Capital allocation for securitizations with uncertainty in loss prioritization

Michael Gordy and David Jones*
Federal Reserve Board

December 6, 2002

This paper sets forth a simple method for estimating the credit risk economic capital associated with securitization exposures, which are defined as credit exposures created by repackaging the cash flows from a pool of assets into various tranches or asset-backed securities. Our approach is motivated by the need for an effective and easily implemented regulatory capital rule for securitization exposures. Consequently, it is designed to be fully compatible with the model underpinning the Basel Committee's (2001) proposed Internal Ratings-Based Approach ("IRBA") to regulatory capital requirements against whole loans and other bank assets. Cost-effective application to a wide variety of securitizations and participating institutions dictates that our approach be parsimonious, in the sense of using minimal information on the contents of the securitized pool and on the contractual design of the securitization, as well as computationally tractable.

The role played by securitizations in unraveling the 1988 Capital Accord demonstrates the need for a regulatory capital regime that is based on an internally consistent approach to quantifying portfolio credit risk. Since 1988, securitizations have become a major funding vehicle and portfolio risk management tool for banks. Concurrently, however, banks also have learned to exploit inconsistencies within the current Accord, under which credit risks assumed through securitization transactions often entail much lower regulatory capital charges than similar risks assumed through traditional loan portfolios (see Jones 2000). Curtailing such regulatory arbitrage, while at the same time encouraging the effective hedging of credit risks through securitization and other techniques, are primary objectives behind the Basel Committee's efforts to revamp the Accord.

The model foundation for the IRBA is a special case of the class of credit-VaR risk models exemplified by CreditMetrics (Gupton, Finger and Bhatia 1997) and KMVs Portfolio Manager (Kealhofer and Bohn 2001). Economic capital is set to cover total mark-to-market credit losses over a one-year horizon with probability q.[1] It is assumed (1) that the credit portfolio is infinitely fine-grained in the sense that any single obligor represents a negligible share of the portfolio's total exposure, and (2) that a single, common systematic risk factor drives all dependence across credit

*The opinions expressed here are those of the authors, and do not reflect the views of the Board of Governors or its staff. We thank Erik Heitfield and William Perraudin for helpful comments.

[1] In this context, credit losses reflect valuation changes that result from credit quality migrations or defaults by obligors, but exclude valuation changes arising from general movements of interest rates and the market price of risk.

losses in the portfolio.

An important implication of this asymptotic single risk-factor ("ASRF") framework is that the economic capital requirement for the portfolio equals the portfolio's expected loss conditional on the systematic risk factor taking a value equal to the q^{th} percentile of its probability distribution. Given the linearity of the expectation operator, this result implies that economic capital for each instrument in the portfolio (whether that instrument is a whole loan or a tranche of a securitization) is its own expected loss conditional on the q^{th} percentile of the systematic risk factor, and thus is independent of the composition of the rest of its portfolio.[2] When applied to securitizations, the ASRF framework – and the model developed herein – implies "capital neutrality" in the sense that the sum of the economic capital charges for the individual tranches of a securitization equals the economic capital for the underlying collateral pool (denoted K_{irb}).

Importantly, our approach does not require that the securitized asset pool itself be infinitely fine-grained. Rather, a sufficient condition is that the bank's total exposure to each securitized pool (that is, through the tranches held by the bank) be small relative to the bank's overall portfolio. Thus, our model can be applied to securitizations of pools ranging from a single loan to infinitely-many loans.

This paper's main innovation is adapting the ARSF framework to permit the economic capital for an individual securitization tranche to be estimated using a relatively simple closed-form expression and parsimonious set of inputs. Computational ease and informational parsimony are especially important practical considerations when attempting to develop a cost-effective regulatory capital treatment for securitizations. The distribution of payouts to participants in a securization (often termed the cash-flow "waterfall") can be quite complex and deal-specific, depending, for example, on the time-profiles of the pool's defaults, recoveries, and principal and interest payments on the underlying loans. For regulatory capital purposes, it is not practical to attempt to account for the myriad of possible deal-specific attributes of the waterfall.

Pykhtin and Dev (2002, forthcoming) propose to cut through these complexities by assuming that, for a particular tranche, the waterfall can be summarized in terms of the tranche's par value or thickness (T) and its credit enhancement level (ζ), defined as the sum of the par values of all more-junior tranches. In practice, such information is readily available to market participants. Pykhtin and Dev also assume that economic losses experienced by the pool over the model horizon are allocated deterministically according to a strict loss prioritization ("SLP") rule, that is, the tranche absorbs pool losses only in excess of ζ, up to a maximum of T. However, when embedded in the ASRF framework, these assumptions imply an unsatisfactory knife-edge property: for an infinitely fine-grained pool, a tranche's economic capital requirement is dollar-for-dollar (100%) if $\zeta + T$ is less than or equal to K_{irb}, and zero if ζ exceeds K_{irb}. Pykhtin and Dev circumvent this problem by assuming pool losses are driven by a systematic risk factor that is correlated imperfectly with the dominant risk factor driving losses on the remainder of the bank's portfolio.

[2]See Gordy (2002) for a derivation of this result under minimal restrictions on the portfolio and very general modeling assumptions.

In contrast, our model retains the basic ASRF setup for characterizing the pool's loss distribution, while assuming that the pool's economic losses over the horizon are distributed among tranches according to a generalized version of SLP that is subject to random errors (referred to as the Uncertainty in Loss Prioritization, or ULP, model). A divergence from the strict prioritization of economic credit losses over the analysis horizon can arise from at least two sources. First, few securitizations actually call for strict prioritization of all cash flows, as subordinated tranches typically are entitled to some cash payouts prior to more-senior investors being paid out in full. Second, even with strict prioritization of cash flows over the life of a securitization, the credit enhancement level ζ generally understates the ability of more-junior tranches to absorb economic losses to the extent their contractual yield is higher than the rate of interest on the underlying loans in the pool.[3]

It should be emphasized that we are *not* suggesting that there is operational or legal risk in the execution of securitization contracts. The new source of uncertainty introduced in this paper instead reflects the potential gap between the accounting representation of the tranche (i.e., its position and thickness relative to other holders of principal) and its vulnerability to economic loss. We draw our intuition from the long vein of econometrics literature on models with hidden parameters.[4] The details of the contractual cash flow waterfall are material but unobservable parameters in the "true" model of the securitization. From the perspective of the econometrician (in our case, the regulator), such parameters act as sources of random error that must be "integrated out" rather than ignored. Were one to have unimpeded access to all details of the securitization contract, a "full information" model such as Duffie and Garleanu (2001) would naturally be preferred.

For a homogeneous pool of one-year loans (equivalent to the simple default-mode structure used by Pykhtin and Dev), the ULP model implies that a tranche's economic capital is closely approximated by a function of six inputs: the economic capital for the pool as a whole (K_{irb}); the number of loans in the pool (n); the expected loss-rate-given-default for these loans (LGD); the tranche's nominal credit enhancement level (ζ) and thickness (T); and a parameter τ that represents the magnitude of uncertainty in loss prioritization. For regulatory capital purposes, the first five parameters would be supplied by a bank, while the τ parameter would be set by supervisors. In addition to its regulatory capital applications, the model also might be used by banks for internal economic capital assessment in situations where an institution's aggregate securitization exposures represent a relatively small fraction of its overall credit portfolio and it would not be cost-effective to develop highly customized models to handle the specific details of each transactions.

