Large Portfolio Losses

Amir Dembo, Jean-Dominique Deuschel, and Darrell Duffie

First Version: March 29, 2002
This Version: March 11, 2003

To appear in *Finance and Stochastics*

Abstract: This paper provide a large-deviations approximation of the tail distribution of total financial losses on a portfolio consisting of many positions. Applications include the total default losses on a bank portfolio, or the total claims against an insurer. The results may be useful in allocating exposure limits, and in allocating risk capital across different lines of business. Assuming that, for a given total loss, the distress caused by the loss is larger if the loss occurs within a smaller time period, we provide a large-deviations estimate of the likelihood that there will exist a sub-period of the future planning period during which a total loss of the critical severity occurs. Under conditions, this calculation is reduced to the calculation of the likelihood of the same sized loss over a fixed initial time interval whose length is a property of the portfolio and the critical loss level.

1 Introduction

We[1] provide a large-deviations approximation of the tail distribution of total financial losses on a portfolio consisting of many positions. Applications include the total default losses on a bank portfolio, or the total claims against an insurer. A key assumption is that, conditional on a common "correlating"

[1]Amir Dembo is with the Department of Statistics, Stanford University. His research was partially supported by NSF grant #DMS-0072331. Jean-Dominique Deuschel is with the Department of Mathematics, Technische Universität, Berlin. His research was partially supported by DFG grant #663/2-3 and DFG FZT 86. Darrell Duffie is with the Graduate School of Business, Stanford University. We are extremely grateful for research assistance by Nicolae Gârleanu and Gustavo Manso, for conversations with Michael Gordy, and for comments from Michael Stutzer, Peter Carr, David Heath, and David Siegmund.

factor Y, position losses are independent. For example, in the case of default losses, Y could be the state of the business cycle. For the case of an insurance portfolio, Y could include indicators of events causing multiple claims, such as epidemics or natural catastrophes. Gordy [3] developed an asymptotic estimate of the portfolio loss distribution that corrects for "granularity" in presence of a one-dimensional source of default correlation.

The results include explicit calculations, conditional on a large portfolio loss, of the probability of loss for each position, and the distribution of the size of the position loss. This information may be useful in allocating exposure limits for each type of position, and in allocating risk capital across different lines of business.

We also address the fact that, for large losses, financial distress costs are more severe if the losses occur over a relatively short period of time. Sudden losses may cause extreme cash-flow stress, and investors may require more favorable terms when offering new lines of financing over short time periods, within which they may have a limited opportunity to gather information about the credit quality and long-term prospects of a distressed financial institution. Our results include conditions under which a large-deviations estimate of the likelihood of a failure-threatening loss during some sub-interval of time during a given planning horizon can be calculated from the likelihood of the same size loss in a certain fixed "key time horizon." From this key time horizon, one can then estimate the conditional distribution of losses on each type of position given the large portfolio loss of concern. Again, this information may be of assistance in structuring a large portfolio so as to withstand severe losses.

2 Motivation

Consider, for example, an insurance company with a large portfolio of property-damage policies, each generating a stream of premiums whose present market value exceeds the present market value of the associated uncertain future damage claims. (This should be the case, at least at origination, for any policy that the insurance company is willing to offer.) In perfect capital markets, the event that total damage claims, net of premiums, exceeds the capital of the insurance firm would trigger a recapitalization of the insurance company, say by selling equity. The alternative of bankruptcy would entail losing the net positive market value of the remaining policies. Losses up to the current date are sunk costs, and therefore irrelevant to the decision to re-capitalize. Likewise, the initial level of capital would be irrelevant in perfect markets, for re-capitalization could occur at any time. In the classical ruin-theory insurance model (for example [5, 6, 9]), however, the insurance company with a given initial capital faces "ruin" whenever the remaining capital, net of gains and losses, reaches zero. We depart from this classical ruin-theory approach, as it ignores the incentive to re-capitalize at low levels of capital.

While classical ruin theory unrealistically ignores re-capitalization, it correctly focuses on the importance, given the reality of imperfect capital markets, of a level of capital that is likely to withstand unexpectedly large aggregate losses. In the event of insufficient capital to cover losses, raising new capital is, in practice, expensive. For example, in order to entice new equity investors to offer a given amount of capital, a firm that is better informed about its prospects than are new investors would generally need to offer investors equity whose value to the firm is worth more than the amount of capital raised. (See, for example, [7].) In sufficiently severe cases, especially when large unexpected losses occur over a relatively short period of time, it may be impossible to raise enough capital to meet obligations, and the ongoing franchise value of the firm may be severely impaired or lost.

This paper provides some analytical guidance on the dependence of large-loss probabilities on the structure of a portfolio with a large number n of positions, and on the "most likely way" that a large loss can occur. For example, given a large loss, we calculate, for each type of position, the conditional likelihood of loss on each type of position as well as the conditional distribution of exposure in the event of loss. These conditional calculations are to be interpreted in the asymptotic (large n) sense of the Gibbs conditioning principle (see, for example, [1, Section 7.3]). For instance, suppose positions are of two types. A given fraction, say high-quality borrowers who are granted large amounts of credit, experience large losses with small probabilities. The remainder, say low-quality borrowers granted less credit, experience smaller losses with higher probabilities. We show how to estimate, under conditions, the expected loss on a policy of each type in the event of large loss. For each given position type, this depends on the size of the total loss considered, and on the probability distributions of losses on all other positions.

Given the costs of raising capital in a distress scenario, this in turn leads to an analytical estimate of the financial distress costs attributable to each policy type, and allows for the structuring of the portfolio so as to balance the effects of each type of policy on the total financial distress costs. For example, we estimate the sensitivity of the large-loss probability to the fraction of the portfolio of any given type, and to the scale of the exposure on a given type of position. With bank portfolios, for instance, we might be interested in how the probability that total default losses exceed a given threshold is sensitive to the exposure limit for counter-parties of a given credit rating.

