptions in financial services around the world, and the loss of employees. For instance, the global financial firm Cantor Fitzgerald, which was headquartered on the top floors of the north tower of the World Trade Center, lost 638 employees. This is an example of a loss from "external events" in the definition of operational risk.

- *Cyber Attacks to Financial Institutions*: In 2005, MasterCard announced a security breach to its cardholder data. At least 68,000 MasterCard account holders' information was accessed. In September 2007, TD Ameritrade, the online trading company, announced that a security breach occurred with its client information database. The database contained personal information such as social security numbers of more than six million customers. These are examples of failed "systems."

As the above examples illustrate, operational risks have indeed become a larger share of the total risk, as Ferguson noted. This is a consequence of the changes that have been occurring in the financial system over the last couple of decades. These changes include: (i) increased globalization and the interdependencies in the world's financial systems, (ii) deregulation such as the Financial Services Act of 1999 in the United States that allows affiliations among different financial services, (iii) the increase of computer-based banking services that bring about a new set of operational risks as illustrated above in different cyber attacks, and (iv) an increase in the variety of types of financial instruments being offered. As a result, the financial industry has been exposed to new kinds of risk that it is still learning how to deal with effectively. Many of these risks fall under the operational risk category. It is fair to say that the topic is still evolving.

Basel II

In response to the changing global banking environment and the enlarged risks banks are facing, the Basel Committee decided to expand Basel I. In 1998, the document *Operational Risk Management* was released. After many changes, the currently accepted form of *Basel II* was finalized in June 2006 (Basel Committee 2006). *Basel II* greatly expanded *Basel I*—*Basel I* is approximately 30 pages and *Basel II* is more than 300 pages. One of the most important changes is the inclusion of the operational risk in addition to credit and market risks.

The *Basel II Capital Accord* uses a three pillar system to guide banks in risk management. The first pillar defines minimum capital requirements banks must hold in order to protect themselves from the risks they face. The second pillar focuses on regulatory policies. For instance, banks should conduct internal reviews involving the board and senior management to ensure adequate risk management, and regulators should oversee the internal risk evaluations. The third pillar complements the first two pillars and focuses on market discipline through public disclosure mechanisms.

Operational Risk in Basel II

The first pillar of *Basel II* proposes three ways for banks to calculate capital to cover losses due to operational risks. The first and easiest of these methods is the *basic indicator approach*, which recommends banks to hold capital equivalent to 15% of their average gross income in the past three years of positive income. The second method, known as the *standardized approach*, provides divides a bank by its business lines and percentages of gross income to hold as reserves within each business line. For instance, *Basel II* recommends setting aside 18% of the gross income for high risk business lines such as corporate finance and sales and trading but only 12% of gross income for low risk business lines such as asset management. Medium level risk items such as commercial banking need a 15% allocation. Then, these are combined to find the minimal capital requirements.

The third approach provides the most flexibility to the banks. It is called the *advanced measurement approach* (AMA). In AMA, banks are allowed to use their own method for operational risk, but these must be approved by the regulators. The AMA models are more complex than the basic indicator or standardized approaches but they typically yield better estimates of risk. This is because they are based directly on the operational loss data rather than the gross income, which is used as a proxy for the first two approaches. Another major advantage of AMA is that it can result in a smaller (but still risk-appropriate) amount of capital to be set aside for operational risk.

The 2008 loss data collection exercise conducted by the Basel Committee shows that AMA results in a capital reserve of approximately 10% of gross income compared to 15% in the basic indicator approach and the 12–18% in the standardized approach (Basel Committee 2009). This frees up capital for banks to use for other profitable business. In addition, *Basel II* allows banks that use AMA to qualify for deductions if they have insurance for operational risk.

Basel II incentivizes the use of AMA to increase self-surveillance of the banks. This way, banks become aware of the risks they face and take action against it. This in turn decreases the cost of regulation and potential bankruptcies due to poor risk management. As a result, banks can allocate more money to other businesses and provide more capital to an economy.