Section 1 develops the ULP model under very general assumptions similar to those in Gordy (2002). A general methodology for simplifying computation of capital charges is set forth in Section

[3]Consider a homogeneous \$100 pool of 8% one-year loans. Suppose the most-junior tranche has an initial value of \$20 and pays 20%, while the \$80 senior tranche pays 5%. From the perspective of the senior tranche, $\zeta = 20$ and $T = 80$. However, suppose that 22% of the loans default, implying a total cash flow available for distribution of \$84.24. In this case, the senior tranche still would be paid in full even though the pool's loss exceeds ζ. Excess spread accounts, which often are seen in securitizations of credit card receivables, create a similar effect.

[4]We also take inspiration from the literature on the potential misalignment of accounting with economic measures (e.g., Fisher and McGowan 1983).

2. The model specification is completed in Section 3, where we impose the specific functional forms and distributional assumptions to which the IRBA is calibrated. This ensures full consistency between the IRBA and our proposed treatment of securitization exposures. We also test the robustness of our simplified computational method, and find it highly accurate across a very broad range of possible pool characteristics. We conclude in Section 4 with discussion of maturity effects and other aspects of application in practice.

1 ULP model

As under the treatment of whole loans, we seek to have capital sufficient to cover credit loss up to some percentile q of the loss distribution. We assume that both the loan originator and tranche investor have asymptotically fine-grained portfolios even when exposure to the securitized borrowers is excluded. The pool itself need not be assumed to be asymptotically fine-grained, but we do require that the tranches held by originator and investor represent trivial shares of their respective total portfolios. We also assume a single systematic risk-factor X. Let L be book-value (or "default-mode") losses incurred within the pool as a share of total pool exposure, and let $H_q(\cdot)$ be the cdf of L conditional on $X = x_q$. We need make only very weak restrictions on the underlying model of portfolio risk. For example, it could make any of a wide variety of distributional assumptions on X and on recovery risk. The technical requirements are those set forth in Gordy (2002).

The pool is securitized into a set of prioritized tranches $1, \ldots, m$, where tranche 1 is most junior, tranche 2 is next most junior, and so on. Let $S = (S_1, \ldots, S_m)$ denote the ownership shares of the tranches (so summing to one). In the standard treatment of securitizations, we take S as specified by contract and fixed ex-ante. In our approach, we recognize that securitization structures are highly complex, and that economic notions of exposure share and priority may not fully align with the legal notions. Therefore, we treat S as a random vector. A natural choice of multivariate distribution for S is the Dirichlet distribution. Let $\omega_1, \ldots, \omega_m$ be the notional stakes of the tranches. We assume that the share vector S is distributed Dirichlet with parameters $(\tau\omega_1, \ldots, \tau\omega_m)$, where $\tau > 0$ is a chosen precision parameter, which implies that the expected value of share S_j is simply ω_j. By making the Dirichlet assumption, we permit the realized share vector to "wiggle" around the expected value $(\omega_1, \ldots, \omega_m)$ and yet still always sum to one. As $\tau \to \infty$, the wiggle room disappears, and the distribution of S becomes degenerate at $(\omega_1, \ldots, \omega_m)$.

Let $G_q(z)$ be the cumulative capital charge (as a share of total pool exposure) on the juniormost share z of the structure. Under the ASRF assumptions, this is given by

$$G_q(z) = \int_0^1 \min\{z, L\}\, dH_q(L) = z - \int_0^z H_q(L) dL \tag{1}$$

where the second equality follows from integration-by-parts. When tranche shares are taken as fixed, this function is sufficient to characterize the allocation of capital across tranches. For notational convenience, let $Z = (Z_1, \ldots, Z_m)$ be the cumulative shares defined by $Z_j = \sum_{i \leq j} S_i$, and let

$\zeta = (\zeta_1, \dots, \zeta_m)$ be the cumulative expected shares defined by $\zeta_j = \sum_{i \le j} \omega_i$. In the case of strict loss prioritization ("SLP"), the capital on tranche j is given by

$$K_j = G_q(Z_j) - G_q(Z_{j-1}) = G_q(\zeta_j) - G_q(\zeta_{j-1}). \tag{2}$$

When S is stochastic and independent of all other risk-factors in the portfolio, the appropriate capital charge is[5]

$$K_j = \mathrm{E}\left[G_q(Z_j)\right] - \mathrm{E}\left[G_q(Z_{j-1})\right]. \tag{3}$$

The linear nature of the expectation operator is extremely convenient here, because it implies that we need only worry about the marginal distributions of the Z_j and not the entire joint distribution. The marginal distribution of Z_j is Beta with parameters $(\tau\zeta_j, \tau(1-\zeta_j))$, which implies that

$$\begin{aligned}
\mathrm{E}\left[G_q(Z_j)\right] &= \mathrm{E}\left[Z_j\right] - \mathrm{E}\left[\int_0^z H_q(L)dL\right] \\
&= \zeta_j - \int_0^1 \frac{z^{\tau\zeta_j - 1}(1-z)^{\tau(1-\zeta_j)-1}}{B(\tau\zeta_j, \tau(1-\zeta_j))} dz \int_0^z H_q(L)dL \\
&= \zeta_j - \int_0^1 (1 - B(z; \tau\zeta_j, \tau(1-\zeta_j)))H_q(z)dz
\end{aligned}$$

where the function $B(y; a, b)$ is the Beta(a, b) cdf evaluated at y.[6] As this expression does not depend on any tranche division point other than ζ_j, we can write the cumulative capital function as a smooth function of the cumulative nominal share ζ:

$$K(\zeta) \equiv \mathrm{E}\left[G_q(Z) | \mathrm{E}\left[Z\right] = \zeta\right] = \zeta - \int_0^1 (1 - B(z; \tau\zeta, \tau(1-\zeta)))H_q(z)dz \tag{4}$$

Intuitively, the stochastic version of the model smooths over all possible share values Z in the SLP model. Equation (4) embeds the SLP rule of equation (1) as a limiting case. As $\tau \to \infty$, the beta distribution for Z becomes degenerate at ζ, so

$$\lim_{\tau\to\infty} \int_0^1 (1 - B(z; \tau\zeta, \tau(1-\zeta)))H_q(z)dz = \int_0^1 1_{z<\zeta} H_q(z)dz = \int_0^\zeta H_q(z)dz.$$

If the securitized pool is itself asymptotically fine-grained, as is a reasonable characterization of most retail securitizations, then equation (4) has a simple analytic solution. In the asymptotic

[5] Here is where we need the assumption that the tranche (or set of tranches of securitizations of the same underlying pool) constitutes a trivial share of the super-portfolio. If this is not the case, the idiosyncratic risk in S would not be diversified away, and thus would demand capital of its own.

[6] The final expression is obtained using integration by parts, where the "du" part is the beta pdf and the "v" part is the integral over H_q.

case, $H_q(z) = 1_{z \geq \mathrm{E}[L|x_q]}$, so

$$\int_0^\zeta H_q(z)dz = \max\{0, \zeta - \mathrm{E}\left[L|x_q\right]\}.$$

Integrating over the distribution of Z for any fixed ζ, we find

$$K(\zeta) = \zeta \cdot B(\mathrm{E}\left[L|x_q\right]; \tau\zeta + 1, \tau(1-\zeta)) + \mathrm{E}\left[L|x_q\right](1 - B(\mathrm{E}\left[L|x_q\right]; \tau\zeta, \tau(1-\zeta))). \quad (5)$$

Note that the expression $\mathrm{E}\left[L|x_q\right]$ is K_{irb} as specified in the New Basel Accord (Basel Committee on Bank Supervision 2001).

