

Operational Risk – Modeling the Extreme

The views expressed in this paper are solely those of the authors and do not reflect official positions of the Office of the Comptroller of the Currency, the Federal Reserve Bank of Boston, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Richmond, the Federal Reserve System, or the National Institute of Statistical Sciences.

Abstract

In February 2009, the Office of the Comptroller of the Currency (OCC) and the National Institute of Statistical Sciences (NISS) organized a workshop titled, *Statistical Issues in Financial Risk Modeling and Banking Regulation*. The main objective of the workshop was to disseminate practical challenges faced in quantitative risk management to the academic community. This white paper highlights a selected list of challenges/open issues in operational risk modeling, particularly with respect to high tail estimation and tail dependence. Presentations from the workshop are available at http://www.niss.org/affiliates/financialrisk200902/financial_risk_home.html.

1 Problem Statement

The new Basel Capital Accord (Basel II)¹ requires financial institutions, for the first time, to reserve capital for operational risk. Under the new framework, banks are allowed to

¹See www.bis.org/list/bcbs/index.htm for various documents on the Basel II framework. For the U.S. implementation of the Basel II framework, see Federal Register [18].

calculate the regulatory capital charge for operational risk by using their own internal models, conditional on meeting the qualifying criteria as specified by Basel II² and as implemented by the home supervisor. The Basel II Accord specifies the use of an extreme quantile (Value-at-Risk, 99.9%, 1 year) as the measure of risk, which makes minimum capital requirements one of the few areas of law with a significant and explicit reliance on statistics.

Recognizing the range of internal models currently under development, the Basel Committee declined to provide any specific direction with regards to the nature of the models that can be used, apart from stating some general soundness standards for AMA:

... a bank must be able to demonstrate that its approach captures potentially severe "tail" loss events ... a bank must demonstrate that its operational risk measure meets a soundness standard comparable to that of the internal ratings-based approach for credit risk (i.e. comparable to a one year holding period and a 99.9th percentile). (Basel Committee [2], p. 151, §§ 667)

Development of the models are left at the discretion of the industry participants, as long as the models meet the qualification standards as outlined in the Accord³.

This white paper provides a brief summary of some of the quantitative challenges faced in operational risk modeling, particularly with respect to high tail estimation and tail dependence.

²Basel Committee [2], p. 149–155, §§ 664–679.

³The qualifying criteria for the AMA are stated at paragraphs 660–679, p. 149–155, Basel Committee [2].

2 Problem Background

Borrowing a terminology from psychiatry,⁴ one can characterize the development of operational risk in a few discrete stages. First came the “*denial*” stage, in which operational risk was dismissed as a quantitative field. For example, in ‘An Academic Response to Basel II’, Danielsson *et al.* [11] argued that operational risk is predominantly idiosyncratic, thus “rendering the need to regulate in order to prevent contagion unjustified”. Such a task was also rendered infeasible due to the complete lack of data at that time.

At the ensuing “*anger*” stage, many of the financial institutions considered operational risk as a purely compliance exercise and resisted to commit sufficient resources for its quantification. In our opinion, banks’ priority was to come up with as simple a framework as possible that would pass regulatory scrutiny.

When an established best practice does not exist, a natural starting point is to find another area where the variable under investigation has similar properties to that of problem at hand. The models that were found to be most feasible by the industry could broadly be classified as the different variants, or indeed combinations of the Loss Distribution Approach (LDA) and Extreme Value Theory (EVT). However, the problem at hand was more complicated than anticipated and required a more rigorous effort than a simplistic implementation of LDA and/or EVT,⁵ which lead to the “*bargaining*” stage.

At the bargaining stage, there was a conscientious effort to bridge the gap between practice and theory in light of the sound modeling expectations in the Basel II framework. With the accumulation of literature on operational risk, both the industry and regulators became more aware of the implementation issues and the limitation of operational risk models.

⁴Kubler–Ross’ five stages of grief. Note that, as in the five stages of grief, these stages of operational risk modeling neither apply to all financial institutions in the same order nor all stages are experienced by all.

⁵See, for example, Dutta and Perry [15] and Nešlehová *et al.* [23] for a critical assessment of the use of EVT in operational risk. Moscadelli [22] showed that operational losses were extremely heavy tailed. Loss data from six out of eight business lines came from an infinite–mean model.

