PART

One

Operational Risk Measurement: Qualitative Approaches

COPYRIGHTED MATERIAL

CHAPTER 1

Modeling Operational Risk Based on Multiple Experts' Opinions

Jean-Philippe Peters and Georges Hübner

ABSTRACT

While the Basel II accord has now gone live in most parts of the world, many discrepancies still remain on advanced modeling techniques for operational risk among large international banks. The two major families of models include the loss distribution approaches (LDAs) that focus on observed past internal and external loss events and the scenario-based techniques that use subjective opinions from experts as the starting point to determine the regulatory capital charge to cover operational risk. A major methodological challenge is the combination of both techniques so as to fulfill Basel II requirements. In this chapter we discuss and investigate the use of various alternatives to model expert opinion in a sound statistical way so as to allow for subsequent integration with loss distributions fitted on internal and/or external data. A numerical example supports the analysis and shows that solutions exist to merge information arising from both sources.

Georges Hübner gratefully acknowledges financial support from Deloitte Luxembourg.

1.1 INTRODUCTION

The revised Framework on Capital Measurement and Capital Standards for the banking sector, commonly referred to as Basel II, has now gone live in most parts of the world. Among the major changes introduced in this new regulatory framework are specific capital requirements to cover operational risk, defined by the Accord as the "risk of loss resulting from inadequate or failed internal processes, people and systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk" (BCBS 2005).

Operational risk management is not new to financial institutions: stability of information technology (IT) systems, client claims, acts of fraud, or internal controls failures have been closely monitored for years. However, these elements have historically been treated separately. Basel II combines all items into one single integrated measurement and management framework.

Three methods are proposed by Basel II to measure the capital charge required to cover operational risk. The two simplest ones—the *basic indicator approach* and the *standardized approach*—define the operational risk capital of a bank as a fraction of its gross income; the *advanced measurement approach* (AMA) allows banks to develop their own model for assessing the regulatory capital that covers their yearly operational risk exposure within a confidence interval of 99.9% (henceforth this exposure is called *operational value at risk,* or OpVaR).

To comply with regulatory requirements, a sound AMA framework combines four sources of information:

1. Internal operational risk loss data
2. Relevant external operational risk loss data
3. Scenario analysis of expert opinion
4. Bank-specific business environment and internal control factors

The relative weight of each source and the way to combine them together are up to the banks; Basel II does not provide a regulatory model.

This chapter mainly relates to the third element—operational risk quantification using experts' opinion—and how it can successfully be addressed so as to produce outcome that can be combined with other elements (i.e., internal and external loss data).

Most of the literature on operational risk modeling focuses either on methods based on actual loss data, the so-called loss distribution approach (Chapelle et al. 2008; Chavez-Demoulin, Embrechts, and Neslehova 2006; Cruz 2002; Frachot, Georges, and Roncalli 2001), or on application of

Bayesian techniques to combine loss-based models and scenario analysis (Alexander 2003; Figini et al. 2007; Lambrigger, Shevchenko, and Wüthrich 2007). To our knowledge, studies focusing on scenario analysis on a stand-alone basis are more scarce, except for Alderweireld, Garcia, and Leonard (2006); Andres and van der Brink (2004); the sbAMA Working Group (2003); and Steinhoff and Baule (2006).

On the contrary, the problems related to obtaining probability measures by experts and combination of opinions from several experts have been extensively studied by statisticians over the last decades. Good reviews of these topics can be found in Garthwaite, Kadane, and O'Hagan (2005) for statistical methods used; Daneshkhah (2004) for the psychological aspects; Clemen and Winkler (2007) and Genest and Zidek (1986) for mathematical aggregation of opinions; and Plous (1993) for behavioral aggregation of opinions.

Starting from a practical business case, our objective is to bridge scenario analysis in operational risk with the vast literature on experts' opinion modeling by studying how operational and methodological challenges can be addressed in a sound statistical way.

In Section 1.2, we present the major families of AMA models and how they can be combined. Narrowing the scope of our analysis, Section 1.3 describes a business case concerning challenges banks face in practice when they desire to rely on experts' opinion. Existing solutions and their limitations are reviewed in Section 1.4. In Section 1.5 we introduce an alternative method based on the so-called supra-Bayesian approach and provide an illustration. Finally, Section 1.6 opens the door for further research and concludes.

1.2 OVERVIEW OF THE AMA MODELS

The Basel Committee released its consultative paper on operational risk in 2001. Since then, banks have started designing and developing their own internal model to measure their exposure to operational risk just as they began implementing consistent and comprehensive operational risk loss data collection processes.