These calculations could also be done by brute-force simulation, and under less restrictive conditions, but the large-deviations approach offers insights that might not be apparent from mere numerical results.

After considering a static setting, we turn in Section 4 to the estimation of the probability that, during some sub-interval of time during a given planning horizon, there will exist financial losses that exceed a threshold, allowing for the fact that financial distress costs are larger when losses are concentrated over a smaller period.

3 Portfolio Calculations

A probability space $(\Omega, \mathcal{F}, \mathbf{P})$ is fixed. For a portfolio with a finite number n of positions, position i will, at the time of revaluation of the portfolio, experience a loss of $Z_i U_i$, where Z_i has outcomes 0 (for no loss) and 1 (for non-zero loss), and $U_i \geq 0$ is the "exposure" (the amount that would be lost in the event there is a loss). For example, with a loan portfolio, $\{Z_i = 1\}$ is the event that i defaults and also has a non-zero exposure, and U_i is the market value of the exposure, net of default recoveries. For the case of insurance claims, [10] analyses associated statistics for fire-insurance policies.

The total loss on the portfolio is thus $L_n = \sum_{i=1}^{n} Z_i U_i$. Our main objective is to characterize, for a given x, the loss probability $\mathbf{P}(L_n \geq nx)$, and, conditional on the event that $L_n \geq nx$, the probability distributions of U_i and Z_i. These conditional distributions are relevant to the structuring of the portfolio from the viewpoint of trading off the expected profit from each position type against the benefits of mitigating large losses.

We suppose that a "macro-environmental" variable Y can be chosen so that, conditional on Y, the loss variables $Z_1, U_1, Z_2, U_2, \ldots, Z_n, U_n$ are independent. For example, in modeling credit risk, Y could incorporate key business-cycle or industry-performance variables. For insurance risk, Y could incorporate the major events such as natural disasters that affect the likelihood of individual losses. For simplicity, we also suppose that, for all i,

U_i and Y are independent and that Y is discretely-valued (merely to avoid numerical integration approximations).

We assume that a portfolio contains k types of positions, in the sense that the distribution of $\{Y, U_i, Z_i\}$ is the same for any position i of a given type α. The types are fixed. We use the notation "U_α" to denote a generic U_i variable of type α, and likewise use "Z_α," and so on. We let $\delta_\alpha(Y) = \mathbf{P}(Z_\alpha = 1 \,|\, Y)$, the Y-conditional loss probability for type α.

For bank portfolios, one may think of a "type" as a credit rating. The common distribution of the loss within a given rating is the distribution of exposures for that rating. For purposes of this analysis, we ignore information that might in practice distinguish among different positions of a given rating. In this sense, the probability distribution of U_i is intended to capture, to some reasonable extent, the effect of variation of exposure within rating. An alternative is to distinguish types by both rating and exposure class.

3.1 Large Deviations

In order to calculate the effect of increasing the number n of positions, we suppose that the number of positions of type α is a fixed fraction $q_\alpha > 0$ of the total number n (in essence it is $q_\alpha(n)$ that is closest to q_α while $nq_\alpha(n)$ is an integer. To simplify notations we suppress this dependence of q_α on n).

For computing the probability of an unexpectedly large loss, the important object to study is the cumulant generating function $L(\cdot \,|\, Y)$ defined by

$$\begin{aligned} L(s \,|\, Y) &= \frac{1}{n} \log \mathbf{E} \left(e^{sL_n} \;\middle|\; Y \right) \\ &= \sum_\alpha q_\alpha \log \left(1 - \delta_\alpha(Y) + \delta_\alpha(Y) M_\alpha(s) \right), \end{aligned} \tag{1}$$

where $M_\alpha(s) = \mathbf{E}(e^{sU_\alpha})$ is the moment generating function for loss exposure of type α, assumed finite for small enough $s > 0$. For example, we might suppose that U_α is distributed exponentially with parameter β_α, in which case

$$L(s \,|\, Y) = \sum_\alpha q_\alpha \log \left(1 - \delta_\alpha(Y) + \delta_\alpha(Y) \frac{\beta_\alpha}{\beta_\alpha - s} \right), \tag{2}$$

when $s < s_0 = \min_\alpha \beta_\alpha$, and $L(s \,|\, Y) = \infty$ otherwise. One then introduces the Legendre Fenchel transform $L^*(\,\cdot\,|\,Y)$ of $L(\,\cdot\,|\,Y)$, by

$$L^*(x \,|\, Y) = \sup_s \{sx - L(s \,|\, Y)\}.$$

That is, one solves for $s = s(x, Y)$ the equation

$$x = L'(s \,|\, Y), \tag{3}$$

and sets

$$L^*(x \,|\, Y) = s(x, Y)x - L(s(x, Y) \,|\, Y).$$

Let $s_0(Y) = \sup\{s : L(s \,|\, Y) < \infty\}$. If $s_0 < \infty$, we assume that $L'(s \,|\, Y)$ is unbounded for $s \uparrow s_0$ (for example, as is true when U_α has an exponential distribution for each α). Under this mild regularity condition, equation (3) has a unique positive solution $s(x, Y)$ provided $x_1(Y) < x < x_*(Y)$, where

$$x_1(Y) = L'(0 \,|\, Y) = \frac{1}{n} E(L_n \,|\, Y) = \sum_\alpha q_\alpha \delta_\alpha(Y) \mathbf{E}(U_\alpha), \tag{4}$$