At the final “*acceptance*” stage, the need for management and quantification of operational risk is, at least to our knowledge, no longer questioned. Many of the financial institutions are at the point of getting a working model and the outstanding issues/challenges are relatively well specified. As a sign of maturity, some are even thinking about going above and beyond of the regulatory requirements by further analyzing the underlying stochastic processes of operational losses with more sophisticated models.

We hope this paper will succinctly specify the outstanding issues on operational risk modeling and point researchers to the areas of practical importance where more work is needed. With this objective in mind, we have identified some issues that are discussed in the more technical/specific formulation section.

3 A Key Issue

A first key issue still concerns the use of VaR as an appropriate risk measure. One knows that VaR is not coherent, in particular VaR is super-additive (and hence defies diversification arguments) typically when the underlying data are either (i) very heavy-tailed; (ii) very skewed, or (iii) have special dependence. No doubt, OpRisk loss data exhibit (i) and (ii). The concern (iii) may at this point be of less importance. These unpleasant properties of VaR will keep on haunting the OpRisk modeler; therefore care is called for and further research on OpRisk relevant examples/counterexamples/consequences is needed; see for example Embrechts *et al.* [17].

Even though there have been rigorous efforts in collecting internal and external operational loss data over the last five years, these datasets are either confidential (internal data) or prohibitively expensive (external data). Academic interest has been slow mainly due to the absence of public data. As a consequence, main modeling issues did not percolate well to the academic and statistical community.

The ultimate object of interest is the 0.999 quantile of the aggregate loss distribution over

a one year time horizon. This presents problems as based on the data available (number and integrity, as well as complexity) such high quantile estimation is far beyond any quantiles that experts would have any ability to assess. See Chavez–Demoulin and Embrechts [9, 10] for a highlight of some of the related statistical problems in risk management.

4 A Remark on Holistic Modeling

An important development is that increasingly holistic OpRisk solutions are becoming available. OpRisk modeling is “not just” about a high quantile estimation from a loss distribution model, a more holistic view of the underlying issues is important here. As an example, see Soprano *et al.* [25]. For a unified approach including OpRisk, see for instance Brammertz *et al.* [7]. Though these publications by no means provide the definitive answers, statisticians ought to be aware of the broader risk management picture. An overly silo–building attitude to risk management is to be avoided.

5 More Technical/Specific Formulation

Basel II defines operational risk as “the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events”⁶. This definition includes legal risk, but excludes strategic and reputational risk. Legal risk includes, but is not limited to, exposure to fines, penalties, or punitive damages resulting from supervisory actions, as well as private settlements.⁷

The aggregate loss distribution itself is built up from distributions in various business lines and/or operational loss event types.⁸ As basic (internal, historical) data structure for (future)

⁶Basel Committee [2], p. 144, §§ 644.

⁷*Ibid.*, p. 144, footnote 97.

⁸Basel II specifies seven major loss event-type categories and eight business lines. These event types are internal fraud; external fraud; employment practices and workplace safety; clients, products and business

year t we have

$$\left\{ X_k^{t-i,b,\ell} : i = 1, \dots, T; \quad b = 1, \dots, r, \quad \ell = 1, \dots, s, \quad k = 1, \dots, N^{t-i,b,\ell} \right\}$$

where $X_k^{t-i,b,\ell}$ stands for the k -th loss of type ℓ for business line b in year $t-i$, $N^{t-i,b,\ell}$ for the number of such losses and T for the total number of years of available loss history. Typically, the loss-severities are truncated from below. The total historical loss amount for business line b in year $t-i$ becomes

$$L^{t-i,b} = \sum_{\ell=1}^s \sum_{k=1}^{N^{t-i,b,\ell}} X_k^{t-i,b,\ell} \quad (1)$$

and the total historical loss amount for year $t-i$ is

$$L^{t-i} = \sum_{b=1}^r L^{t-i,b}. \quad (2)$$

The Basel II standard corresponds to $r = 8$, $s = 7$.