Operational Risk Data

Data! Data! Data! he cried impatiently. I can't make bricks without clay.

(Sherlock Holmes, *The Adventure of the Copper Beeches*)

Although banks do want to use AMA because of the aforementioned advantages, one of the biggest challenges they face in its implementation is the *availability of data* for operational losses. Credit risk and market risk models have been developed over the years and extensive data for these models are readily available. For example, stock prices and interest rates are updated almost instantaneously. However, operational risk is a newer concept and thus, data for it has not been collected over the years.

Losses that result from operational risk events can be either high frequency and low impact (e.g., minor data entry errors), or low frequency and high impact (e.g., terrorist attacks, rogue trading). Banks can easily recover from high frequency and low impact losses but they must protect themselves from the high impact losses. The *low frequency* of these events poses a major difficulty in the statistical analysis of operational risk. First, there are not enough data points to allow for reliable statistical estimation. A common approach adapted by banks is to supplement the data with plausible scenarios for these tail events. In addition, because rare events that happen in the tail of a distribution are of interest, statistical analysis requires a large sample size.

Another problem with operational loss data is that it is recorded only above a certain threshold. That is, small values of losses are not typically recorded. First, some of these small losses are easy to hide and may not be detected. Second, given the high frequency of these events, data recording of *all* events might be quite costly. Third, because these are usually low impact events, financial institutions tend to ignore them. However, small losses should be appropriately addressed in a thorough statistical analysis.

Table 1 Databases for operational loss data

Database	By	Information
ORX	Operational Riskdata eXchange Association, a not-for-profit association based in Zurich, Switzerland	Founded in 2002. Has 53 members. Contains 170,081 losses of €58 billion as of 31 March 2010. www.orx.org
Algo OpData	Algorithmics, part of Fitch Group	Contains 12,000 publicly reported losses in excess of \$1 million. www.algorithmics.com
GOLD (Global Operational Loss Database)	British Bankers' Association	Initiated in 2000. Data from UK, Europe and Australia. www.bba.org
DIPO (Database Italiano Perdite Operative)	DIPO Association, previously under the Italian Banking Association	(An example of a national database.) Initiated in 2003. In 2008, had 34 members and 200 entries. www.dipo-operational-risk.it

Given the scarcity of loss data for reliable operational risk modeling, the Basel Committee has conducted three international data collection exercises since 2001 called the *loss data collection exercise* (LDCE). The last LDCE was conducted in 2008 with the participation of 121 banks across 17 countries (Basel Committee 2009). The 2004 exercise in the United States and the 2007 exercise in Japan were LDCEs at the national level. The results of these studies show differences across geographical regions and countries. For instance, data indicate that the Japanese banks have lower frequency of losses compared to banks in the United States.

As operational risk gained more importance, starting in the early 2000s, a number of international and national operational loss databases have been initiated. Table 1 provides examples and information about some of these databases. These databases contain invaluable information for business practitioners.

A Bank

Banks in the G-10 countries were to implement Basel II by the end of 2008. This is where A Bank stands right now.

Remarks: Note that A Bank is facing this issue at around 2006–2007, when our case study takes place. At the time of writing, new updates to the Basel Accords are underway. These updates are referred to as Basel III and include revised definitions of capital, introduction of a leverage ratio, countercyclical capital buffers, mitigations to counterparty credit risk,