As under the SLP case, total required capital sums to K_{irb} across the tranches, that is, $K(1) = \mathrm{E}\left[L|x_q\right]$. In the SLP case, any tranche of a securitization of an asymptotically fine-grained pool that is senior to the K_{irb} threshold requires zero capital, and any tranche of such a securitization that is strictly junior to the K_{irb} threshold requires dollar-for-dollar capital. In this model, such senior tranches always require some capital because of the possibility that the "realized" exposure of the senior tranches exceeds $1 - K_{\mathrm{irb}}$, and such junior tranches require less than dollar-for-dollar capital because of the possibility that their realized exposure is less than K_{irb}. Thus, the extended model unambiguously increases capital for tranches senior to K_{irb}, and unambiguously reduces capital for tranches junior to K_{irb}. For tranches that straddle this breakpoint, the effect is ambiguous but typically small.

Observe that τ is the only additional parameter added by this extension to the standard model. The effect of τ is best understood by examining equation (5). As $\tau \to \infty$, uncertainty in the division of risk vanishes, and we find

$$\lim_{\tau \to \infty} K(\zeta) = \zeta 1_{\mathrm{E}[L|x_q] \geq \zeta} + \mathrm{E}\left[L|x_q\right](1 - 1_{\mathrm{E}[L|x_q] \geq \zeta}) = \min\{\zeta, \mathrm{E}\left[L|x_q\right]\} = G_q(\zeta),$$

which is the result of the SLP model for an asymptotic portfolio. As $\tau \to 0$, the contractual exposure shares become less and less informative of the actual division of risk, and we find

$$\lim_{\tau \to 0} K(\zeta) = \mathrm{E}\left[L|x_q\right]\zeta,$$

which implies a proportional sharing of K_{irb} across the tranches (i.e., the tranches are treated as *pari passu*). Intermediate values of τ correspond to greater or lesser degrees of smoothing between these extremes.

For less fine-grained pools, uncertainty in Z has a smaller effect. Figure 1 shows the effect of τ and n using the model specification set forth in Section 3. The case of $n = \infty$ is shown in the bottom panels. When $\tau = \infty$, marginal capital is dollar-for-dollar up to K_{irb} and then zero thereafter. Setting $\tau = 1000$ provides a modest degree of smoothing, and much lower values of τ a much greater degree of smoothing. For a portfolio of $n = 128$ (middle row of panels), which is representative of the concentration typically seen in CDO pools, undiversified idiosyncratic risk is sufficient to smooth away the cliff effect at K_{irb}. In this case, the capital and marginal capital curves for $\tau = 1000$ are indistinguishable from those of $\tau = \infty$. Uncertainty in Z has no material effect on capital unless τ is below, say, 100. When $n = 4$ (upper row of panels), idiosyncratic risk within the pool has a dominant effect on the distribution of losses across tranches, and τ must be extremely low (on the order of 10) for uncertainty in Z to have any additional smoothing effect.

Figure 1: Effect of τ on capital and marginal capital

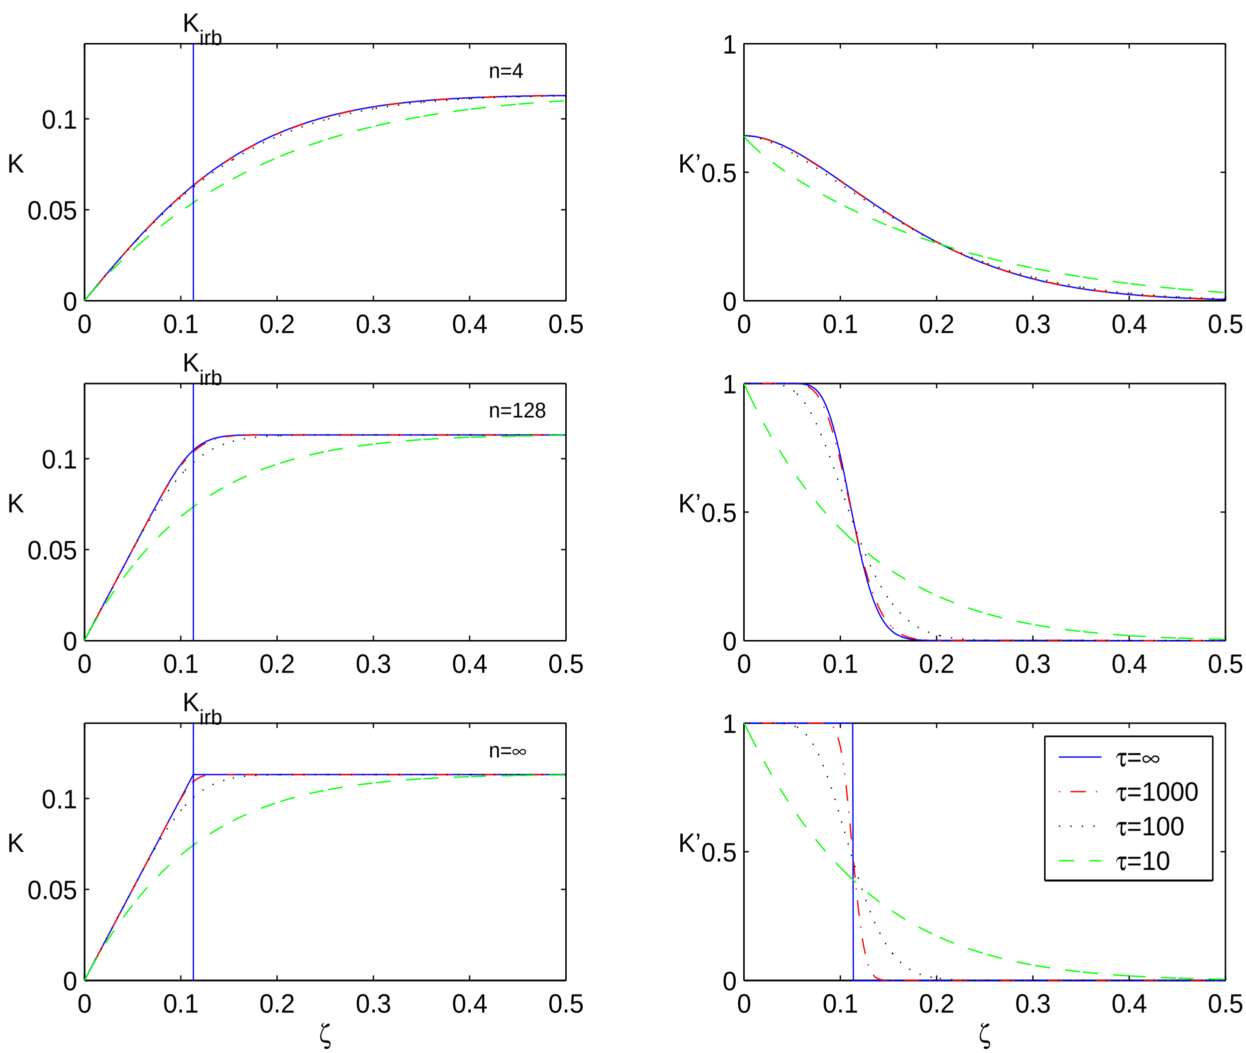

Note: Model specified as in Section 3 with pool parameters PD = 0.02, ELGD = 0.5 and $\rho = 0.2$. We set VaR target quantile $q = 0.999$ and recovery risk parameter $\gamma = 0.25$.