The application of AMA is in principle open to any proprietary model, but methodologies have converged over the years and standards have appeared. The result is that most AMA models can now be classified into two categories:

1. Loss distribution approaches
2. Scenario-based approaches[1]

1.2.1 Loss Distribution Approach

The loss distribution approach (LDA) is a parametric technique primarily based on past observed internal loss data (potentially enriched with external data). Based on concepts used in actuarial models, the LDA consists of separately estimating a frequency distribution for the occurrence of operational losses and a severity distribution for the economic impact of the individual losses. Both distributions are then combined through n-convolution of the severity distribution with itself, where n is a random variable that follows the frequency distribution.[2]

When applied to operational risk, LDA is subject to many methodological issues that have been extensively treated in the literature over the past few years.[3]

All these issues (and others) are specifically challenging when developing LDA models for operational risk because of a major root cause: the availability of loss data points, specifically large ones. Because conducting financial activities implies interaction with many other parties (clients, banking counterparties, other departments, etc.), operational errors are often spotted relatively fast by ex-post controls, such as monitoring reports, reconciliation procedures, or broker confirmations. As such, most of the operational errors to which financial institutions are exposed have acceptable financial impact (i.e., high-frequency/low-severity events). From time to time, however, the impact can be huge, as when fraudulent purposes are combined with a weak control environment and large underlying operations (e.g., cases in Barings, Daiwa, Allied Irish Bank or, more recently, Société Générale).

Having enough relevant data points to implement a sound LDA model can prove to be difficult in many cases, including:

- Small or midsize banks with limited and/or automated activities
- New business environment due to merger/acquisition
- Specific operational risk classes that are not related to "everyday business"

In such cases, banks often rely on experts' opinions in what is sometimes referred to as scenario-based AMA.

1.2.2 Scenario-Based AMA

The scenario-based AMA (or sbAMA) shares with LDA the idea of combining two dimensions (frequency and severity) to calculate the aggregate loss

distribution (ALD) used to derive the OpVaR. It differs, however, from LDA in that estimation of both distributions builds on experts' opinion regarding various scenarios.[4]

Based on their activities and their control environment, banks build scenarios describing potential adverse events actualizing the operational risks identified as relevant. Then experts are asked to give opinions on probability of occurrence (i.e., frequency) and potential economic impact should the events occur (i.e., severity).

As human judgment of probabilistic measures is often biased, a point discussed in more detail in Section 1.3, a major challenge with sbAMA is to obtain sufficiently reliable estimates from experts. One might wonder why experts should be required to give opinions on frequency and severity distributions (or their parameters). The official reason lies in the regulatory requirement to combine four sources of information when developing an AMA model. But the usefulness of this type of independent, forward-looking information should not be overlooked, provided it is carefully integrated in an actuarial, backward-looking setup. The challenge of modeling adequately experts' opinions is thus unavoidable when quantifying operational risk and is discussed in Section 1.4.

Outcome from sbAMA shall be statistically compatible with that arising from LDA so as to enable a statistically tractable combination technique. As we show in next section, the most adequate technique to combine LDA and sbAMA is Bayesian inference, which requires experts to set the parameters of the loss distribution.

1.2.3 Integrating LDA and sbAMA

When it comes to combining sbAMA and LDA,[5] banks have adopted integration methods that greatly diverge and that can be split in two groups: *ex post* combination and *ex ante* combination:

1. *Ex post* combination consists in merging the various sources at the aggregate loss distribution (ALD) level. Typically, separate ALDs are computed independently, and they are combined on the final distribution. As stated in Section 1.2, the n-fold convolution of the severity distribution with itself is very difficult to derived analytically so numerical solutions, such as Monte Carlo simulations, are used in practice. As a consequence, ALD is usually discretized and combinations are applied to each of the points defining the distribution (often 10,000 to 50,000 data points).

2. *Ex ante* combination focuses on the frequency and the severity distributions, prior to the simulations leading to the ALD. More specifically, sources of information are combined to derive the parameters of both distributions.

Whichever combination approach used,[6] a widely accepted solution in the academic literature is to merge information from the various components of the AMA using Bayesian inference.