In a first step (*i.e.*, disregarding other data sources) towards the calculation of minimum capital requirement for operational risk, one needs to estimate the 99.9% quantile of the distribution function $G_{L^t}(x) = P(L^t \leq x)$ of the total loss L^t ,

$$G_{L^t}^{-}(0.999) = VaR_{0.999}(L^t) = \inf \{x \in \mathbb{R} \mid G_{L^t}(x) \geq 0.999\},$$

where in the above notation, L^t consists of a severity (the X 's) – frequency (the N 's) model. Models of the type (1) and (2) are very well known from the (non-life) actuarial literature; see Chapter 10 in McNeil *et al.* [20]. The regulatory requirement is that a bank must use a sufficiently granular model that is commensurate with the bank's range of business activities and the variety of operational loss events to which it is exposed to, and that it does not combine business activities or operational loss events with demonstrably different risk profiles within the same loss distribution.⁹

practices; damage to physical assets; business disruption and system failure; and execution, delivery and process management. The business lines are corporate finance; trading & sales; retail banking; commercial banking; payment and settlement; agency services and custody; asset management; and retail brokerage.

⁹Federal Register [18], p. 69317.

The U.S. supervisors define granularity of operational risk models as a “unit of measure”,¹⁰ the level at which the bank’s operational risk quantification system generates a separate distribution of potential operational losses. The granularity of an operational risk model would be a function of the size, nature, scale and complexity of the financial institution and, naturally, availability of operational loss data. For a large international bank, one typically has up to a 56 cell–structure problem to model (based on the 8×7 matrix that results from combining the aforementioned 8 business lines and 7 event-types).

For notational convenience, we denote the chosen granular model (aggregate operational losses over a year) from (1), (2) as

$$S_D = S_1 + \cdots + S_d,$$

and its distribution function by $G_{S_D}(x) = P(S_D \leq x)$.

The firm–wide minimum capital requirement for operational risk (MRC^{OR}) is then estimated as the 99.9% quantile of the distribution G_{S_D} of the aggregate loss S_D

$$MRC^{OR} = VaR_{0.999}(S_D) = VaR_{0.999} \left(\sum_{k=1}^d S_k \right)$$

However, in general the structure of the total loss distribution G_{S_D} is complicated/unknown. To bypass this issue, banks should instead aggregate quantiles at unit–of–measure level and allow for a diversification benefit δ (≥ 0) such that

$$MRC^{OR} = VaR_{0.999} \left(\sum_{k=1}^d S_k \right) = (1 - \delta) \sum_{k=1}^d VaR_{0.999}(S_k). \quad (3)$$

However, unless the dependence assumptions underlying $\delta(> 0)$ are sound, and are robust to a variety of scenarios, implemented with integrity, and allow for the uncertainty surrounding the estimates,¹¹ regulators require the diversification benefit δ to be equal to zero. It is expected that assumptions regarding dependence will be conservative given the uncertainties surrounding dependence modeling for operational risk.¹²

¹⁰ *Ibid.*, p. 69317.

¹¹ Basel Committee [2], p. 152, §§ 669(d), and the Federal Register [18], p. 69317.

¹² Federal Register [18], p. 69317.

We have identified the following issues with regards to the above:

I1: Multivariate extremal behavior. For the class of elliptical models, the standard questions concerning risk measurement and capital allocation are well understood and behave much as in the exact multivariate normal case. However, as soon as one deviates from this class of elliptical models, the task becomes much more complicated.¹³ For instance, in the elliptical world, VaR as a risk measure is sub-additive (*i.e.*, $\delta \geq 0$ in (3)) meaning that the VaR of a sum of risks is bounded above by the sum of the individual VaRs. The multivariate extremal behavior needs to be understood and is perhaps more critical than the asymptotic behavior of the particular marginals in determination of the aggregate loss. There are several research topics worth exploring in this context. For instance, aggregation properties of quantiles at one level may not hold at other levels and this can result in incoherent models. One can indeed have loss models for which VaR is sub-additive at the very high quantile levels and super-additive at lower levels. In particular, the diversification benefit δ in (3) is not constant but depends on the quantile level; small changes in the quantile level may lead to large changes of δ . Mathematically, these aggregation properties of quantiles depend very much on the second-order properties of the underlying models; see Degen *et al.* [12] for a recent discussion of such issues. Possible further research includes questions on the sensitivity of δ with respect to different dependence structures. Further, in Embrechts *et al.* [17] several (counter-)examples are given leading to a better understanding of why coherence may fail when one moves away from ellipticality. In Degen *et al.* [14] and more recently Degen *et al.* [12] the issue of “phase-transition” from sub- to super-additivity (or vice versa) as a function of the quantile level α chosen in the VaR_α calculation is highlighted, and this in particular for the within OpRisk popular g -and- h example. The statistical (EVT-based) issues related to scaling between confidence levels is also discussed in Degen and Embrechts [13]. The probabilistic theory underlying the above papers is Multivariate Extreme Value Theory (MEVT) and Multivariate Regular Variation (MRV); see Mikosch [21], Beirlant *et al.* [3] or Resnick [24] for some basic references.