the mean loss per position given Y, and where

$$x_*(Y) = \sum_{\{\alpha :\, \delta_\alpha(Y) > 0\}} q_\alpha \operatorname{ess\,sup}\{U_\alpha\} \tag{5}$$

is the maximal value of the average loss per position. For example, for the case above of exponentially distributed U_i, we have a unique positive solution for any $x > x_1(Y) = \sum_\alpha q_\alpha \delta_\alpha(Y)/\beta_\alpha$. (Here, $x_* = \infty$ assuming $\delta_\alpha(Y) > 0$ for some α.) Then, the probability of a "large" loss given Y is given by the precise large-deviations (LD) approximation

$$\mathbf{P}(L_n > nx \,|\, Y) = p_n(Y)(1 + \mathrm{o}(1)), \tag{6}$$

where, for $x < x_*(Y)$,

$$p_n(Y) = (2\pi n s^2 L''(s))^{-1/2} e^{-nL^*(s \,|\, Y)}, \tag{7}$$

using for $s = s(x, Y)$ the solution of (3), and $p_n(Y) = 0$, for $x > x_*(Y)$. To adapt the derivation of such approximations in [1, Theorem 3.7.4] to our setting, one just shows that the distribution function of $(L_n - nx)/\sqrt{nL''(s \,|\, Y)}$ is uniformly within $o(n^{-1/2})$ of $\Phi(u) + n^{-1/2}\epsilon(u)$, where $\Phi(\,\cdot\,)$ is the standard

Normal distribution function and $\epsilon'(\,\cdot\,)$ is some uniformly bounded, absolutely integrable function, such that $|\epsilon(u)| \to 0$ as $|u| \to \infty$. Using the independence of $\{Z_i U_i\}$, it is not hard to establish the latter estimate, starting from the Berry-Esseen bound for the contribution of each fixed type α.

If $x < x_1(Y)$, then, for large n, we have by the law of large numbers (LLN), that (6) holds for $p_n(Y) = 1$, while, obviously, $\mathbf{P}(L_n > nx \,|\, Y) = 0$ if $x > x_*(Y)$. Thus,

$$\mathbf{P}(L_n > nx) = p_n(1 + \mathrm{o}(1)), \tag{8}$$

where $p_n = \sum_y \mathbf{P}(Y = y) p_n(y)$ is the weighted average of the above approximations, with weights given by $\mathbf{P}(Y = y)$.

In a somewhat different setting for modeling default-risk correlation and portfolio heterogeneity, Gordy [4] also computed the cumulant generating function of the portfolio loss distribution, using a saddle-point approximation of the higher-order terms represented in (8) by $p_n \mathrm{o}(1)$).

For large n, conditioning on the the excessive loss $\{L_n > nx\}$ and Y, we have the following estimates of the conditional distribution of the losses on each position, letting Q_α denote the original law of a loss exposure U_α of type α:

- The individual losses $Z_i U_i$ remain independent and of identical law within each type α, although Z_i and U_i themselves are not independent (as they were, unconditionally).

- The loss probability for a position is changed to the larger value

$$\widehat{\delta}_\alpha(Y) = \frac{\delta_\alpha(Y)}{\delta_\alpha(Y) + (1 - \delta_\alpha(Y))/M_\alpha(s)}. \tag{9}$$

- For those i of type α for which $Z_i = 1$, the conditional law of U_i is $M_\alpha(s)^{-1} \exp(su)\, dQ_\alpha(u)$, where $s = s(x, Y)$ is the solution of (3).

- For those i of type α for which $Z_i = 0$, the conditional law of U_i is the original law, Q_α.

For example, if U_i is unconditionally exponentially distributed with parameter β_α, then a position i causing a loss ($Z_i = 1$) has a loss exposure U_i that is conditionally exponential with parameter $\beta_\alpha - s(x, Y)$. This applies for cases

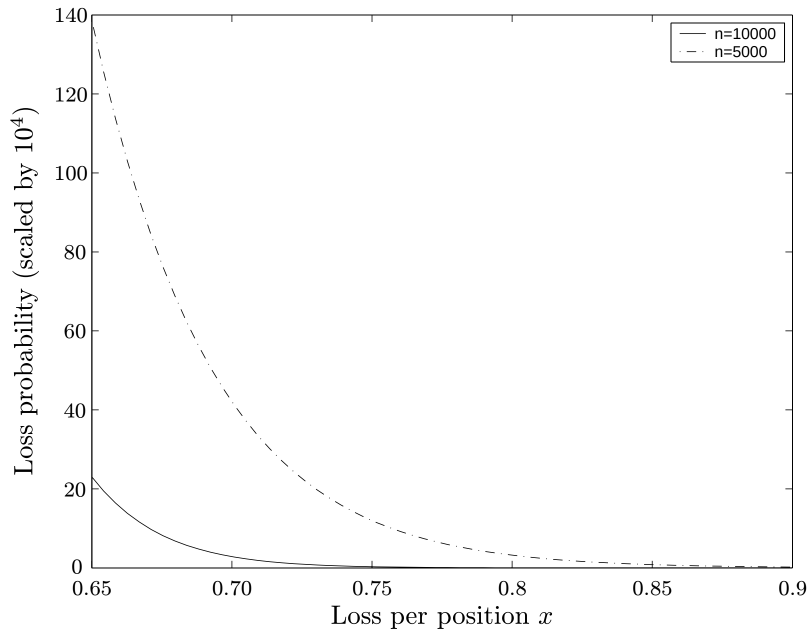

Figure 1: Probability of a loss per position exceeding x, for two portfolio sizes.

in which the environmental variable Y satisfies $x > x_1(Y)$, that is, where large deviations apply. (A rigorous, precise statement of this "principle" and its derivation, is along the lines of [1, Sections 7.3.1 and 7.3.3].) Otherwise, that is, with $x < x_1(Y)$, because the mean loss exceeds x, the impact of conditioning on $\{L_n > nx\}$ causes no change in the law of $Z_i U_i$.