Consider a random vector of observations $X = (X_1, X_2, \ldots, X_n)$ whose density, for a given vector of parameters $\theta = (\theta_1, \theta_2, \ldots, \theta_K)$, is $\ell(X \mid \theta)$ and called the *sample likelihood*. Suppose also that θ itself has a probability distribution $\pi(\theta)$, called the *prior distribution*. Then Bayes' theorem (Bayes 1763) can be formulated as

$$\ell(X, \theta) = \ell(X \mid \theta)\pi(\theta) = \hat{\pi}(\theta \mid X)\ell(X) \tag{1.1}$$

where $\hat{\pi}(\theta \mid X) =$ density of parameters θ given observed data X and is called the *posterior* distribution
$\ell(X, \theta) =$ joint density of observed data and parameters
$\ell(X) =$ a marginal density of X

Equation 1.1 can also be expressed as

$$\hat{\pi}(\theta \mid X) = \frac{\ell(X \mid \theta)\pi(\theta)}{\ell(X)} \tag{1.2}$$

where $\ell(X) =$ a normalization constant

Thus, the posterior distribution can be viewed as a product of a prior knowledge with a likelihood function of observed data. Mathematically,

$$\hat{\pi}(\theta \mid X) \propto \pi(\theta) \times \ell(X \mid \theta) \tag{1.3}$$

In the context of operational risk we consider:

- The posterior distribution $\hat{\pi}(\theta \mid X)$ provides the input to define the distributions used in the Monte Carlo simulations to derive the OpVaR.
- The prior distribution $\pi(\theta)$ should be estimated by scenario analysis (i.e., based on expert opinions).
- The sample likelihood $\ell(X \mid \theta)$ is obtained when applying LDA on actual losses (whether internal losses only or a mixture of internal and external data).

Bayesian inference has many great features, which is why it fits well with operational risk modeling:

1. It provides a structural and sound statistical technique to combine two heterogeneous sources of information (subjective human opinions and objective collected data).
2. It provides transparency for review by internal audit and/or regulators as both sources of information can be analyzed separately.
3. Its foundations rest on assumptions that fit well with operational risk, as both observations and parameters of the distributions are considered to be random.

1.3 USING EXPERTS' OPINIONS TO MODEL OPERATIONAL RISK: A PRACTICAL BUSINESS CASE

Imagine this business case, commonly observed in the banking industry nowadays: A bank wishes to use AMA to measure its exposure to operational risk (be it for the regulatory requirements of Pillar I or within its internal capital adequacy assessment process [ICAAP], under Pillar II). It has set up a loss data collection process for several years and has enough loss events to develop an LDA model (at least for its major business lines).

To ensure compliance of its model with regulatory requirements, the bank decides to associate its LDA model with scenario-based experts' opinion, adopting the *ex ante* combination technique with Bayesian inference, as described in Section 1.2.3. The operational risk manager (ORM) faces implementation and methodological challenges.

First, the ORM should decide who should serve as "experts" to assess the various scenarios. The subjective nature of this exercise might cause the ORM to consult multiple experts in an attempt to beef up the information base and reduce the uncertainty surrounding the assessment. (Subjective information is often viewed as being "softer" than "hard scientific data" such as actual loss events.) In most banks, scenario analysis is carried out by department or business lines. The ORM would thus ask the head of each business line to provide a list of staff members having sufficient skills, experience, and expertise to take part in the scenario analysis exercise.

Internal risk governance implies that the head of a business line has the ultimate responsibility for managing operational risk in the daily activities of the unit and for ensuring that a sound and efficient control environment is in place. As a consequence, this person will probably be requested to review, validate, or approve results of the estimation process. This leads to

the second challenge: How should the ORM organize the scenario analysis work sessions? Often workshops are held, with all experts including the head of business line attending. Participants discuss and debate to get a final common estimate. Such a solution has advantages: It is less time consuming than individual sessions, and participants can share information to enhance accuracy of these estimates. But workshops also present some risks: A group decision-making process implies psychological factors, such as, for instance, the emergence of leaders in the discussion who strongly influence the others, who in turn may give in for the sake of group unity even though they do not agree. This is especially true when various hierarchical levels are represented (e.g., the head of business line). Another psychological factor can be linked to game theory concepts, with some experts intentionally providing estimates that are not their real beliefs in an attempt to influence others and drive the consensus toward a particular value (Fellner 1965). To reduce these risks, we assume in our business case that the ORM decides to gather the opinions from each individual expert and from the head of business line separately.

Before opinions can be combined, they must be elicited and expressed in some quantified form, which is the root of the third challenge. Indeed, as stated by Hogarth (1975), "man is a selective, sequential information processing system with limited capacity . . . ill-suited for assessing probability distributions." Unfortunately, many of the solutions to combine opinions require that each individual's opinion to be encoded as a subjective probability distribution. Consequently, the way scenarios and related questions are built is vital, as many potential biases may happen during the extraction process.[7] These biases are reflected in the main heuristics used by experts when assessing probabilities under uncertainty:

- **Availability.** People overestimate the frequency of events similar to situations they have experienced and underestimate probabilities of less familiar events (Tversky and Kahneman 1974).
- **Anchoring.** When people are asked to provide a range for the impact of a potential uncertain, severe event or assess its frequency, they use a starting point, often the perceived expected value, and adjust upward or downward. Unfortunately, this adjustment is often not sufficient and produces systematic underestimation of the variability of the estimates (O'Hagan 1998; Winkler 1967).
- **Representativeness.** When people have to assess the probability of an event, they tend to link this event with another similar event and derive their estimate from the probability of the similar event. One of the errors produced by this heuristic is called the law of small numbers, as people

typically ignore the variability of the mathematical laws of probability in small samples (Kahnemann, Slovic, and Tversky 1982).