Table 2 The largest banks in the United States as of March 31, 2008

Rank	Name (city, state)	Assets (in millions)
1	Citigroup (New York, NY)	2,199,848
2	Bank of America Corp. (Charlotte, NC)	1,743,478
3	J. P. Morgan Chase & Company (New York, NY)	1,642,862
4	Wachovia Corp. (Charlotte, NC)	808,575
5	Taunus Corp. (New York, NY)	750,323
6	Wells Fargo & Company (San Francisco, CA)	595,221
7	HSBC North America Inc. (Melrose, IL)	493,010
8	U.S. Bancorp (Minneapolis, MN)	241,781
9	Bank of the New York Mellon Corp. (New York, NY)	205,151
10	Suntrust, Inc. (Atlanta, GA)	178,986
11	Citizens Financial Group, Inc. (Providence, RI)	161,759
12	National City Bank (Cleveland, OH)	155,046
13	State Street Corp. (Boston, MA)	154,478
14	Capital One Financial Corp. (McLean, VA)	150,608
15	Regions Financial Corp. (Birmingham, AL)	144,251
16	PNC Financial Services Group, Inc. (Pittsburg, PA)	140,026
17	BB&T Corp. (Winston-Salem, NC)	136,417
18	TD Bank North, Inc. (Portland, MA)	118,171
19	Fifth Third Bankcorp (Cincinnati, OH)	111,396
20	Keycorp (Cleveland, OH)	101,596

Notes. The assets are listed in millions of dollars. Data from: Federal Reserve System, National Information Center.

and quantitative liquidity ratios. Current Basel III proposed regulations do not significantly change operational risk regulations as defined in Basel II. Financial institutions are tentatively planning Basel III implementation in early 2013.

A Bank is one of the nation's largest bank-based financial services companies with assets valued more than \$100 billion. (See Table 2 for a list of the 20 largest banks in terms of their assets in the United States as of May 30, 2008. A Bank is one these banks but the name of the bank is hidden and the example data was artificially generated for confidentiality reasons.) A Bank provides investment management, retail and commercial banking, consumer finance, and investment banking products and services to individuals and companies throughout the United States and, for certain businesses, internationally.

Like all financial institutions in the United States, A Bank is subject to periodic reviews by federal regulators. These regulatory reviews require, among many things, that A Bank sets aside cash reserves to offset the potential risk of loss that it faces every day. A Bank wants to hold the correct amount of capital (cash reserves) required to cover its risk. If this capital is too small, A bank could suffer a large loss and go bankrupt. If this capital is too large, the bank will lose out on earning money on the excess capital, hence, its profitability will go down.

A Bank has credit and market risk models that have been used for many years and are well calibrated. These models allow A Bank to allocate an appropriate amount of capital to cover credit and market risks.

However, operational risk is a relatively new concept. A Bank realizes that with sensitive risk models it can lower its operational risk capital and increase its earnings. However, A Bank does not want to go bankrupt. To find the appropriate level of capital, they want to build risk models sensitive to operational risk.

Loss Distribution Approach

After some initial studies, A Bank has decided to use the AMA approach. Within this approach, they have narrowed down their focus to a method that has been in use in the actuarial risk calculations, known as the *loss distribution approach* (LDA). In LDA, operational risk is composed of two basic components:

- (i) frequency of loss events, and
- (ii) severity of each loss.

Frequency of loss events refers to the number of losses that occur within a given time period. Typically (and because of Basel II) this is taken as one year. The severity of each loss then refers to the amount of money lost because of the event that occurred.

The bank is ultimately interested in the *aggregate loss*, that is, the total amount of money that the bank may lose because of operational risk in any given year, which is found by combining the two basic components above. Let X be the random variable corresponding to aggregate loss, Y be the random variable corresponding to the number of loss events (frequency) in a given time period (e.g., a year), and Z be the random variable corresponding to the amount of money lost on each loss event (severity). Then, the aggregate loss is given by

$$X = \sum_{i=1}^Y Z_i.$$

To find the capital to hold, A Bank wants to cover 99.5% of the losses it may face. Then, there is only a 0.5% chance that total losses in a given year may cause A Bank to go bankrupt. Let $F_X(x)$ be the cumulative distribution function of the aggregate loss; $F_X(x) = P(X \leq x)$. Then, A Bank is interested in finding the quantity

$$\text{VaR}_{0.995}(X) = \min\{x: F_X(x) \geq 0.995\}.$$

This resulting quantity is known as the *value at risk* (VaR) with confidence level 0.995 and gives the capital reserves for operational risk. If the bank has capital for 99.5% of losses, then it expects that its capital will only be exceeded one out of every 200 years.