3.2 Numerical Example

We illustrate with a simple numerical example that will be extended as we later consider additional calculations.

Suppose a bank has $n = 10,000$ borrowers or other forms of counter-parties of two types, high rated ($\alpha = 1$) and low rated ($\alpha = 2$). There are two macro-environments, "growth" (Y has outcome g) and "recession" (Y has outcome b). High-rated counter-parties constitute half ($q_1 = q_2 = 0.5$) of the positions. In a growth environment, the high-rated default probability is $\delta_1(g) = 0.001$, and the low-rated default probability is $\delta_2(g) = 0.004$. In a recessionary environment, we suppose that $\delta_1(b) = 0.0015$, and $\delta_2(b) = 0.10$.

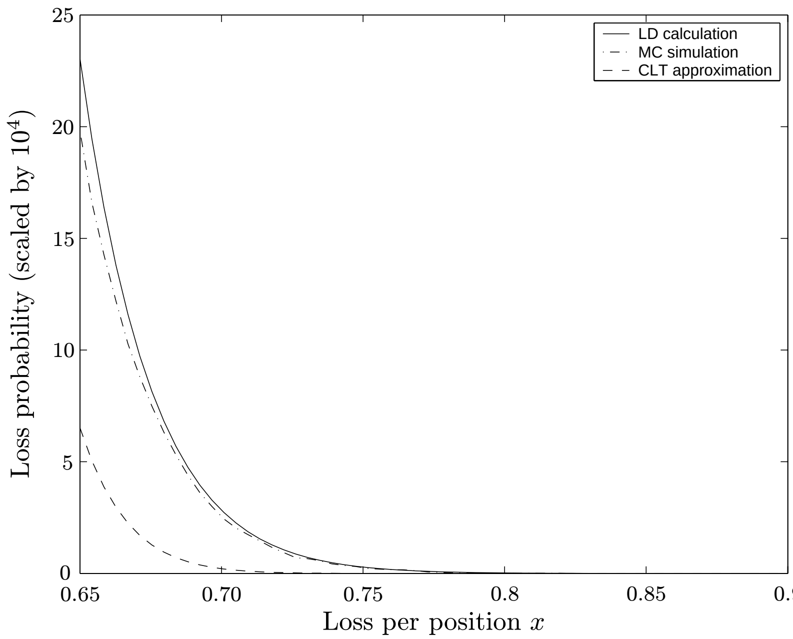

Figure 2: Comparison of approximations based on large deviations, brute-force Monte Carlo simulation, and Central Limit Theorem, of the unconditional probability that the loss per position exceeds x, at base-case parameters and with $n = 10,000$.

These illustrative parameters are very roughly consistent with the data for high-grade and speculative-grade debt, as for example in [8]. We suppose that growth occurs with probability $\mathbf{P}(Y = g) = 0.7$ (so recession has probability 0.3). We will suppose that the exposures given loss are exponential with means (in units of, say, \$10,000) of 100 and 10 for high-rated and low-rated counter-parties, respectively. (That is, $E(U_1) = 100$ ($\beta_1 = 0.01$) and $E(U_2) = 10$ ($\beta_2 = 0.1$).) The LLN average loss per position is $x_1(g) = 0.07$ in a growth economy and $x_1(b) = 0.575$ in a recessionary economy.

Figure 1 shows the large-loss probability, as a function of $x > x_1(b)$ for $n = 10,000$, and also for a less diversified portfolio ($n = 5,000$) that is otherwise identical. Figure 2 illustrates a comparison between three methods of the calculating the large-loss probabilities, Monte Carlo simulation (with 1 million scenarios), the large-deviations analytical method developed here, and the approximation based on the central limit theorem. As this example

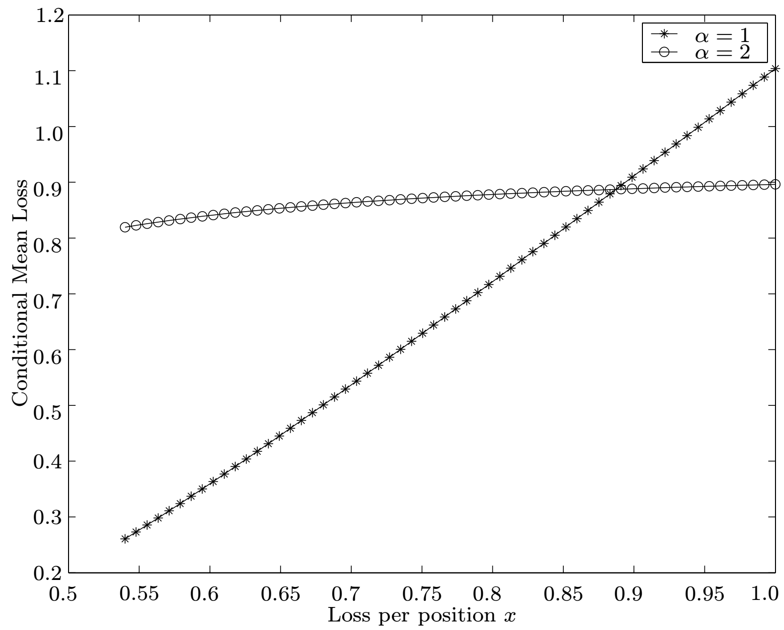

Figure 3: Large-Deviation estimates of large-loss-conditional expected position loss of each type ($\alpha = 1$ and $\alpha = 2$), given that the total loss per position exceeds x.

shows, approximation based on the central limit theorem is not as well suited for an examination of the extreme loss distribution.