- **Framing.** Outcomes from questionnaires (i.e., probability estimates) are sensitive to the phrasing and the order of questions used (Pahlman and Riabacke 2005).

To avoid direct probabilistic judgments by the experts, the ORM should prepare questions that fit better with the way nonstatisticians think. A solution is to give experts a list of potential loss impacts for each scenario and to ask: How many years will we have to wait, all things being equal, to observe such a scenario happening with loss impact of x or above? Such a question focuses on the notion of duration (i.e., the mean time one should expect to wait until the occurrence of a given operational risk event exceeding a certain severity), which is easier for experts to handle in practice (Steinhoff and Baule 2006).

Assuming the frequency distribution is a Poisson distribution with parameter λ, the duration until an event of magnitude x occurs can be expressed as

$$d(x) = \frac{1}{\lambda\left(1 - F\left(x; \Theta\right)\right)} \tag{1.4}$$

where $F(x; \Theta) =$ any parametric severity distribution

Hence, based on duration estimates provided by experts for each of the potential financial impacts, it is easy to construct the related probability measures.

Then comes the final challenge: What method should the ORM use to combine these individual opinions? The quantification method shall meet these conditions:

Condition #1. Able to combine several opinions on a parameter of interest θ[8]

Condition #2. Accounting for the assessment of the head of each business line who has the ultimate responsibility of the process

Condition #3. Building on sound and robust statistical methods to get supervisory approval

Condition #4. Producing an outcome that can be plugged in the subsequent (Bayesian) combination with LDA

The next section analyzes potential solutions proposed in the literature.

1.4 COMBINING EXPERTS' OPINIONS

In many other areas where sample data are rare or nonexisting (e.g., the nuclear industry, seismology, etc.), more than one expert is requested to provide an opinion. In other words, a decision maker (DM) relies on opinions collected from a pool of experts to assess a variable of interest θ. In operational risk, θ could be the probability P that the severity of a loss event is equal to the amount A.[9]

The problem to be considered is the combination of a number of probabilities representing their respective judgments into a single probability to be used as an input (the *prior* distribution) for the Bayesian combination with LDA, as described in Section 1.2.3. This problem is called the "consensus problem" by Winkler (1968). In this section, we provide a brief overview of existing solutions used to solve the consensus problem that could fit in an operational risk context.

Note that pure group decision-making situations (i.e., workshops where one single agreed estimate is obtained) would not fit our model as they rely on aggregation approaches that are more "behavioral" in the sense that they "attempt to generate agreement among the experts by having them interact in some way" (Clemen and Winkler 2007).

1.4.1 Weighted Opinion Pool

The problem can be tackled by relying on some weighting technique to pool all experts' opinions. Mathematically, if we consider k experts and a linear weighting scheme (called *linear opinion pool*), we have

$$P = \sum_{i=1}^{k} \lambda_i P_i \tag{1.5}$$

where $P =$ probability of interest
$P_i =$ opinion of the ith expert

Of course, we have $\sum_{i=1}^{k} \lambda_i = 1$.

The most obvious solution is simply to take the equally weighted average of the opinions—that is, $\lambda_i = 1/k$—for all i. If we want to use different weights, a ranking between experts should be established based on some metrics measuring the "goodness" or "reliability" of each expert. This ranking could be set up by the line manager based on his or her level of confidence in each expert. The major drawback of this alternative is that the choice of the weights is subjective and hardly justifiable.

While meeting conditions #1 and #4 introduced in the previous section, both the equal-weight approach and the manager-based ranking approach might not fit condition #3. A more convincing method is to base the ranking on measured accuracy of previous assessments made by the experts. This can be done using the classical model of Cooke (1991) by testing the level of expertise of each expert using a questionnaire that includes questions directly related to θ (target questions) but also questions to which the assessor (in our case, the ORM) already knows the answers (seed questions). The seed questions are then used to assess the degree of expertise of each expert through a scoring rule, which combines the notions of calibration and information to estimate the weight of each expert. The scoring rule requires setting up a specific parameter called the cutoff level. As the final weighted estimates can be scored just as those from each expert, this cutoff level is chosen so that this weighted estimate's score is maximized. In the field of operational risk, this method has been applied in Bakker (2004). The major difficulty with this model is defining the scoring rule.