2 Fitting a simple functional form to $K(\zeta)$

Unless we restrict ourselves to analysis of securitizations of asymptotically fine-grained pools, the conditional loss distribution $H_q(z)$ is likely to be analytically intractable, and the solution for $K(\zeta)$ in equation (4) requires numerical integration or simulation.[7] For regulatory purposes, a simpler and more transparent functional solution is required, even if it comes at slight expense of precision.

We define the *fitting function* $F(\cdot)$ by

$$F(\zeta) = 1 - \frac{K'(\zeta)}{K'(0)}. \tag{6}$$

This definition is useful because lets us exploit three known properties of the first derivative of $K(\zeta)$: $K'(\zeta)$ is nonincreasing on the unit interval, and we have $K'(1) = 0$ and $K'(0) = 1 - H_q(0)$.[8] From these properties, we see that $F(\zeta)$ is nondecreasing on the unit interval and that $F(0) = 0$ and $F(1) = 1$. Thus, F behaves like a cumulative distribution function for a random variable with support on the unit interval. Although this cdf is typically of intractable form, we might expect that it can be closely approximation by the cdf of a simple distribution such as the beta. In this section, we derive the mean (μ) and standard deviation (σ) of F.

To get the mean parameter, we rearrange equation (6) as $K'(\zeta) = K'(0)(1 - F(\zeta))$, and integrate to get

$$K(\zeta) = K'(0) \int_0^\zeta (1 - F(y))dy.$$

At $\zeta = 1$, we can integrate by parts to get

$$\int_0^1 F(y)dy = 1 - \int_0^1 yf(y)dy = 1 - \mu,$$

which implies that

$$\mu = \frac{K(1)}{K'(0)} = \frac{\mathrm{E}\left[L|x_q\right]}{1 - H_q(0)}. \tag{7}$$

The variance of F is more challenging. By definition,

$$\sigma^2 = \int_0^1 y^2 f(y)dy - \mu^2 = \frac{-1}{K'(0)} \int_0^1 y^2 K''(y)dy - \mu^2 \tag{8}$$

[7] Monte Carlo simulation of $K(\zeta)$ is straightforward but computationally intensive. Since $K(\zeta) = \mathrm{E}\left[\min\{Z, L\}|x_q\right]$ for $Z \sim B(\tau\zeta, \tau(1-\zeta))$, we need only draw a sample of $\{z_1, \ldots, z_T\}$ for Z and a sample of $\{\ell_1, \ldots, \ell_T\}$ for L (from the H_q distribution). $K(\zeta)$ is estimated by $(1/T)\sum \min\{z_i, \ell_i\}$. A fresh sample of the z_i must be drawn for each value of ζ.

[8] To obtain the derivative of $K(\zeta)$ at $\zeta = 0$, note that the expression $(1 - B(z; \tau\zeta, \tau(1-\zeta)))$ in equation (4) behaves like a step function at ζ in the neighborhood of $\zeta = 0$. Therefore, its derivative with respect to ζ at $\zeta = 0$ is the Dirac delta function $\delta_0(z)$.

In Appendix A, we apply integration by parts to arrive at

$$\sigma^2 = \frac{1}{K'(0)}\left(1 - 2\int_0^1 \Xi_\tau(z)H_q(z)dz\right) - \mu^2 \tag{9}$$

where the function $\Xi_\tau(z)$ is defined by

$$\Xi_\tau(z) \equiv \int_0^1 B(z; \tau y, \tau(1-y))dy. \tag{10}$$

The Ξ_τ function can be understood as the unconditional cdf of a random variable Z that has conditional distribution $Z|(Y = y) \sim \text{Beta}(\tau y, \tau(1-y))$, where $Y \sim U[0,1]$.

The difficulty remains in the integral over $\Xi_\tau(z)H_q(z)$. In the special case of $\tau = \infty$ we have $\Xi_\infty(z) = z$ and the integral can be solved analytically:

$$\begin{aligned}\int_0^1 \Xi_\tau(z)H_q(z)dz &= \left[\frac{1}{2}z^2 H_q(z)dz\right._0^1 - \int_0^1 z^2 h_q(z)dz \\ &= \frac{1}{2}(1 - \mathrm{E}\left[L^2|x_q\right]) = \frac{1}{2}(1 - (\mathrm{V}[L|x_q] + \mathrm{E}\left[L|x_q\right]^2))\end{aligned}$$

where $\mathrm{V}[L|x_q]$ is the conditional variance of pool loss. In this case, equation (9) simplifies to

$$\sigma^2 = \frac{1}{K'(0)}(\mathrm{V}[L|x_q] + \mathrm{E}\left[L|x_q\right]^2) - \mu^2. \tag{11}$$

For our purposes, the limiting case of $\tau = \infty$ is not terribly interesting because it eliminates uncertainty in loss prioritization. The ULP model collapses back to the SLP model, and we again have a cliff effect at K_{irb} when $n = \infty$.[9] For large but finite values of τ, the $\Xi_\tau(z)$ function is close to z, but the difference cannot be ignored. Figure 2 shows how $\Xi_\tau(z)$ varies with τ.

For finite τ, the $\Xi_\tau(z)$ function does not have an analytical solution. Nonetheless, as shown in Appendix B, it can be closely approximated by

$$\hat{\Xi}_\tau(z) = z + \xi\left(\frac{1}{2} - z\right)\frac{(z(1-z))^{\alpha-1}}{B(\alpha,\alpha)} \tag{12}$$

where

$$\alpha = \frac{3(\tau^2 + 6\tau + 6)}{3\tau^2 + 13\tau + 18} \qquad \text{and} \qquad \xi = \frac{2\alpha + 1}{3(\tau + 1)}.$$

The approximation is exact at both $\tau = 0$ and $\tau = \infty$, and extremely precise at any τ in between.

For reasonably large values of τ, say $\tau = 100$ or larger, we have $\alpha \approx 1$ and $\xi \approx 1/\tau$. If we

[9] Note that when $n = \infty$, we have $H_q(0) = 0$, so $K'(0) = 1$, so $\mu = \mathrm{E}\left[L|x_q\right]$. By the law of large numbers, we also have $\mathrm{V}[L|x_q] = 0$. Thus, equation (11) yields $\sigma^2 = 0$.