Figure 3 shows how the high-rated and low-rated mean losses, conditional (in the sense of (9)) on a large average loss x, depend on x in the bad state ($Y = b$).

Because the exposures given default are exponential in this example, the event of an average position loss of $x = 0.6705$ that occurs with LD-estimated probability 0.001 is associated with conditional exposures for defaulting high-rated counter-parties that remain exponentially distributed but with a substantial increase in mean exposure from the unconditional mean of 100 to a conditional mean of

$$p_g \times (0.01 - s(x, g))^{-1} + p_b \times (0.01 - s(x, b))^{-1} = 155.53,$$

where $p_g = p_n(g)\mathbf{P}(Y = g)/p_n$ is the large-deviations estimate of $\mathbf{P}(Y = g \,|\, L_n > nx)$ based on (7). The low-rated counter-party mean exposures are of course less dramatically affected by this large-loss event, from 10 to 10.37, similarly calculated.

3.3 Sensitivity Analysis

We now consider the marginal impact of a change in the structure of a portfolio on large-loss probabilities. For example, in the case of a bank portfolio, one may be interested in the impact of a change in policy toward increasing the fraction q_α of borrowers or over-the-counter (OTC) derivative counter-parties of given credit rating. In practice, exposure limits are based on credit rating (the higher the rating, the larger the allowed exposure). For example, banks typically limit each of their OTC counter-parties, based on rating, to positions whose probability distribution of exposure given loss has a given high (say 95%) confidence level of no more than a given number. When considering marginal adjustments in exposure limits, we will let a_α be a scaling parameter for the distribution of U_α. For example, U_α may be distributed uniformly on the interval $[0, a_\alpha]$, or may be distributed exponentially with parameter $\beta_\alpha = 1/a_\alpha$.

Recalling that p_n is the precise LD approximation of $\mathbf{P}(L_n > nx)$ given by (8), we get

$$\frac{d \log p_n}{dq_\alpha} = \frac{1}{p_n}\frac{dp_n}{dq_\alpha} = n\mathbf{E}[L_\alpha(s \,|\, Y) \,|\, L_n > nx] + \mathrm{O}(1) \tag{10}$$

and

$$\frac{d \log p_n}{da_\alpha} = \frac{1}{p_n}\frac{dp_n}{da_\alpha} = n\frac{q_\alpha}{a_\alpha}\mathbf{E}[sL'_\alpha(s \,|\, Y) \,|\, L_n > nx] + \mathrm{O}(1), \tag{11}$$

where $L_\alpha(s \,|\, Y) = \log \mathbf{E}(e^{sZ_\alpha U_\alpha} \,|\, Y)$ is the cumulant generating function for loss in a loan of type α (so that $L(s \,|\, Y) = \sum_\alpha q_\alpha L_\alpha(s \,|\, Y)$). The effect of conditioning on $\{L_n > nx\}$ in the weighting of the different outcomes of Y is computed via Bayes formula using the LD approximation $p_n(Y)$ of (7) for $\mathbf{P}(L_n > nx \,|\, Y)$. The parameter $s = s(x, Y)$ that solves $x = L'(s \,|\, Y)$ is as above.

In deriving the above formulas, we assume that $x_* > x > x_1(y)$ for every outcome y of Y, so that the LD approximation applies in all economy states. The derivatives above typically grow linearly with n. The extra constants (independent of n) marked in these formulas as "O(1)" are explicit and computable, but complicated. In general, one sets $s(x, y) = 0$ for those y such that $x < x_1(y)$, reflecting the fact that these outcomes do not affect the linear growth of the derivatives of p_n. We repeat that these formulas are not necessarily derivatives of $\mathbf{P}(L_n > nx)$, but rather of the LD approximation p_n of (8).

In our illustrative example, the average position loss x that is exceeded with an LD estimated probability of 0.001 is $x = 0.7343$. The estimated sensitivity of this large-loss probability to increasing the fraction of low-rated counter-parties is

$$\begin{aligned} \frac{d \log p_n}{dq_2} &= p_g \times 3.11 + (1 - p_g) \times 29.58 \\ &= 29.58, \end{aligned}$$

where $p_g = p_n(g)\mathbf{P}(Y = g)/p_n$ is the large-deviations estimate of $\mathbf{P}(Y = g \,|\, L_n > nx)$ based on (7). The estimated sensitivity to increasing exposure limits of high-rated counter-parties is

$$\begin{aligned} \frac{d \log p_n}{da_1} &= p_g \times 0.4676 + (1 - p_g) \times 0.0863 \\ &= 0.0863. \end{aligned}$$

One can likewise estimate other sensitivities.

4 Time Evolution of Losses

For estimating the risk of large loss in a portfolio, the conventional value-at-risk approach is to calculate the probability distribution of losses at a fixed time horizon, for example as explained in [2, Sections 2.4.1 and 13.4]. The ability of a financial institution to replace lost capital (or otherwise restructure its portfolio) in order to return to a safe condition depends, however, on the period of time over which the loss occurs. For example, re-capitalizing for a given large loss that occurs over one year causes a smaller financial distress cost than would be caused by the same-sized loss occurring over a one-week period. Rather than attempting to treat the high-dimensional joint distribution of losses at each time horizon, we will estimate the likelihood that there exists some sub-interval of time over a given planning period, say one year, over which financial distress costs exceed some critical level. Under conditions, we will show that the probability of this relatively complicated event can be approximated with a calculation that is in the spirit of a conventional value-at-risk calculation.