Alternative nonlinear methods (such as the logarithmic opinion pool) have also been proposed, but, as mentioned by Genest and McConway (1990):

> *it had been hoped initially that a careful analysis . . . would point to the "better" solution, but this approach is gradually being abandoned as researchers and practitioners discover that equally compelling (but irreconcilable) sets of axioms can be found in support of different pooling recipes.*

While more advanced linear solutions, such as the classical model of Cooke (1991) or the logarithmic opinion pool, meet condition #3, they still fail to provide a satisfying answer to condition #2. In the next section, we introduce an approach that fills the gap.

1.4.2 Supra-Bayesian Approach

Relying on ideas detailed in Section 1.2.3, we could assume that each opinion is the data input to a single decision maker (the head of business line in our case) who updates his or her prior views with these opinions. Such an approach would meet condition #3. Originally introduced in Winkler (1968), this method was called "supra Bayesian" by Keeney and Raiffa (1976) and has since then been studied by, among others, Morris (1974), Lindley (1983), French (1985), Genest and Schervish (1985), and Gelfand et al. (1995).

In the supra-Bayesian approach, the pooling process is: The decision maker (called the "supra Bayesian") defines his or her prior probability P for the occurrence of E, and the collected opinions from the k experts $P_1, P_2, \ldots, P_k$ constitute the sample likelihood. Using these opinions, the supra-Bayesian beliefs can be updated using Bayes' formula:

$$P(E \mid P_1, P_2, \ldots, P_k) \propto P(E) \times \ell(P_1, P_2, \ldots, P_k \mid E) \tag{1.6}$$

Here the posterior distribution $P(E \mid P_1, P_2, \ldots, P_k)$ (or P^* in the remainder of this chapter) can be seen as the consensus. It can then be used as *prior* distribution in the subsequent phase of the modeling processes (i.e., the combination with LDA).

The supra-Bayesian approach thus nicely fits with all four conditions listed in Section 1.3. Furthermore, it does not require much additional information technology tools or specific statistical skills from the modeler, as it relies on the same concept and calculation process as those presented in Section 1.2.3. But a major difficulty remains as Equation 1.6 implies that the likelihood function is known (i.e., the probability distribution of experts' opinion given that E occurs). The supra Bayesian (or whoever has to assess the likelihood function) shall evaluate the experts, their respective prior information sets, the interdependence of these sets, and so on. In summary, this approach requires the pooling of a very substantial number of expert opinions.

Fortunately, solutions exist to address this problem, by assuming more realistically that the likelihood function cannot be specified fully. In particular, Genest and Schervish (1985) provide a model for such situation that could be applied in an operational risk context.[10] This model is explicated in the next section.

1.5 SUPRA-BAYESIAN MODEL FOR OPERATIONAL RISK MODELING

1.5.1 Model

Suppose that θ is the severity distribution we want to model, which is a continuous univariate distribution. In scenario analysis applications for operational risk, it is often assumed that θ is a member of a standard parametric family, to which the expert need only supply its parameters. The most widely used distribution in such cases is the lognormal distribution for which experts provide the mean and standard deviation. Other solutions include the

triangular distribution and the PERT distribution for which experts provide minimum, most likely (mode), and maximum values.

This assumption seems too restrictive and likely inappropriate. Severity of operational risk events indeed presents a highly leptokurtic nature, which is inadequately captured by these distributions, even the lognormal one. Many studies suggest using mixture distributions or the concepts of extreme value theory (i.e., the generalized Pareto distribution).[11] Experts are unable to provide the full specifications of such complex distributions, so we assume that each expert expresses his/her opinion on θ in a partial way. That is, probabilities are provided for a small collection of disjoint exhaustive intervals in the domain of θ. In practice, we use the duration approach of Steinhoff and Baule (2006) described in Section 1.3; experts are requested to provide duration for various financial impacts, and probabilities for the disjoint ranges are built based on these duration estimates. In this case, Equation 1.4 becomes

$$d(x) = \frac{1}{\lambda\left(F(y;\Theta) - F(x;\Theta)\right)} \tag{1.7}$$

In statistical terms, we assume the domain of θ, Θ, is an interval in R that has lower and upper bounds, namely $a_0 = \inf\{\theta \in \Theta\}$ and $a_k = \sup\{\theta \in \Theta\}$.[12] We further assume that Θ is partitioned in n intervals determined by the points $a_0 < a_1 < \ldots < a_n$ and let $I_j = (a_{j-1}, a_j)$. Let $p_i^{\mathrm{T}} = (p_{i1}, p_{i2}, \ldots, p_{in})$, where p_{ij} is the opinion of the ith expert regarding the chance that $\theta \in I_j$. If we have k experts participating, this results in a $n \times k$ matrix $P = (p_l, p_2, \ldots, p_k)$. Like the experts, the supra Bayesian is not in a position to provide the fully specified distribution for θ. Thus he or she will answer the same questions as the experts. Mathematically, the supra Bayesian provides a vector of probabilities ρ for the same sets of I's. Equation 1.6 then becomes