Data and Risk Modeling

A Bank has collected data on how many losses have occurred in the last 15 years and the amount of money lost on each loss event. The data set has been expanded with plausible scenarios for very large loss

Table 3 Number of Losses in the Past 15 Years (Frequency Data)

Year	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
No. of losses	8	7	14	16	7	23	21	12	6	7	11	5	9	12	6

events and also adjusted for small losses that might have been ignored. The resulting data are given in Tables 3 and 4, and plots of the frequency data and severity data are provided in Figure 1. As can be seen in the plots, the severity data has a large number of small losses but a small number of very large losses. That is, it has a high peak at a small value and is heavily skewed to the right. Distributions of this type of data are called *leptokurtic*.

The bank plans to model the loss frequency with a discrete probability distribution and model the loss severity with a continuous probability distribution. By combining these two distributions, the bank can come up with an estimate of the capital required to cover 99.5% of losses that it may face. Some initial studies suggest that the frequency distribution is best modeled by either the Poisson or the negative binomial distributions. The Poisson distribution is often used to model the number of events in a given time. In this case, it can be used to model the number of losses within a year. The negative binomial is sometimes referred to as the “over dispersed Poisson” because its variance is greater than its mean. Recall that Poisson has its mean equal to its variance. Initial studies also indicate that the severity distribution is best modeled by either the lognormal, or by the generalized Pareto distribution (GPD), or a combination of the two. The severity of the losses are *heavily* skewed to the right, which makes modeling difficult. The lognormal models the body of the losses well but does not capture the extreme losses. The GPD does not model the body well, but it does an excellent job of modeling the extreme losses. Sometimes the lognormal is spliced together with the GPD to model both the body and tail of the loss distribution. Appendices A and B provide information about these probability distributions.

Unfortunately, the distribution of X (the aggregate loss) cannot be written in closed form; instead, it is usually simulated to make inferences. *Monte Carlo simulation* can be easily done by first generating a realization of the frequency random variable to simulate the number of losses in a year. Then, that many losses are generated from the severity distribution and are summed to get *one* observation from the aggregate loss distribution. After enough observations of X are generated, the observations create a “picture” of what the true distribution looks like. Figure 2 shows the histogram of 100,000 simulated yearly total losses from the aggregate loss distribution. As can be seen

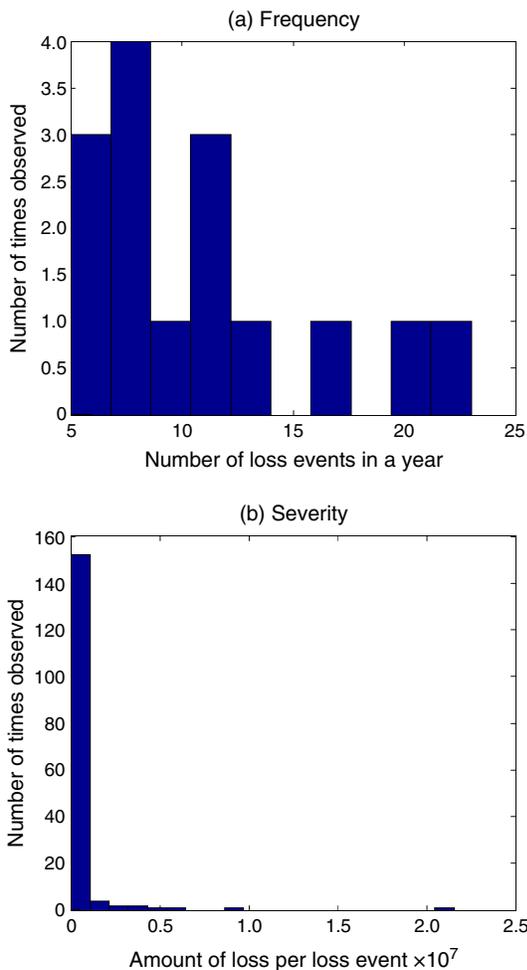
Table 4 Severity of Losses (Severity Data)