Figure 2: Dependence of $\Xi_\tau(z)$ on τ

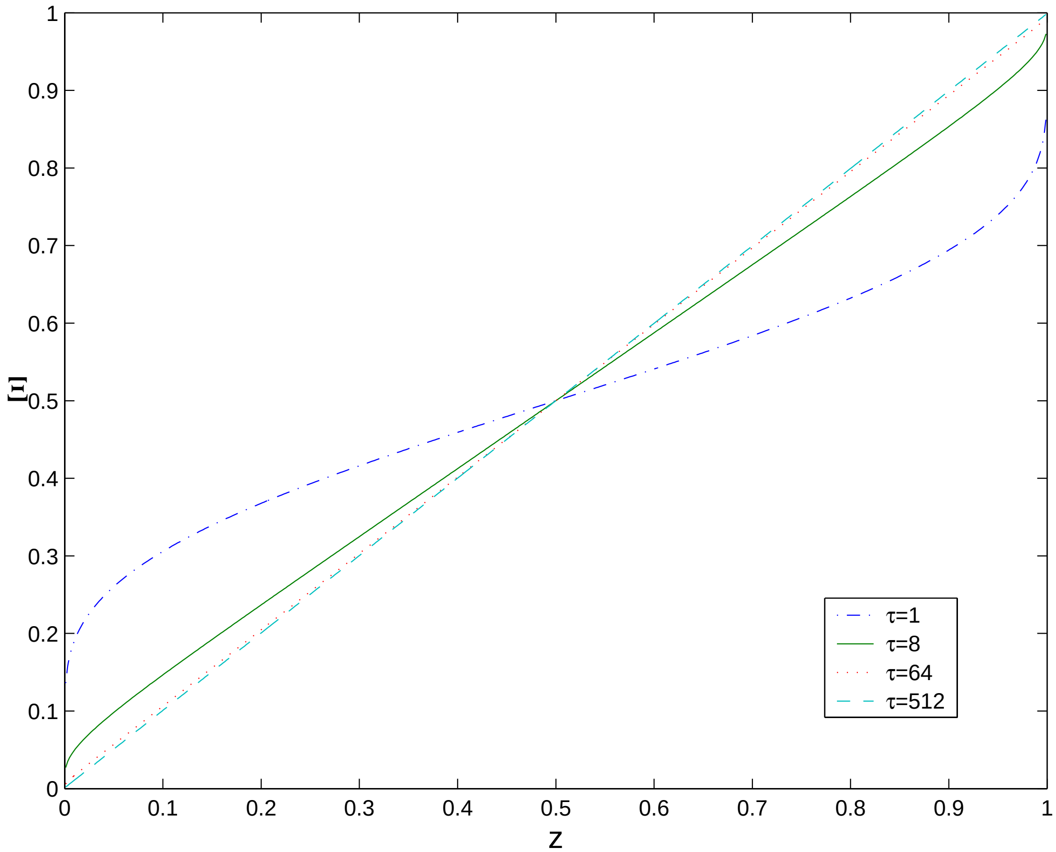

impose these approximations, equation (12) simplifies to

$$\hat{\Xi}_\tau(z) = z + \frac{1}{\tau}\left(\frac{1}{2} - z\right). \tag{13}$$

Applying integration by parts, we find

$$\begin{aligned} 2\int_0^1 \hat{\Xi}_\tau(z) H_q(z) dz &= 1 - \int_0^1 \left(z^2 + \frac{1}{\tau}(z - z^2)\right) h_q(z) dz \\ &= 1 - \left(\mathrm{V}[L|x_q] + \mathrm{E}\left[L|x_q\right]^2 + \frac{1}{\tau}(\mathrm{E}\left[L|x_q\right](1 - \mathrm{E}\left[L|x_q\right]) - \mathrm{V}[L|x_q])\right). \end{aligned}$$

This approximation leads to a closed-form solution for σ^2:

$$\sigma^2 = \frac{1}{K'(0)}\left(\mathrm{V}[L|x_q] + \mathrm{E}\left[L|x_q\right]^2\right) - \mu^2 + \frac{1}{\tau}\frac{1}{K'(0)}\left(\mathrm{E}\left[L|x_q\right](1 - \mathrm{E}\left[L|x_q\right]) - \mathrm{V}[L|x_q]\right). \tag{14}$$

The expression for σ^2 has been arranged to show that it naturally decomposes into two components: The first is the contribution of undiversified idiosyncratic risk in the underlying pool (i.e., the impact of pool granularity), and equals the formula for σ^2 when $\tau = \infty$ (equation 11). The second is the contribution of uncertainty in loss prioritization and is inversely proportional to τ.

From an operational point of view, equation (14) is no more complicated than equation (11). Thus, so long as the uncertainty in loss prioritization is fairly small (i.e., so long as τ is reasonably large), we can accurately and conveniently capture its contribution in the variance of the fitting function.

The calculations simplify further in the special case of $n = \infty$, which includes most securitizations of retail pools. When $n = \infty$, $K'(0) = 1$ and $\mathrm{V}[L|x_q] = 0$, so we have $\mu = \mathrm{E}\left[L|x_q\right]$ and $\sigma^2 = (1/\tau)\mathrm{E}\left[L|x_q\right](1 - \mathrm{E}\left[L|x_q\right])$.

3 A Complete Specification

We thus far have not needed to specify a model for portfolio loss or to choose a cumulative distribution function to assign to the fitting function. In order to arrive at an implementable formulation, we now complete our specification.

We assume that the securitized pool is homogeneous and that the conditional loss distribution H_q comes from a single-factor default-mode model with idiosyncratic recovery risk. This implies that the number of defaults in a portfolio of n loans is distributed Binomial(p_q, n), where p_q is the conditional probability (given $X = x_q$) of default for a single loan in the pool. If LGD for a single default has a continuous distribution, then H_q is continuous on unit interval support, except that there is probability mass at $L = 0$. The probability of zero loss is the probability that every borrower performs, so $H_q(0) = (1 - p_q)^n$. If LGD has mean ELGD and standard deviation VLGD, then the mean and variance of the conditional loss distribution are given by $\mathrm{E}\left[L|x_q\right] = \mathrm{ELGD} \cdot p_q$

and

$$\mathrm{V}[L|x_q] = \frac{1}{n}\left(\mathrm{ELGD}^2 p_q(1-p_q) + p_q \mathrm{VLGD}^2\right) \tag{15}$$

To retain consistency with the IRB treatment of whole loans, we adopt the CreditMetrics model of obligor dependence and LGD volatility. That is, we assume that X has standard normal distribution, that the conditional probability of default is given by

$$p_q = \Phi\left(\frac{\Phi^{-1}(\mathrm{PD}) + \Phi^{-1}(q)\sqrt{\rho}}{\sqrt{1-\rho}}\right) \tag{16}$$

where Φ is the standard normal cdf and ρ is the correlation in asset returns (see Gordy (2000) for a derivation of p_q from the CreditMetrics model). Loss rates given default are drawn as independent beta random variables. Following the convention in CreditMetrics and KMV Portfolio Manager, we assume the variance of loss given default is given by

$$\mathrm{VLGD}^2 = \gamma \cdot \mathrm{ELGD} \cdot (1 - \mathrm{ELGD}) \tag{17}$$

where γ is a parameter in $[0,1]$. Special cases include $\gamma = 0$, which corresponds to fixed LGD rates (no recovery risk), and $\gamma = 1$, which arises when LGD is distributed Bernoulli (i.e., 100% LGD with probability ELGD, zero LGD otherwise).

A variety of two-parameter distributions for approximating the fitting function would lead to a closed-form solution for the capital function. The beta distribution would be a natural choice given its unit interval support, and we have found that it provides excellent fit as well under a wide range of parameter values.[10] Let θ be the precision of F defined by

$$\theta \equiv \frac{\mu(1-\mu)}{\sigma^2} - 1.$$

The parameter θ measures the precision of F in the same manner as τ measures the precision of the distribution for tranche cutoff Z.