$$P^* \propto \rho \times \ell\left(P \mid \theta \in I_j\right) \tag{1.8}$$

To model the likelihood function ℓ, we rely on the method proposed by Genest and Schervish (1985), who provide a Bayesian answer requiring a minimum of a priori assessments. For each interval I_j, the supra Bayesian is required only to provide the first moment (i.e., the expected value) of the *marginal* distribution F_j for $\pi_j = (p_{1j}, p_{2j}, \ldots, p_{kj})$, the vector of experts' probabilities for the occurrence $\theta \in I_j$. This mean vector is denoted $M_j = (\mu_{1j}, \mu_{2j}, \ldots, \mu_{kj})$. Note that no specification on the *conditional* distribution is thus required.

In our operational risk business case, this is equivalent to the head of the business line having to answer this question: What duration do you think

your experts will most probably associate with the given loss amount? The head of the business line should be able to identify the more conservative people on his or her team (or the opposite).

Genest and Schervish (1985) test pooling processes against a consistency condition that guarantees "it is consistent with the unspecified F, in the sense that the joint distribution for $\theta \in I_j$ and π_j that is compatible with F can always be found such that the pooling process is the posterior probability P^*."

When only M_j is specified, Genest and Schervish (1985) show that the only formula to pass this test is

$$P_j^* = \rho_j + \sum_{i=1}^{k} \lambda_{ij} \left(p_{ij} - \mu_{ij}\right) \tag{1.9}$$

with possibly negative weights λ_{ij} that expresses the coefficients of multiple linear regressions (or the amount of correlation) between each p_{ij} and $\theta \in I_j$. These weights must satisfy $n^* 2^{k+1}$ inequalities when no prior restrictions on the sign of the λ_{ij}'s exist. In the most common case (i.e., λ_{ij}'s are positive), the inequalities take the form

$$\max\left\{\sum_{i=1}^{k} \frac{\lambda_{ij}\mu_{ij}}{\rho_j}; \sum_{i=1}^{k} \frac{\lambda_{ij}\left(1-\mu_{ij}\right)}{\left(1-\rho_j\right)}\right\} \leq 1, \quad \forall j = 1, \ldots, n \tag{1.10}$$

1.5.2 Illustration of the Model

To illustrate this model, consider the simple case where a head of business line H and three experts A, B, and C are required to provide frequency of occurrence of an event-type E and the duration for two severity amounts €1,000,000 and €5,000,000. We further assume that the total available owned funds of the bank is €100,000,000 and that the upper bound for severity is fixed at this level. Finally, H is requested to provide expectations in terms of the experts' answers. Knowing the staff sufficiently well, H anticipates that A would provide a more severe picture of the event but remains neutral for B and C. Table 1.1 summarizes the answers received.

Using Equation 1.7 to transform these answers into probability measures leads to the figures reported in Table 1.2 and Table 1.3. In Table 1.2, the opinion of H is equivalent to ρ in Equation 1.8.

TABLE 1.1 Experts' Opinions

Question A: What is the frequency of occurrence for event E?

Frequency (F)	***H***	***A***	***B***	***C***
(# events/year)	3	5	3	3.5

Question B: how many years will we have to wait, all things being equal, to observe the event E of severity higher than X?

Severity (in EUR)	**Duration (in years)**			
S	***H***	***A***	***B***	***C***
1,000,000 (= S_1)	5	3	6	6.5
5,000,000 (= S_2)	20	15	22	20

Question C (only for H): What answers do you think the experts will provide for Questions A and B?

Variable	***H***	***A***	***B***	***C***
F	—	3.5	3	3
S_1	—	4	5	5
S_2	—	18	20	20

TABLE 1.2 Probability Measures Associated with Experts' Opinions

Interval (I)		**$P = \Pr(\theta \in I \mid E)$**			
Lower Bound	**Upper Bound**	***H***	***A***	***B***	***C***
0	1,000,000	0.9333	0.9333	0.9444	0.9560
1,000,000	5,000,000	0.0500	0.0533	0.0404	0.0297
5,000,000	100,000,000	0.0167	0.0133	0.0152	0.0143

TABLE 1.3 Beliefs of the Supra-Bayesian

Interval (I)		**$P = \Pr(\theta \in I \mid E)$**			
Lower Bound	**Upper Bound**	***H***	***A***	***B***	***C***
0	1,000,000	—	0.9286	0.9333	0.9333
1,000,000	5,000,000	—	0.0556	0.0500	0.0500
5,000,000	100,000,000	—	0.0159	0:0167	0.0167

TABLE 1.4 Prior and Posterior Probability Measures

Interval (I)	ρ	P^*
I_1	0.9333	0.9499
I_2	0.0500	0.0354
I_3	0.0167	0:0147

Inequalities of Equation 1.10 are satisfied with $\lambda_1 = \lambda_2 = \lambda_3 = 0.3$. When plugging these coefficients in Equation 1.9, we obtain the posterior probability measures $P*$ reported in Table 1.4.