2,666,000	686,000	21,544,000	14,000	7,000	16,000	1,331,000	39,000	3,000
1,000	500	16,000	3,000	91,000	440,000	6,000	1,000	1,000
100	16,000	29,000	50,000	370,000	31,000	92,000	201,000	3,378,000
16,000	311,000	551,000	47,000	39,000	232,000	232,000	57,000	130,000
1,000	125,000	718,000	105,000	85,000	1,000	33,000	4,000	900
195,000	122,000	933,000	10,000	92,000	2,000	5,000	11,000	11,000
1,000	11,000	23,000	300	181,000	379,000	47,000	1,000	374,000
5,000	347,000	7,000	22,000	69,000	2,000	173,000	5,515,000	45,000
124,000	28,000	22,000	230,000	3,000	3,000	21,000	581,000	2,000
200	14,000	130,000	6,000	1,000	5,000	81,000	271,000	1,000
1,741,000	6,000	24,000	147,000	618,000	1,000	499,000	133,000	3,000
4,000	8,000	32,000	800	5,000	1,000	2,000	100	242,000
30,000	9,000	8,000	21,000	3,000	131,000	35,000	37,000	926,000
1,000	59,000	26,000	98,000	3,000	7,000	6,000	1,009,000	262,000
17,000	39,000	1,000	20,000	247,000	1,088,000	1,000	16,000	3,000
12,000	533,000	38,000	600	88,000	13,000	4,782,000	9,069,000	4,145,000
300	69,000	296,000	5,000	9,000	195,000	1,000	24,000	3,043,000
8,000	39,000	270,000	18,000	7,000	137,000	117,000	438,000	1,983,000
34,000	13,000							

from Figure 2, the distribution of the aggregate loss is also leptokurtic.

The bank needs to come up with an estimate of the capital required to cover 99.5% of losses that it may

Figure 1 Histograms of (a) Frequency of Loss Events Each Year, and (b) Severity of Loss from Each Loss Event

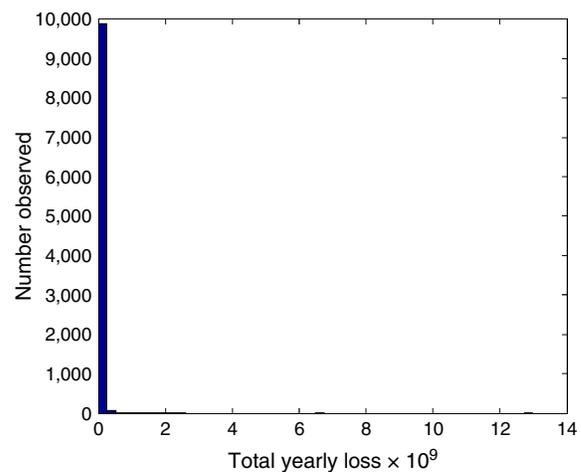


face. This can be accomplished by taking the 99.5% quantile by sorting the list of simulated observations and taking the $n(0.995)$ th observation, where n is the total number of simulated values. For example, if you have 1,000 observations, you take the 995th sorted observation (i.e., the 995th order statistic).

When banks do not have sensitive risk models, regulators force banks to hold larger amounts of capital than actually required. Therefore, A Bank needs to show the regulators that its model is adequate and valid. A Bank is fairly confident that its data is representative but the regulators need to see if the approach and model of A Bank is statistically valid.