¹³Balkema and Embrechts [1].

I2: Adding statistical uncertainty. Whereas the probabilistic background is better understood by now, adding statistical model uncertainty may change the resulting additivity properties. One should therefore combine the probabilistic results together with the statistical ones (in a first instance coming from EVT) and check/quantify the influence on capital calculations.

I3: Model uncertainty. In line with I1 and I2 above, more work is needed on the general issue of model uncertainty and model robustness in financial risk management in general and operational risk more in particular. Possible research topics include:

(RT1) Quantify numerically bounds on risk measures of financial positions as a function of the typically incomplete model input. For instance by only specifying marginal distributions of the underlying risk factors together with “some” notion of dependence.

(RT2) Stochastic ordering of loss models: Flexibility inherent in AMA creates a broader range of practices in modeling. Understanding this dispersion is critical for regulators: Is it a result of differences in the underlying risk profiles or a result of the wide range of practice in modeling alternatives and/or combination mechanics of the alternative sources of data? Which particular aspects of the model determine the first-order outcome? For example, how sensitive is the reported capital to the alternative aggregation schemes? How can a bank minimise its regulatory capital by choosing an optimal aggregation path in the Basel II 8×7 business line/event-type matrix?

On both (RT1) and (RT2) there exists a considerable literature; look for instance for joint papers of Embrechts and Puccetti on www.math.ethz.ch/~embrechts.

- (RT3) Investigate the potential of robust optimization as for instance discussed in Ben-Tal and Nemirovski [5] and Ben-Tal *et al.* [4]. Aharon Ben-Tal has written lecture notes on the topic which can be downloaded from his (Technion) website (<http://iew3.technion.ac.il/Home/Users/morbt/rom.pdf>)

- (RT4) The theory of robust statistics is well-developed by now. Recently, several papers have appeared combining this theory with rare event estimation (which at first may seem somewhat contradictory), an example is Mancini and Trojani [19]. A more careful analysis with respect to operational risk modeling would be useful. In particular, stability of capital calculation must be taken into account: models leading to highly unstable capital estimations even when fitted with synthetical data drawn from a fixed distribution cannot be regarded as acceptable.

- (RT5) One needs to investigate further the applicability of resampling techniques (*i.e.*, the bootstrap) and simulation methodology (*i.e.*, importance sampling) in the context of operational risk. Some references can be found in Chavez-Demoulin and Embrechts [10]. A word of warning: traditional resampling methods typically fail for heavy-tailed loss models and rare event estimation.

I4: Some further research topics we like to stress are:

(RT6) **Benchmarking.** Developing credible benchmarks for operational risk exposure would help in better understanding the dispersion of the estimated capital charges. (This is of course related to topics under I3). In particular, one has to check the applicability of the analytic versions of the “One claim causes ruin” formula; see Embrechts *et al.* [16], Section 8.3.3. The formula mainly goes back to results for subexponential distribution functions as discussed in Embrechts *et al.* [16], A3.2 and has been popularized and extended in the context of operational risk by Böcker and Klüppelberg [6]. One definitely needs a careful analysis of the formula’s accuracy for a wide range of models. A possible starting point might be Degen *et al.* [12] whose second-order results may be shown to carry over to the Böcker-Klüppelberg framework. In particular it seems necessary to check its accuracy in relation with frequency (for high frequency cells, the formula seems not to be accurate) and which proportion of the distribution is taken into account for the modelling (POT).