1.6 CONCLUSION

Operational risk quantification using scenario analysis is a challenging task, both methodologically and organizationally. But its informative value can hardly be ignored in any sound operational risk measurement framework; in addition, regulatory requirement exist regarding the use of experts' opinions in the AMA approach.

This chapter has presented the practical business case of a banking institution wishing to adopt a scenario analysis to model operational risk based on opinions from several experts who are subject to validation or review by a single hierarchical superior. While most of the existing solutions fail to meet all constraints faced by such a situation, the supra-Bayesian model presented adequately copes with the major challenges identified in our business case by providing a sound, robust, yet tractable framework to model the consensus of experts' opinions.

NOTES

1. Some alternatives exist, but they are scarcely used in the industry. As observed by the Financial Supervisory Authority (FSA 2005) in the United Kingdom, "the institutions represented on the AMA Quantitative Expert Group currently all use either Loss Distribution or Scenario approaches."
2. See, for instance, Cruz 2002 or Frachot, Georges, and Roncalli 2001 for theoretical background; see Aue and Kalkbrener 2006 or Chapelle et al. 2008 for practical illustrations.
3. See, for instance, Crama, Hübner, and Peters 2007 on impact of the loss data collection threshold; Moscadelli 2004 on tail modeling; King 2001 on the usage

of mixed distribution to model severity; Di Clemente and Romano 2004 on dependence among operational risks.

4. See sbAMA Working Group (2003) for an introduction to the general concepts of this method.
5. Similar techniques are also applied to combine internal loss data with relevant external loss data. The *ex post* combination sometimes is applied only over a given (high) threshold.
6. In the remainder of this chapter, we consider only the *ex ante* combination.
7. See Daneshkhah (2004) for a review of the literature on the subject.
8. In our case, θ are the answers to the questions linked to a given scenario.
9. Note that the probability measure P could be substituted by a probability distribution p throughout this discussion without loss of accuracy.
10. Other alternatives have been proposed by Genest and Schervish (1985); West (1988); West and Crosse (1992); and Gelfand, Mallick, and Dey (1995).
11. See, for instance, Chavez-Demoulin, Embrechts, and Neslehova (2006) or Chapelle et al. (2008).
12. In operational risk, a_0 could be 0 and a_k could be equal to a loss so important it leads the bank to bankruptcy (e.g., a loss equal to its total own funds).

REFERENCES

Alderweireld, T., J. Garcia, and L. Leonard. 2006. A practical operational risk scenario analysis quantification. *Risk* 19, no. 2:93–95.

Alexander, C. 2003. *Operational risk: Regulation, analysis and management*. London: Prentice Hall-FT.

Andres, U., and G. J. Van Der Brink. 2004. Implementing a Basel II scenario-based AMA for operational risk. In *The Basel handbook*, ed. K. Ong. London: Risk Books.

Aue, F., and M. Kalkbrener. 2006. LDA at work: Deutsche Bank's approach to quantifying operational risk. *Journal of Operational Risk* 1, no. 4:49–93.

Bakker, M. R. A. 2004. Quantifying operational risk within banks according to Basel II. Master's thesis. Delft Institute of Applied Mathematics, Delft, Netherlands.

Bayes, T. 1783. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society* 53:370–418.

Basel Committee on Banking Supervision. 2005. Basel II: International convergence of capital measurement and capital standards—A revised framework. Basel Committee Publications No. 107, Bank for International Settlements, Basel, Switzerland.

Chapelle, A., Y. Crama, G. Hübner, and J.-P. Peters. 2008. Practical methods for measuring and managing operational risk in the financial sector: A clinical study. *Journal of Banking and Finance* 32, no. 6:1049–1061.

Chavez-Demoulin, V., P. Embrechts, and J. Neslehova. 2006. Quantitative models for operational risk: Extremes, dependence and aggregation. *Journal of Banking and Finance* 30, no. 10:2635–2658.

Clemen, R. T., and R. L. Winkler. 2007. Aggregating probability distributions. In *Advances in decision analysis: From foundations to applications*, ed. R.F. Miles and D. von Winterfeldt. New York: Cambridge University Press.