For simplicity, we will suppose that financial distress costs are proportional to the size of the loss, with a proportionality constant that is monotone with respect to the length of the time interval during which the loss occurs. That is, a loss of ℓ incurred over an interval of time of length Δ generates a financial distress cost of $G(\Delta)\ell$, for a given non-negative, monotone non-increasing function $G(\,\cdot\,)$. For tractability, we also suppose that $G(\,\cdot\,)$ is bounded and Hölder continuous.

The loss on position i occurs at a random time T_i. We let $Z_i(t) = 0$ for $t < T_i$, whereas $Z_i(t) = 1$ for $t \geq T_i > 0$. We assume that, given Y, the times $\{T_i\}$ of loss events are independent of each other and of the loss exposures U_i, where the distribution of T_i is the same for any position i of a given type α. We further assume that, for each outcome y of Y and for each type α, the distribution function $F_\alpha(t \,|\, y) = \mathbf{P}(T_\alpha \leq t \,|\, Y = y)$ is sub-additive in t. That is, whenever $0 \leq t \leq 1 - \Delta \leq 1$,

$$F_\alpha(t + \Delta \,|\, y) - F_\alpha(t \,|\, y) \leq F_\alpha(\Delta \,|\, y)\,. \tag{12}$$

This condition is satisfied whenever each loss time T_i has a monotone non-increasing density, as for example when T_α is distributed exponentially with parameter $\eta_\alpha(Y)$, or uniformly.

Now, the total loss on the portfolio by time t is $L_n(t) = \sum_{i=1}^n Z_i(t) U_i$. Looking at a fixed planning period, say $[0, 1]$, we wish to approximate

$$\mathbf{P}_n(Y) = \mathbf{P}\left(\sup_{0 \leq t \leq 1-\Delta \leq 1} G(\Delta)(L_n(t+\Delta) - L_n(t)) > nx \;\middle|\; Y \right), \tag{13}$$

for large n. We will eventually show, under conditions, that there is a key event time horizon Δ^* that allows us to approximate $\mathbf{P}_n(Y)$ with

$$\mathbf{P}_n(Y, \Delta^*) = \mathbf{P}\left(G(\Delta^*) L_n(\Delta^*) > nx \;\middle|\; Y \right), \tag{14}$$

which is in the spirit of a value-at-risk calculation, at a time horizon Δ^* that depends on the portfolio. (We emphasize, however, that Δ^* also depends on Y and x.)

Moreover, the distribution of the individual position losses conditional on the large portfolio loss can be similarly approximated based on the key event horizon time Δ^*. That is, given Y and the actual event of concern,

$$\mathcal{B} = \left\{ \sup_{0 \le t \le 1-\Delta \le 1} G(\Delta)(L_n(t+\Delta) - L_n(t)) > nx \right\},$$

the conditional distribution of the individual losses and their times, (U_i, T_i), can be approximated by conditioning instead on Y and the simpler event

$$\mathcal{A}(\Delta^*) = \{G(\Delta^*) L_n(\Delta^*) > nx\}.$$

Given Y and $\mathcal{A}(\Delta^*)$, the individual losses and their times, (U_i, T_i), remain independent and of identical law within each type α. The conditional default probability per loan, $\mathbf{P}(T_\alpha \le \Delta^* \,|\, Y = y, \mathcal{A}(\Delta^*))$, is given by $\widehat{\delta}_\alpha(\Delta^*, Y)$, corresponding as with (9), to the unconditional probability $\delta_\alpha(\Delta^*, y)$. The law of T_i conditional upon $\mathcal{A}(\Delta^*)$ and $\{T_i \le \Delta^*, Y = y\}$ is the same as the law of T_i conditional only upon $\{T_i \le \Delta^*, Y = y\}$. For i of type α, the law of U_i conditional upon $\mathcal{A}(\Delta^*)$ and $\{T_i \le \Delta^*, Y = y\}$ is now $M_\alpha(s)^{-1} \exp(su)\, dQ_\alpha(u)$, for s as in (9). In contrast, the law of (T_i, U_i) conditional upon $\mathcal{A}(\Delta^*)$ and $\{T_i > \Delta^*, Y = y\}$ is the same as the law of (T_i, U_i) conditional only upon $\{T_i > \Delta^*, Y = y\}$, that is, T_i and U_i are independent, the law of U_i is unchanged, and the law of T_i is simply $\frac{F_\alpha(t|Y=y) - F_\alpha(\Delta^*|Y=y)}{1 - F_\alpha(\Delta^*|Y=y)}$. The above statements are to be understood also in form of the "Gibbs conditioning principle" as in [1, Sections 7.3.1 and 7.3.3].

Now, we provide the arguments and additional technical conditions that allow us to replace the more complicated severe-loss event $\mathcal{B}$ with the simpler event $\mathcal{A}(\Delta^*)$, and define the key event horizon time Δ^*.

With each choice of Δ we associate $\delta_\alpha(\Delta, Y) = \mathbf{P}(T_\alpha \le \Delta \,|\, Y)$ and the corresponding $L(s \,|\, \Delta, Y)$ and $x_\Delta(Y)$, as in (1) and (4), respectively. For example, with T_i exponential as above, $\delta_\alpha(\Delta, Y) = 1 - \exp(-\eta_\alpha(Y)\Delta)$. We then have the (rough) LD approximation,

$$\log \mathbf{P}_n(Y) = -n \inf_{0 \le \Delta \le 1} I\left(\frac{x}{G(\Delta)} \,\middle|\, \Delta, Y \right) + \mathrm{O}(\log n), \tag{15}$$

where, for $x_* > z > x_\Delta(Y)$, we set $I(z \,|\, \Delta, Y) = sL'(s \,|\, \Delta, Y) - L(s \,|\, \Delta, Y)$ for the unique $s = s(\Delta, Y) > 0$ such that $L'(s \,|\, \Delta, Y) = z$, while $I(z \,|\, \Delta, Y) = \infty$ if $z \ge x_*$, and $I(z \,|\, \Delta, Y) = 0$ if $z \le x_\Delta(Y)$. The LD approximation (15) is relevant when