To distinguish between the "true" fitting function, which is of intractable form, and our approximation based on the beta cdf, let $\hat{F}$ denote the approximation. Similarly, let $\hat{K}$ denote the approximation implied by $\hat{F}$ to the true K function. The solution to $\hat{K}(\zeta)$ is given by

$$\begin{aligned}\hat{K}(\zeta) &= \int_0^\zeta \hat{K}'(y)dy = K'(0)\int_0^\zeta (1-\hat{F}(y))dy \\ &= (1-H_q(0))\cdot(\zeta\cdot(1-B(\zeta;\theta\mu,\theta(1-\mu))) + \mu\cdot B(\zeta;\theta\mu+1,\theta(1-\mu)))\,.\end{aligned} \tag{18}$$

[10] If we ignore the upper bound on loss at $\zeta = 1$, then the gamma and lognormal distributions also are reasonable choices and yield closed-form $K(\zeta)$. However, the beta provides much the best fit overall.

Note that when $n = \infty$, we have $\theta = \tau - 1$ and $\hat{K}(\zeta)$ simplifies to

$$\hat{K}(\zeta) = \zeta \cdot (1 - B(\zeta; (\tau-1)K_{\text{irb}}, (\tau-1)(1-K_{\text{irb}}))) + K_{\text{irb}} \cdot B(\zeta; (\tau-1)K_{\text{irb}} + 1, (\tau-1)(1-K_{\text{irb}})). \quad (19)$$

We have examined the robustness of our fitted $\hat{K}(\zeta)$ to the "true" $K(\zeta)$ given by equation (4). For each combination of parameters $(n, \text{PD}, \text{LGD}, \rho, \tau)$, we calculate the theoretical K function by Monte Carlo and the fitted $\hat{K}$ function by equation (18). Throughout the exercise, we fix constant the VaR target quantile q at 0.999, and set the recovery risk parameter to $\gamma = 0.25$. We then measure the relative root-mean-squared-error as

$$\text{RMSE} = \frac{1}{K_{\text{irb}}} \sqrt{\int_0^1 \left(K(\zeta) - \hat{K}(\zeta)\right)^2 d\zeta}.$$

We performed these calculations for each combination of

$$\begin{aligned}
n &\in \{1, 4, 16, 64, 256, \infty\} \\
\text{PD} &\in \{0.1, 0.2, 0.5, 1, 2, 4, 6, 10, 15\} \quad \text{(in percentage points)} \\
\text{LGD} &\in \{0.05, 0.20, 0.35, 0.5, 0.65, 0.8, 0.95\} \\
\rho &\in \{0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28, 0.32\} \quad \text{(asset correlation), and} \\
\tau &\in \{100, 200, 400, 600, 800, 1000, 1600, 3200\}.
\end{aligned}$$

In total, this exercise covered 24,192 parameter combinations.

We find that the fitted $\hat{K}$ function performs extremely well nearly everywhere in the parameter space. The only exception arises when the pool is comprised of a *single* loan to an investment grade borrower with expected LGD of 5% and asset-correlation under 12%. In this case, relative RMSE reaches as high as 10.3% of K_{irb}. For loans with very low PD and very low expected LGD, simulation noise is of large relative magnitude, so the high RMSE may be due in part to error in estimating K rather than in $\hat{K}$. Furthermore, K_{irb} is always quite small in these cases (under 0.005), so a fitting error of 10% has a negligible impact on absolute capital requirements. Perhaps most importantly, this peculiar combination of parameters does not arise under the proposed New Basel Accord. Low PD borrowers cannot be assigned low asset-correlations unless they are retail or small corporate borrowers, in which case the loans are invariably far too small to comprise an entire securitized pool.

Excluding the exceptional case, the median relative RMSE is 0.15%. That is, at the median, root-mean-squared error is less than one-sixth of one percent of K_{irb} for the parameter combination. The maximum relative RMSE is under 5.5% of K_{irb}. The performance of the fitting function is shown in Figure 3 for four very different underlying pools. K_{irb} varies across these examples by a factor of six, and n varies from one to infinity, yet in each case the fit is excellent. Indeed, in the bottom two panels, it is impossible to distinguish the two curves. Figure 4 shows the Monte Carlo and fitted marginal capital curves ($K'(\zeta)$) for the same four examples. The simulations produce somewhat jagged estimates of marginal capital, so the fitted function may indeed provide the more accurate curve.