(RT7) **Bayesian approaches.** In the Appendix to this paper we have included the **White Paper on Bayesian Operational risk**. Beyond these comments, we would like to add the wish for further research on Bayesian estimators for extreme events. A place to start is Chapter 11 in Beirlant *et al.* [3] and references therein. In particular, for the combination of external, internal and expert opinion data, one can apply the standard actuarial empirical Bayes approach known under the name of Credibility Theory; see Bühlmann and Gisler [8]. Concrete applications in the context of operational risk are to be found mainly through the work of Mario Wüthrich; see www.math.ethz.ch/~wueth.

(RT8) Comparing models. An issue which comes up over and over again in the context of operational risk is the competition between parametric, non-parametric and semi-parametric models. The recent work in statistics on the topic of, for instance, targeted maximum likelihood estimation and super learning may be useful here. In order to get the flavor, visit the website of Mark J. van der Laan, www.stat.berkeley.edu/~laan/.

References

- [1] G. Balkema and P. Embrechts (2007). *High Risk Scenarios and Extremes: A Geometric Approach*. Zurich Lectures in Advanced Mathematics, European Mathematical Society Publishing House.
- [2] Basel Committee (2006). *International Convergence of Capital Measurement and Capital Standards: A Revised Framework – Comprehensive Version*. June. Available at: <http://www.bis.org/publ/bcbs128.pdf>.
- [3] J. Beirlant, Y. Goegebeur, J. Segers and J. Teugels (2005). *Statistics of Extremes: Theory and Applications*. Wiley.
- [4] A. Ben-Tal, S. Boyd and A. Nemirovski (2006). Extending scope of robust optimization: comprehensive robust counterparts of uncertain problems. *Math. Program Ser. B*, 107, 63–89.
- [5] A. Ben-Tal and A. Nemirovski (2002). Robust optimization–methodology and applications. *Math. Program Ser. B* 92, 453–480.
- [6] K. Böcker and C. Klüppelberg (2005). Operational VaR: a closed-form approximation. *RISK* December 2005, 90–93.
- [7] W. Brammertz, I. Akkizidis, W. Breymann, R. Entin and M. Rustmann (2009). *Unified Financial Analysis. The Missing Link of Finance*. Wiley.

- [8] H. Bühlmann and A. Gisler (2005). *A Course in Credibility Theory and its Applications*. Springer.
- [9] V. Chavez–Demoulin and P. Embrechts (2009). Revisiting the edge, ten years on. *Communications in Statistics–Theory and Methods*, to appear.
- [10] V. Chavez–Demoulin and P. Embrechts (2009). An EVT primer for credit risk. In A. Lipton and A. Rennie (Eds.) *Handbook of Credit Derivatives*, Oxford University Press, to appear.
- [11] J. Danielsson, P. Embrechts, C. Goodheart, C. Keating, F. Muennich, O. Renault and H.S. Shin (2005). *An academic response to Basel II*. Special Paper No 130, Financial Markets Group, LSE. Available at <http://www.bis.org/bcbs/cacomments.htm>
- [12] M. Degen, D. D. Lambrigger and J. Segers (2009). Risk concentration and diversification: second-order properties. Preprint ETH Zurich.
- [13] M. Degen and P. Embrechts (2009). Scaling of high–quantile estimators. Preprint ETH Zurich.
- [14] M. Degen, P. Embrechts and D. D. Lambrigger (2007). The quantitative modelling of operational risk: between g -and- h and EVT. *ASTIN Bulletin* 37(2), 265–29.
- [15] K. Dutta and J. Perry (2006). *A Tale of Tails: An Empirical Analysis of Loss Distribution Models for Estimating Operational Risk Capital*. FRB of Boston Working Paper No. 06-13. Available at SSRN: <http://ssrn.com/abstract=918880>
- [16] P. Embrechts, C. Klüppelberg and T. Mikosch (1997). *Modelling Extremal Events for Insurance and Finance*. Springer.
- [17] P. Embrechts, D. D. Lambrigger and M.V. Wüthrich (2009). Multivariate extremes and the aggregation of dependent risks: examples and counter-examples. *Extremes* 12, 107–127.