Given the difficulties in data collection with respect to operational losses, issues with respect to fitting a distribution to severity data, and the difficulties in determining the distribution of aggregate loss, A bank is considering some nonparametric statistical methods to validate its models. Appendices C and D provide

Figure 2 Histogram of 100,000 Simulated Yearly Total Losses from the Aggregate Loss Distribution



background information on selected parametric and nonparametric statistical methodology.

Acknowledgments

This case won the 2008 INFORMS Case Competition.

Appendix A. Selected Frequency Distributions

• *Poisson.* The Poisson distribution is used to model the number of things happening within a certain interval. Here, A Bank is interested in modeling the number of loss events per year. Let $\lambda > 0$. The probability mass function is given by

$$p(x | \lambda) = P(X = x) = \frac{\lambda^x e^{-\lambda}}{x!}, \quad x = 0, 1, 2, \dots$$

The mean and variance of a Poisson random variable are equal; $E[X] = \text{var}(X) = \lambda$.

The Poisson distribution has the following property: if X is distributed Poisson with mean λ_1 and Y is distributed Poisson with mean λ_2 , and X and Y are independent from one another, then, $X + Y$ has a Poisson distribution with mean $\lambda_1 + \lambda_2$. This can be useful when modeling the number of loss events from different lines of business of a bank or a financial institution, assuming the lines of business are independent from one another.

• *Negative binomial.* There is another form of this distribution; however, the bank requires that you use the one given below. Let $r \geq 1$ be an integer and $0 < p < 1$. The probability mass function is given by

$$p(x | r, p) = P(X = x) = \binom{r+x-1}{x} p^r (1-p)^x, \quad x = 0, 1, 2, \dots$$

The mean of a negative binomial random variable is $E[X] = r(1-p)/p$ and variance is $\text{var}(X) = r(1-p)/p^2$. Hence, its variance is larger than the mean.

The relationship between the Poisson and negative binomial distributions is as follows: let X be such that for a given value of λ , it is a Poisson random variable with mean λ . However, λ is not constant and is assumed to be a random variable that has a Gamma distribution with parameters r and β , where $\beta = p/(1-p)$. Then, X is a negative binomial random variable.

The negative binomial also has the property that if two independent random variables, X and Y , are distributed as negative binomials with same p but with respective parameters r_1 and r_2 , their sum, $X + Y$, has a negative binomial distribution with parameters $r_1 + r_2$ and p .

Appendix B. Selected Severity Distributions

• *Lognormal.* If Y is a normally distributed random variable with mean μ and variance σ^2 , then, $X = e^Y$ has a lognormal distribution with the probability density function

$$f(x | \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma x}} e^{-1/2(\ln x - \mu)^2/\sigma^2}, \quad x > 0.$$

The lognormal distribution has a large spike at a small value close to zero and has moderately heavy tails.

The lognormal distribution can be used to model quantities that are products of other independent quantities. (This is due to the central limit theorem.) For instance, if certain loss types can be viewed as a product of many independent losses, then the lognormal distribution may be appropriate to model such losses.

• *Generalized Pareto distribution (GPD).* Let $\xi > 0$ and $\sigma > 0$. Then, the probability density function of GPD is given by

$$f(x | \xi, \sigma) = \frac{1}{\sigma} \left(1 + \frac{\xi x}{\sigma}\right)^{-1-1/\xi}, \quad x > 0.$$

(Note that we assumed the shape parameter ξ to be positive; $\xi > 0$. In this case, the GPD has a heavy tail. However, there are more general forms of the GPD where ξ can take other values. For instance, when $\xi = 0$, the GPD is equivalent to an exponential distribution, when $\xi = -1$ the GPD is the uniform distribution.)

The GDP is important in *extreme value theory*, which deals with rare events that lie in the tail of a probability distribution. Consider Y , a random variable, and define the *excess distribution function*, $F_\tau(y)$, as the conditional distribution of Y over a certain threshold τ :

$$F_\tau(y) = P(Y - \tau \leq y | Y > \tau).$$

The *Pickands-Balkema-de Haan theorem* states that, for sufficiently large values of τ , under certain regularity conditions, the excess distribution function $F_\tau(y)$ is approximately distributed as GPD (in its more general form than we presented).