$$x > \sup_{0 \le \Delta \le 1} G(\Delta) x_\Delta(Y), \tag{16}$$

while otherwise $\mathbf{P}_n(Y)$ is bounded away from zero, uniformly in n. Indeed, the lower bound on $\mathbf{P}_n$ is obtained by fixing $t = 0$ and then applying (6) and (7) for a $\Delta \in [0, 1]$ that minimizes $I(x/G(\Delta) \,|\, \Delta, Y)$. The corresponding upper bound on $\mathbf{P}_n$ requires us to exchange in (13) the supremum over (t, Δ) and the computation of the conditional probability. To this end, by monotonicity of $G(\,\cdot\,)$ and $L_n(\,\cdot\,)$, we upper bound $\mathbf{P}_n(Y)$ while restricting the supremum in (13) to some t and $\Delta \geq 2\theta$ that are integer multiples of some fixed $\theta > 0$, provided we change there $G(\Delta)$ to $G(\Delta - 2\theta)$. Since $G(\,\cdot\,)$ is Hölder continuous, with $r > 0$ sufficiently large and $\theta = \theta(n) = n^{-r}$, we have, for all n large enough,

$$\sup_{2\theta \leq \Delta \leq 1} \{G(\Delta - 2\theta) - G(\Delta)\} \leq n^{-1}.$$

Consequently, after some algebra,

$$\begin{aligned} \mathbf{P}_n(Y) \;\; &\leq \;\; \theta^{-2} \sup_{0 \leq t \leq 1 - \Delta \leq 1} \mathbf{P}\left(L_n(t + \Delta) - L_n(t) > n\frac{x}{G(\Delta)} - \log n \;\middle|\; Y \right) \\ &\qquad + \theta^{-2} \, \mathbf{P}(G(0)L_n(1) > n \log n \,|\, Y). \end{aligned}$$

(In the above, the factor θ^{-2} which is the number of possible (t, Δ)-pairs to consider, results from the union bound.) Applying Chebyshev's inequality (of the form used in deriving [1, (2.2.12)]), we see that the second term in the above is of order $O(\exp(-cn \log n))$ for some $c > 0$, hence negligible. Condition (12) implies that $\mathbf{P}(t < T_i \leq t + \Delta \,|\, Y)$ is maximal for each i at $t = 0$, hence it suffices to consider only $t = 0$ (with $L_n(0) = 0$). To complete the derivation, recall that for any z and Δ, by Chebyshev's inequality,

$$\mathbf{P}(L_n(\Delta) > nz - \log n \,|\, Y) \leq n^s \exp(-nI(z \,|\, \Delta, Y)),$$

with $s = s(\Delta, Y)$. (See [1, (2.2.12)] for a similar bound, but without the $-\log n$ correction to nz.)

Fix an outcome y of Y for which (16) holds and suppose the minimum of $I(x/G(\Delta) \,|\, \Delta, y)$ is obtained at a unique $0 \leq \Delta^* \leq 1$. Set $s = s(\Delta^*, y) > 0$ as above. Then, for large n, the "most likely way" for the effectively excessive loss to happen is during the time interval $[0, \Delta^*]$ (assuming also that in (12) we have strict inequality for $\Delta = \Delta^*$, for at least one type α per $t > 0$). That is, under the above conditions, conditioning on Y and the actual event $\mathcal{B}$ of concern is equivalent, in the sense of the Gibbs conditioning principle already mentioned, to conditioning on Y and the simpler fixed-horizon loss event $\mathcal{A}(\Delta^*)$.

Continuing our illustrative example, we ignore for simplicity the growth macro-economic state (now taking $P(Y = b) = 1$), and suppose exponential default arrival times, with parameters chosen so as to match the one-year default probabilities of the original example, which is otherwise preserved. We take the proportional financial distress cost factor to be $G(\Delta) = e^{0.9(1-\Delta)}$, so that a given loss that occurs instantaneously generates a multiple of $e^{0.9}$ of the distress cost that would be generated by the same size loss over one year. (This is arbitrary, and merely for the purpose of a simple illustration.)

For the average position loss x that is associated with a one-year LD-approximate excess-loss probability of $p_n(b) = 0.001$, the "key time horizon" for purposes of analyzing large financial distress costs is $\Delta^* = 0.90$ years. Defaulting high-rated positions have an associated large-loss conditional default probability over the period $[0, \Delta^*]$ of approximately $\hat{\delta}_1(\Delta^*) = 0.0023$, compared to the unconditional default probability of 0.0014. For low-rated positions, we have $\hat{\delta}_2(\Delta^*) = 0.0587$, compared to the unconditional probability of 0.0566. High-rated defaulting positions have a conditional expected exposure of $(0.01 - s(\Delta^*))^{-1} = 163.54$, compared to the unconditional expected exposure of 100. For low-rated positions; the conditional expected exposure is 10.40; the unconditional is 10.