- [18] Federal Register (2007). *Risk-Based Capital Standards: Advanced Capital Adequacy Framework – Basel II; Final Rule, Federal Register*, Vol. 72, No, 235, December 7, 2007, p. 69287–69445.
- [19] L. Mancini and F. Trojani (2008). Robust value at risk prediction. Preprint, University of Zurich.
- [20] A.J. McNeil, R. Frey and P. Embrechts (2005). *Quantitative Risk Management: Concepts, Techniques, Tools*. Princeton University Press.
- [21] T. Mikosch (2004). Modeling Dependence and Tails of Financial Time Series. In: *Extreme Values in Finance, Telecommunications and the Environment*, B. Finkenstädt and H. Rootzén (eds.), Chapman and Hall, 185–286.
- [22] M. Moscadelli (2004). The modeling of operational risk: experience with the analysis of the data collected by the Basel Committee. Technical Report 517, Banca d’Italia.
- [23] J. Nešlehová, P. Embrechts and V. Chavez–Demoulin (2006). Infinite–mean models and the LDA for operational risk. *Journal of Operational Risk* 1, 3–25.
- [24] S. Resnick (2006). *Heavy-Tail Phenomena - Probabilistic and Statistical Modeling*. Springer, New York.
- [25] A. Soprano, B. Crielgaard, F. Piacenza and D. Ruspantini (2009). *Measuring Operational and Reputational Risk. A Practitioner’s Approach*. Wiley.

6 Appendix

White Paper on Bayesian Operational Risk

Problem Statement

Calculation of operational risk capital under Basel II involves estimation of the 0.999 quantile of the annual aggregate operational loss distribution for the bank. The aggregate operational loss distribution is built up from distributions in various business lines and/or loss types. Given that actual data are scarce, coherent incorporation of expert opinion is important. While there most likely are no experts who understand the bank-level aggregate loss distribution, there are experts who understand the distribution in individual business line/loss type cells (e.g., experts in credit card fraud, legal exposure). The Bayesian approach assures coherence.

Problem Background

The efficiency of resource allocation in a modern economy depends crucially on the quality of numerous decentralized decisions on credit allocation. These decisions depend on inference about small probabilities and rare events, where data information can be sparse but expert information is clearly available. Large, internationally active banks must follow internationally negotiated guidelines. The Basel II (B2) framework (Basel Committee on Banking Supervision (2006)) for calculating minimum capital requirements provides for banks to use models to assess the variety of risks to which they are exposed. All aspects of these models (specification, estimation, validation) will have to meet the scrutiny of national supervisors.

This paper will focus on operational risk; see Sections 1–7. We will in particular focus on incorporating expert information in estimating the severity distribution for operational losses.

Our approach begins with the observation that there is some information available about the loss severity distribution in addition to the available historical data. The simple fact that a business is ongoing shows that some risk assessment is occurring and that exposure to operational risk (together with other relevant risks) in that business are deemed to be less than (in some sense) the potential profits to be made. This additional information should be organized and incorporated in the analysis in a sensible way, specifically represented in a probability distribution. It can be shown that beliefs that satisfy certain consistency requirements, for example that the believer is unwilling to make sure-loss bets, lead to measures of uncertainty that combine according to the laws of probability: convexity, additivity and multiplication. See for example DeGroot (1970). This information should be combined with data information as represented in the relevant likelihood function. This combination of information is easy to do once the information is represented in probability distributions. The final distribution should represent both data and expert information about the loss severity or exposure.

Current Obstacles to Solution

Obstacles are both practical and technical. The first practical obstacle is identifying relevant experts and training them to think about uncertain severities or exposures in terms of probability distributions. This can be addressed by training and accumulating experience in probability assessment. A second practical obstacle is that staff at many banks may not currently have the tools/skills to perform a Bayesian analysis. The principal difficulty in applying the Bayesian approach is that unfamiliar thinking is required. It is not easy to quantify uncertainty. Quantification of uncertainty requires comparison with a standard. One standard for measuring uncertainty is a simple conceptual experiment, such as drawing balls from an urn at random as above, or sequences of coin flips. Assign probabilities by comparison. For example A is about as likely as seeing three heads in 50 throws of a fair coin. Sometimes it is easier to assign probabilities by considering the relative likelihoods of events and their complements. Thus, either A or "not A" must occur. Some prefer to recast this