Appendix C. Maximum Likelihood Estimation

Maximum likelihood is a widely used method of estimating the unknown parameters of a probability distribution. It has several important properties that make it useful for statistical inference. For instance, maximum likelihood estimators are *asymptotically normal*, which is often used to make inferences. They yield consistent estimators under mild regularity conditions and these estimators have the *invariance property*, that is, if $\hat{\theta}_{MLE}$ is the maximum likelihood estimator of θ , then, given a function h , $h(\hat{\theta}_{MLE})$ is the maximum likelihood estimator of $h(\theta)$. We briefly describe this method here.

Let $f(x | \theta)$ be the pdf or pmf of a random variable X with parameter θ . Although θ could represent a vector of parameters, for ease of exposition, let us assume there is only one unknown parameter. Given an independent and identically distributed sample of observations, $\mathbf{x} = (x_1, x_2, \dots, x_n)$, the likelihood function is defined as

$$\mathcal{L}(\theta | \mathbf{x}) = \prod_{i=1}^n f(x_i | \theta).$$

The likelihood function gives the likelihood of a parameter θ for the observed sample \mathbf{x} . For instance, if θ_1 yields a higher value than θ_2 , then, θ_1 has a higher likelihood of being the unknown parameter given the observations \mathbf{x} . So, the idea of maximum likelihood estimation (MLE) is to find the value of θ that maximizes the likelihood function. The resulting estimator is

$$\hat{\theta}_{MLE} = \arg \max_{\theta \in \Theta} \mathcal{L}(\theta | \mathbf{x}).$$

Downloaded from informs.org by [47.88.87.18] on 02 April 2017, at 15:20. For personal use only, all rights reserved.

Additional information, including supplemental material and rights and permission policies, is available at <http://ite.pubs.informs.org>.

This is often done by first taking the logarithm of the likelihood function, which turns the product into a sum, and then maximizing the loglikelihood function over the set of possible parameter values, Θ . Because the logarithm is a monotonically increasing function, maximizing the loglikelihood function is equivalent to maximizing the likelihood function. When this function is differentiable, the value of θ that yields a first derivative equal to zero is typically the MLE. For some distributions, an analytical solution is not available and some numerical optimization technique must be used.

Example. Let X_1, X_2, \dots, X_n be a random sample from a normal distribution with unknown mean μ and known variance σ^2 . (So, here, $\theta = \mu$ is the unknown parameter.) The likelihood function is

$$\begin{aligned} \mathcal{L}(\mu | \mathbf{x}) &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} e^{-(x_i - \mu)^2/2} \\ &= \frac{1}{(2\pi)^{n/2}} e^{-\sum_{i=1}^n (x_i - \mu)^2/2}. \end{aligned}$$

Now, maximizing this function with respect to θ is equivalent to maximizing the power of the exponent (because the exponent is a monotonically increasing function). Therefore

$$\frac{\partial -\sum_{i=1}^n (x_i - \mu)^2/2}{\partial \mu} = 0$$

yields $\sum_{i=1}^n (x_i - \mu) = 0$. So, the maximum likelihood estimator of μ is

$$\hat{\mu}_{MLE} = \frac{\sum_{i=1}^n x_i}{n},$$

which is the well-known sample mean.

Appendix D. Goodness of Fit

In this section, we summarize some commonly used goodness-of-fit tests. In these tests, the alternatives are as follows:

- H_0 = the data follow the specified distribution, and
- H_1 = the data does *not* follow the specified distribution.