In this sense, given the occurrence of a large financial distress cost sometime during the year, the conditional expected loss associated with α-rated positions is

$$\overline{x}_\alpha = \hat{\delta}_\alpha(\Delta^*) M_\alpha(s)^{-1} \int_0^\infty e^{su} u \, dQ_\alpha(u).$$

The dependence of $\overline{x}_1$ and $\overline{x}_2$ on portfolio diversification (n) is shown in Figure 4 for exponential Q_α, and in Figure 5 for uniform Q_α of the same respective means. For each portfolio size n, the threshold loss x_n is chosen for an estimate of 0.001 (according to $p_n(b)$) of the large-loss probability $\mathbf{P}(L_n > nx_n)$. We have in mind a range of financial institutions that have been structured and capitalized so as to sustain a loss that would occur during a given year with probability 0.001. We are interested in the contribution of individual counter-parties of each type to expected losses in the event of large loss. The information in these figures shows that, with exponential distributions of loss exposures, at our base-case parameters, as one increases the size of the portfolio beyond about $n = 1100$ names, the high-rated-high-exposure

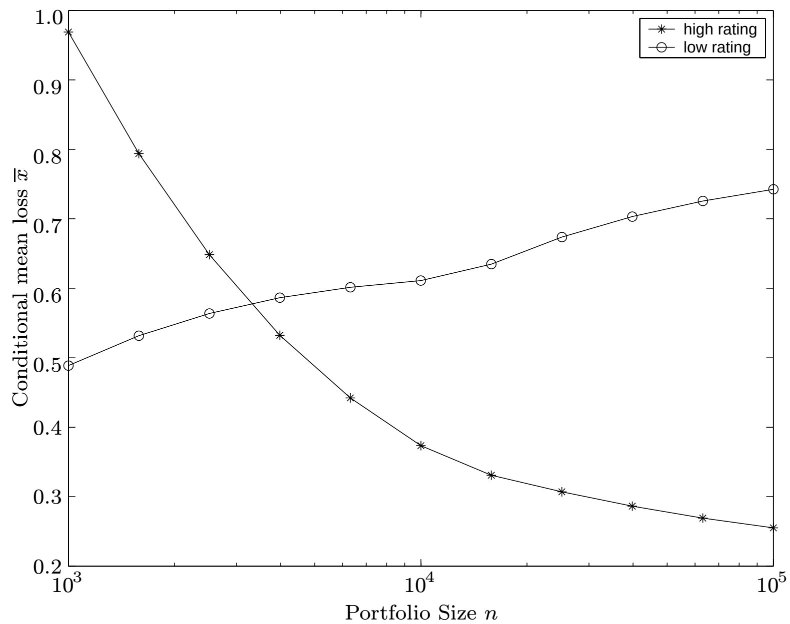

Figure 4: Conditional expected losses, with exponential exposure distribution.

counter-parties cease being the greater contributors to large losses, in this sense, and the low-rated-low-exposure counter-parties become the greater expected contributors to large losses. From the viewpoint of estimated distress losses, this information may be useful when choosing exposure limits.

References

[1] A. Dembo and O. Zeitouni, *Large Deviation Techniques and Applications: Second Edition*, Springer-Verlag, New-York, (1998).

[2] D. Duffie and K. Singleton (2003) *Credit Risk: Pricing, Measurement, and Management*, Princeton University Press.

[3] M. Gordy (2001) "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules," Working Paper, Board of Governors of the Federal Reserve System.

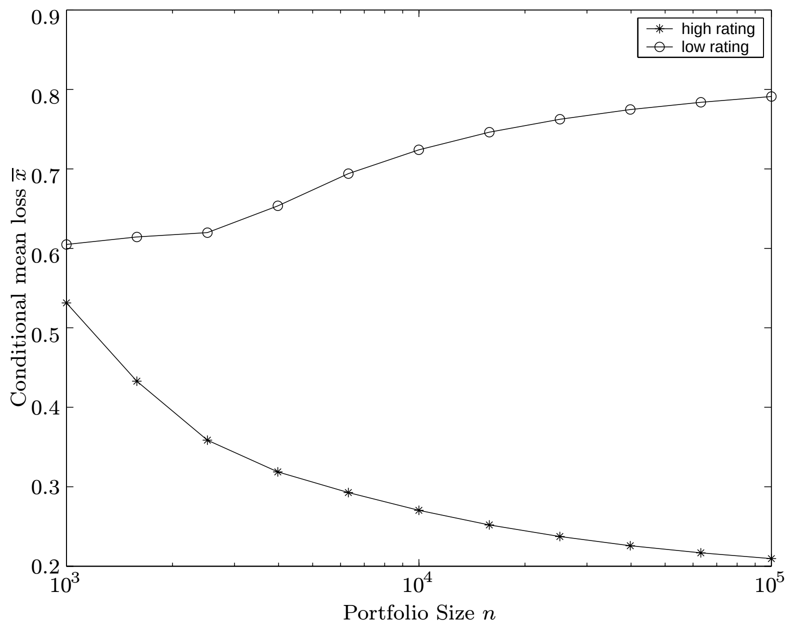

Figure 5: Conditional expected losses, with uniform exposure distribution.

[4] M. Gordy (2002) "Saddlepoint Approximation of Credit Risk," *Journal of Banking and Finance* **26**, 1335-1353.

[5] J. M. Harrison (1977) "Ruin Problems with Compounding Assets," *Stochastic Processes and Their Applications* **5**, 67-79.

[6] V. Kalashnikov and R. Norberg (1999) "Power Tailed Ruin Probabilities in the Presence of Risky Investments," forthcoming, *Stochastic Processes and Their Applications*.

[7] S. Myers and N. Majluf, (1984) "Corporate Financing and Investment When Firms Have Information Shareholders Do Not Have," *Journal of Financial Economics* **13**, 187-221.

[8] P. Nickell, W. Perraudin, and S. Varotto (2000) "Stability of Ratings Transitions," *Journal of Banking and Finance*, January.

[9] R. Norberg (1999) "Ruin Problems with Assets and Liabilities of Diffusion Type," *Stochastic Processes and Their Applications* **83**, 319-330.

[10] S. Corradin and B. Verbrigge, (2001) "Economic Risk Capital and Reinsurance: An Application to Fire Claims of an Insurance Company," *RAS, Pianificazione Redditività di Gruppo*, December.