assessment in terms of betting. These possibilities and others are discussed in Berger (1980). Thinking about uncertainty in terms of probabilities requires effort and practice (possibly explaining why it is so rarely done). Nevertheless it can be done once experts are convinced it is worthwhile. Indeed, there is experimental evidence in game settings that elicited beliefs about opponents' future actions are better explanations of responses than empirical beliefs – Cournot or fictitious play – based on weighted averages of previous actions. For details see Nyarko and Schotter (2002) and Kiefer (2009). O'Hagan et al. (2006) discusses elicitation techniques and several applications. A technical issue is combining the experts' views on the marginals for each cell into an aggregate loss distribution. There is virtually no expert opinion on dependence, and very little data information either. The usual current approaches rely heavily on parametric assumptions, allowing inference about extreme events from data on common events. This dependence problem is difficult, but we will not have a solution to propose at this point. The ultimate object of interest is the 0.999 quantile of the aggregate loss distribution. This presents problems, as such a high quantile is far beyond any quantiles that experts would have any ability to assess.

Potential Avenues to Solution and Data Requirements

We suggest that when performing probability assessment, it is more practical to elicit expert opinion regarding quantiles rather than moments. Given the assessments, we need to represent the information in a probability distribution. Here, a maximum entropy approach might be applicable, as parametric approaches may inappropriately add additional information that was not contained in the experts' assessments.

More Technical Formulation

Prior information assessed would consist of various quantiles of the loss distribution. Next, get the maximum entropy density, which is piecewise uniform as a first pass. Then smooth using kernel smoothing, as that is more consistent with experts' beliefs. The final smoothed

prior is shared with the expert to make sure this represents what he/she really believes. This work will also touch on ongoing questions relating to the combination of elements. Current industry practice has experts consulting both internal and external data in the process of providing their quantile estimates. We note that the former is problematic within the Bayesian context, because what they would be providing in this situation would be a posterior rather than a true prior distribution. In layman’s terms, this would be double-counting. In the latter case, use of external data as a scenario input and also as a separate direct input would be considered double-counting for similar reasons. A key consideration will be whether to estimate exposure or severity. This remains open. Exposure may be difficult to think about in terms of prior assessment, although it is easy to write down a nonparametric (partial) likelihood. Severity itself might be easier to think about, and perhaps the assessed quantiles of a severity distribution (rather than a prior distribution on exposure) could be combined with a fairly rich parametric model to yield an estimated severity distribution. If you run the model on an annual basis, you would use the posterior from last year as this year’s prior. You would need to consult with the experts only to the extent that there were significant changes in exposure that were not captured in internal data.

References

Basel Committee on Banking Supervision (2005): Basel Committee Newsletter No. 6: Validation of Low- Default Portfolios in the Basel II Framework, Discussion paper, Bank for International Settlements.

Berger, J. O. 1980): *Statistical Decision Theory: Foundations, Concepts and Methods*. Springer.

Garthwaite, P. H., J. B. Kadane, and A. O’Hagan (2005): Statistical Methods for Eliciting Probability Distributions, *Journal of the American Statistical Association*, 100, 680–700.

Kiefer, N. (2009): Default Estimation, Correlated Defaults, and Expert Information, forthcoming in *Journal of Applied Econometrics*.

Nyarko, Y., and A. Schotter (2002): An Experimental Study of Belief Learning Using Elicited Beliefs, *Econometrica*, 70(3), 971–1005.

O’Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow (2006): *Uncertain Judgements: Eliciting Experts’ Probabilities*. Wiley.

7 Group Members

Emre Balta, Office of the Comptroller of the Currency

Santiago Carrillo Menéndez, Universidad Autonoma de Madrid

Patrick de Fontnouvelle, Federal Reserve Bank of Boston

Paul Embrechts, ETH Zurich

Bakhodir Ergashev, Federal Reserve Bank of Richmond

Kam Hamidieh, Rice University

Nicholas M. Kiefer, Cornell University and Office of the Comptroller of the Currency

Ken Swenson, Federal Reserve Bank of Chicago