- *Chi-Squared test.* In the chi-squared test, the distribution specified in the null hypothesis as well as the data are divided into cells. Then, the number of observed values in each cell is compared to the expected number of observations under the distribution specified in the null hypothesis. If these values are close to one another, then we fail to reject the null hypothesis, and the specified distribution is a “good fit” for the data. Let k be the number of cells, O_i be the number of observations, and E_i be the expected number of observations under the null distribution in cell i , $i = 1, 2, \dots, k$. The test statistic is

$$\chi^2 = \sum_{i=1}^k \frac{(O_i - E_i)^2}{E_i}.$$

Larger values of χ^2 indicate that the null distribution may not be a good fit for the data. When the expected number of observations are sufficiently large for each cell, χ^2 is approximately distributed as a chi-squared random variable with $k - c$ degrees of freedom, where c is the number of estimated parameters.

The chi-squared test can be applied to both continuous and discrete distributions. The test is sensitive to the

number of cells used and the size of each cell. A common rule of thumb is to select the cells such that each cell has at least five observations in it.

- *Kolmogorov-Smirnov (KS) test.* The KS test is a nonparametric test that compares the empirical cumulative distribution function (cdf) with the cdf of the null distribution. Let $F_n(x)$ be the empirical cdf and let $F(x)$ be the null distribution. Then, the KS test looks at the largest distance between $F_n(x)$ and $F(x)$. The test statistic is

$$D = \sup_x |F_n(x) - F(x)|.$$

Large values of D indicate a poor fit. The KS test is only applicable to continuous distributions but its test statistic is distribution free. When the parameters of the null distribution are estimated, the critical values need to be simulated. Another important characteristic of the KS test is that it is more sensitive to the deviations in the center of the distribution rather than its tails.

- *Anderson-Darling (AD) test.* The AD test is similar to the KS test in the sense that it also uses the empirical cdf. However, unlike the KS test, which puts more weight in the center of the distribution, the AD test puts more weight on the tails of the distribution. This can be important for the severity distributions. Although there is a supremum type AD test, A Bank prefers the more commonly used quadratic type AD test, for which the test statistic is

$$A^2 = n \int_{-\infty}^{+\infty} \frac{(F_n(x) - F(x))^2}{F(x)(1 - F(x))} dF(x).$$

Different critical values for some distributions (e.g., normal, lognormal, Weibull) have been tabulated. When one of these distributions is tested, Anderson-Darling is a powerful goodness-of-fit test. Critical values for other distributions can be simulated.

Above, we presented the formal definitions for the KS and AD test statistics. Formulas to calculate these statistics based on the empirical cdf can be found in many statistics books, see also §§1.3.5.16 and 1.3.5.14 of the *Engineering Statistics e-Handbook (NIST/SEMATECH 2011)*.

References

Basel Committee on Banking Supervision. International convergence of capital measurement and capital standard. Bank of International Settlements, Basel, Switzerland. <http://www.bis.org/publ/bcbsc111.htm>. Last accessed September 25, 2011.

Basel Committee on Banking Supervision. Working paper on the regulatory treatment of operational risk. Bank of International Settlements, Basel, Switzerland. http://www.bis.org/publ/bcbs_wp8.pdf. Last accessed September 25, 2011.

Basel Committee on Banking Supervision. Basel II: Revised international capital framework. Bank of International Settlements, Basel, Switzerland. <http://www.bis.org/publ/bcbsca.htm>. Last accessed September 25, 2011.

Basel Committee on Banking Supervision. Results from the 2008 loss data collection exercise for operational risk. Bank of International Settlements, Basel, Switzerland. <http://www.bis.org/publ/bcbs160a.pdf>. Last accessed September 25, 2011.

NIST/SEMATECH E-Handbook of Statistical Methods. 2011. <http://www.itl.nist.gov/div898/handbook/>. Last accessed September 25, 2011.

Perry, J., P. de Fontnouvelle. 2005. Measuring reputational risk: The market reaction to operational loss announcements. Technical report, Federal Reserve Bank of Boston, Boston.