Figure 3: Performance of fitted $\hat{K}(\zeta)$

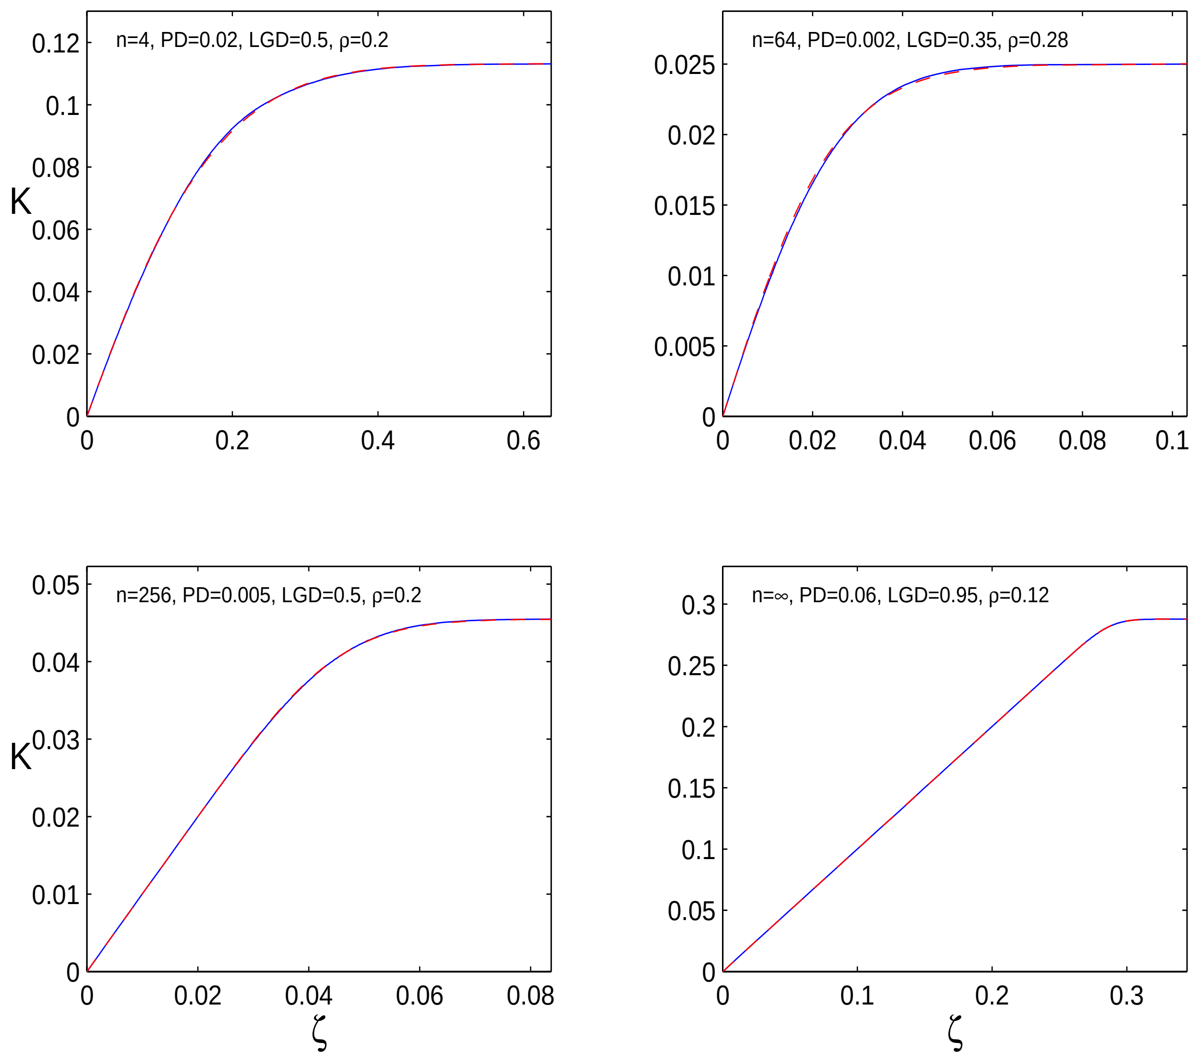

Note: Monte Carlo $K(\zeta)$ and fitted $\hat{K}(\zeta)$ are plotted with a solid line and a dashed line, respectively. In each panel, we set $\tau = 1000$.

4 Application to Regulatory Capital Treatment

The ULP capital charge implied by equation (18) on a tranche with credit enhancement level ζ and thickness T is $(\hat{K}(\zeta + T) - \hat{K}(\zeta))/T$ per dollar of tranche par value. The calculations make

Figure 4: Performance of fitted marginal capital

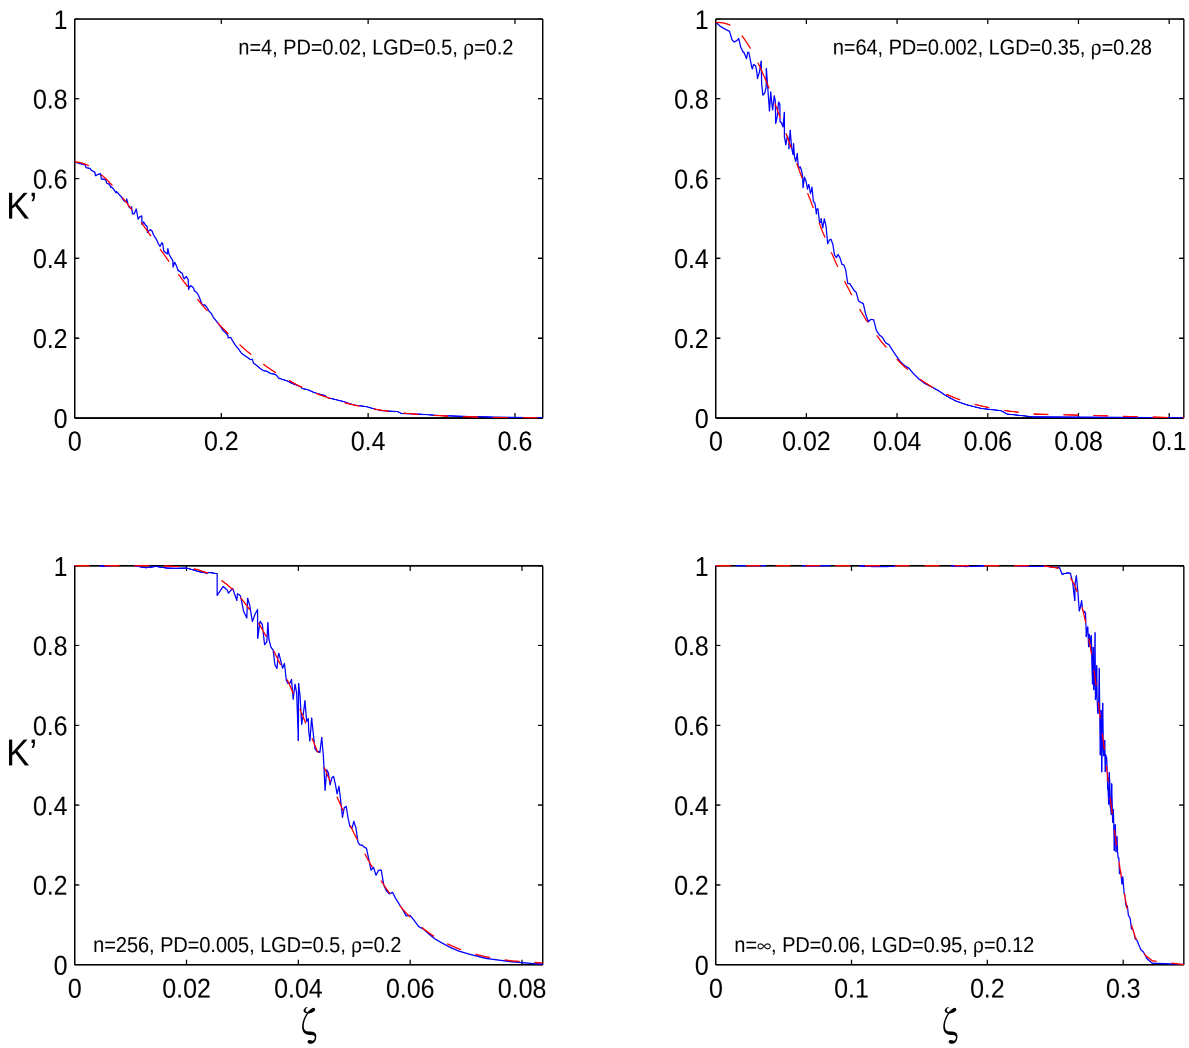

Note: Monte Carlo $K'(\zeta)$ and fitted $\hat{K}'(\zeta)$ are plotted with a solid line and a dashed line, respectively. In each panel, we set $\tau = 1000$.

use of a number of intermediate quantities, such as K_{irb}, that depend ultimately on the underlying pool parameters n, PD, ELGD and asset-correlation ρ, as well as on regulatory parameters τ and γ. Further simplification can be obtained by noting that PD and ρ enter the calculations only via p_q, and that $p_q = K_{\text{irb}}/\text{ELGD}$. Thus, a sufficient set of pool parameters is n, K_{irb} and ELGD. These are the pool-level inputs to the K function in the Basel Committee's (2002) Second Working Paper on Securitisation. In Appendix C, we show how the ULP model (equation (18) specifically) is embedded in the proposed Supervisory Formula Approach.

As a practical matter, parameterization of the ULP in terms of K_{irb} compensates for some of the limitations of the model's default-mode notion of credit loss.[11] Strictly speaking, the model assumes that the underlying assets are of one-year maturity. If the pool were of longer maturity, there would be no mechanism in the model for recognition of economic losses in the pool due to rating migrations short of default. Absence of arbitrage implies that economic losses in the pool must be equaled by the sum of economic losses to the tranches.[12] Under the IRBA, capital charges for the underlying pool incorporate maturity effects, so parameterization in terms of K_{irb} may be more robust than parameterization in terms of PD and ρ.