Cooke, R. M. 1991. *Experts in uncertainty*. New York: Oxford University Press.

Crama, Y., G. Hübner, and J.-P. Peters. 2007. Impact of the collection threshold on the determination of the capital charge for operational risk. In *Advances in Risk Management*, ed. G. Gregoriou. London: Palgrave-MacMillan.

Cruz, M. G. 2002. *Modeling, Measuring and hedging operational risk*. Hoboken, NJ: John Wiley & Sons.

Daneshkhah, A. R. 2004. Psychological aspects influencing elicitation of subjective probability. Working paper, University of Sheffield, U.K.

Di Clemente, A., and C. Romano. 2004. A copula–extreme value theory approach for modelling operational risk, In *Operational risk modelling and analysis: Theory and practice*, ed. M. Cruz. London: Risk Books.

Fellner, W. 1965. *Probability and profits*. Homewood, IL: Irwin.

Figini, S., P. Guidici, P. Uberti, and A. Sanyal. 2007. A statistical method to optimize the combination of internal and external data in operational risk measurement. *Journal of Operational Risk* 2, no. 4:87–99.

Frachot, A., P. Georges, and T. Roncalli. 2001. Loss distribution approach for operational risk. Working paper, Groupe de Recherche Opérationnelle, Crédit Lyonnais, Paris.

French, S. 1985. Group consensus probability distributions: A critical survey. In *Bayesian statistics 2*, ed. J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, Amsterdam: North-Holland.

Financial Supervisory Authority. 2005. AMA soundness standard. Working paper. FSA AMA Quantitative Expert Group, London.

Garthwaite, P. H., J.B. Kadane, and A. O'Hagan. 2005. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* 100, no. 470:680–700.

Gelfand, A. E., B. K. Mallick, and D. K. Dey. 1995. Modeling expert opinion arising as a partial probabilistic specification. *Journal of the American Statistical Association* 90, no. 430:598–604.

Genest, C., and K. J. McConway. 1990. Allocating the weights in the linear opinion pool. *Journal of Forecasting* 9, no. 1:53–73.

Genest, C., and M. J. Schervish. 1985. Modeling expert judgments for Bayesian updating. *Annals of Statistics* 13, no. 3:1198–1212.

Genest, C., and J. V. Zidek. 1986. Combining probability distributions: A critique and annotated bibliography. *Statistical Science* 1, no. 1:114–148.

Hogarth, R. M. (1975) Cognitive processes and the assessment of subjective probability distributions. *Journal of the American Statistical Association* 70, no. 350:271–294.

Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Keeney, D., and H. Raiffa. 1976. *Decisions with multiple objectives: Preferences and value trade-offs*. New York: John Wiley & Sons.

King, J. L. 2001. *Operational risk: Measurement and modeling*. New York: John Wiley & Sons.

Lambrigger, D., P. Shevchenko, and M. Wüthrich. 2007. The quantification of operational risk using internal data, relevant external data and expert opinions. *Journal of Operational Risk* 2, no. 3:3–27.

Lindley, D. V. 1983. Reconciliation of probability distributions. *Operations Research* 31, no. 5:866–880.

Morris, P. A. 1974. Decision analysis expert use. *Management Science* 20, no. 9:1233–1241.

Moscadelli, M. 2004. The modelling of operational risk: Experience with the analysis of the data collected by the Basel Committee. Working paper 517, Banca d'Italia, Rome.

O'Hagan, A. 1998. Eliciting expert beliefs in substantial practical applications. *The Statistician* 47, no. 1:21–35.

Pahlman, M., and A. Riabacke. 2005. A study on framing effects in risk elicitation. *Proceedings of the International Conference on computational intelligence for modelling, control and automation* 1, no. 2:689–694, Vienna, Austria.

Plous, S. 1993. *The psychology of judgment and decision making*. New York: McGraw-Hill.

sbAMA Working Group. 2003. Scenario-based AMA. Working paper, London.

Steinhoff, C., and R. Baule. 2006. How to validate OpRisk distributions. *OpRisk and Compliance* 1, no. 8:36–39.

Tversky, A., and D. Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185, no. 4157:1124–1131.

West, M. 1988. Modelling expert opinion. In *Bayesian statistics* 3, ed. J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith. Amsterdam: North-Holland.

West, M., and J. Crosse. 1992. Modelling probabilistic agent opinion. *Journal of the Royal Statistical Society*, Series B, 545, no. 1:285–299.

Winkler, R. L. 1967. The assessment of prior distributions in Bayesian analysis. *Journal of the American Statistical Association* 62, no. 319:776–800.

Winkler, R. L. 1968. The consensus of subjective probability distributions. *Management Science* 15, no. 2:361–375.