In a similar vein, the model also assumes that the pool contains only simple whole loans or bonds. In practice, it is now not uncommon for securitized pools to include, for example, tranches of other securitizations. Modeling the performance of such pool assets requires more than the simple CreditMetrics-based function of PD, ELGD and asset-correlation. However, K_{irb} can be obtained for a wide variety of assets, so we can avoid the hazardous task of assigning PD and asset-correlation parameters to complex asset types.[13]

By parameterizing the ULP model in terms of the K_{irb} of the underlying pool, we at least impose the appropriate total economic capital requirement on the securitization as a whole. An open question, and topic of ongoing research, is whether this convenient approach to extending the applicability of the model can lead to significant misallocation of economic capital across the tranches.

[11]Pykhtin and Dev's (2002, forthcoming) model suffers from the same limitations.

[12]Much as a decline in the asset value of a firm will be spread among shareholders and bondholders, a decline in the value of the pool will not be allocated by strict prioritization, but rather will be spread to some degree across all tranches.

[13]The problem of assigning ELGD remains, but the effect of this parameter is of secondary order when K_{irb} is held fixed. Setting ELGD to 0.5 maximizes VLGD, and so yields a conservative treatment.

A Variance of the fitting function

In order to obtain the variance σ^2 of the fitting function, we need to solve the integral $\int_0^1 y^2 K''(y)dy$. We apply integration-by-parts twice:

$$\begin{aligned}\int_0^1 y^2 K''(y)dy &= \Big[_0^1 y^2 K'(y) - 2\int_0^1 yK'(y)dy \\ &= (0-0) - 2\left(\Big[_0^1 yK(y) - \int_0^1 K(y)dy\right) \\ &= -2\left(K(1) - \int_0^1 K(y)dy\right).\end{aligned}$$

Drawing on equation (4) and again integrating by parts, we get

$$\begin{aligned}\int_0^1 K(y)dy &= \int_0^1 \left(y - \int_0^1 (1 - B(z;\tau y, \tau(1-y)))H_q(z)dz\right) \\ &= \frac{1}{2} - \int_0^1 H_q(z)dz + \int_0^1 B(z;\tau y,\tau(1-y))H_q(z)dz \\ &= \frac{1}{2} - (1 - K(1)) + \int_0^1 \left(\int_0^1 B(z;\tau y,\tau(1-y))dy\right) H_q(z)dz \\ &= K(1) - \left(\frac{1}{2} - \int_0^1 \Xi_\tau(z)H_q(z)dz\right)\end{aligned}$$

We combine these results are substitute into equation (8) to get equation (9).

B Approximation of $\Xi_\tau(z)$

In order to obtain a reasonably simple expression for the variance of the fitting function, we need to be able to provide a simple approximation to the function $\Xi_\tau(z)$ defined by

$$\Xi_\tau(z) \equiv \int_0^1 B(z;\tau y,\tau(1-y))dy. \tag{20}$$

This function itself does not have a tractable analytical solution, except at $\tau = 0$ and $\tau = \infty$. For all $z \in (0,1)$ and $y \in (0,1)$, we have $\lim_{\tau\to 0} B(z;\tau y,\tau(1-y)) = \frac{1}{2}$, so $\Xi_0(z) = 1/2$. We also have $\lim_{\tau\to\infty} B(z;\tau y,\tau(1-y)) = 1_{z\geq y}$, which implies $\Xi_\infty(z) = z$. It is desirable that our approximation to Ξ_τ take on the same limiting forms.

For positive finite τ, $\Xi_\tau(z)$ weaves around the 45° line in a regular symmetric pattern. As shown in Figure 5, the function $\Xi_\tau(z) - z$ starts at zero, rises sharply, levels off quickly, then becomes linear with negative slope and hits zero at $z = 1/2$. The pattern above $z = 1/2$ is the mirror image of the pattern below $z = 1/2$; i.e., the function displays rotational symmetry around $z = 1/2$.

A simple function that displays the same properties is $(1/2 - z)(z(1-z))^{(\alpha-1)}$ for $\alpha \geq 1$.

Figure 5: Cyclical component of $\Xi_\tau(z)$

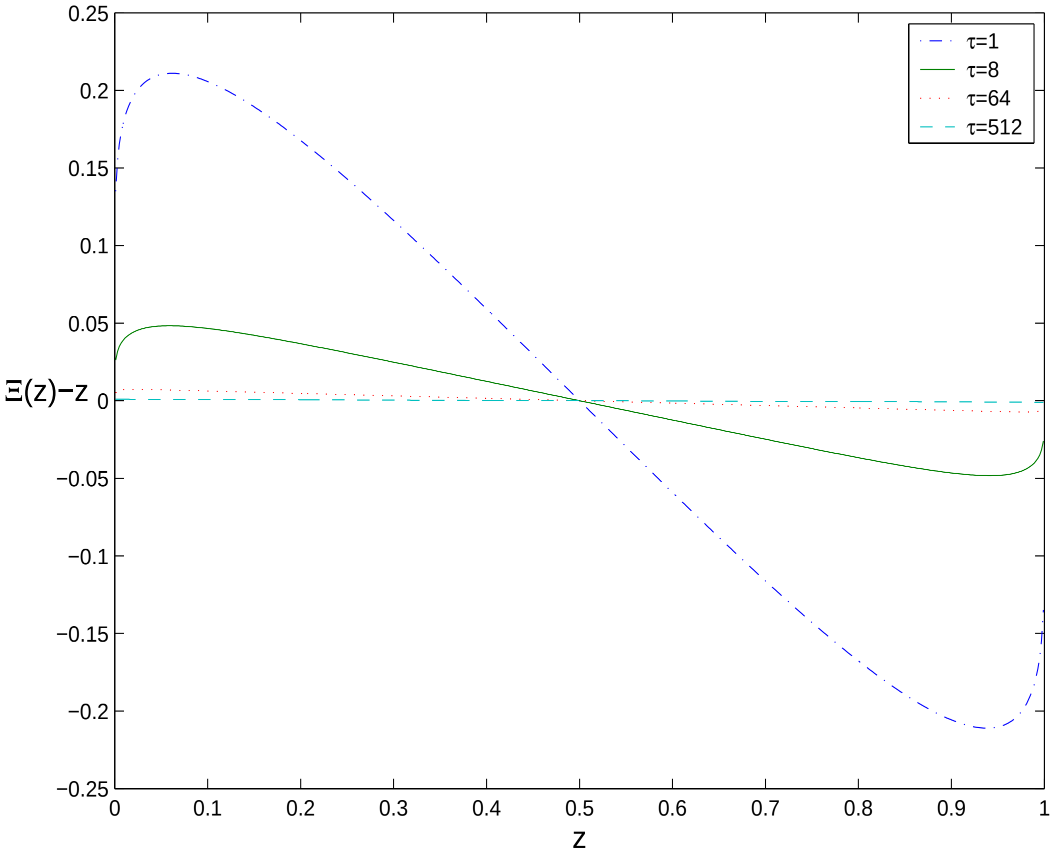

Therefore, we propose to approximate $\Xi_\tau(z)$ by

$$\hat{\Xi}_\tau(z) \equiv z + \xi \left(\frac{1}{2} - z\right) \frac{(z(1-z))^{\alpha-1}}{B(\alpha,\alpha)}, \tag{21}$$

where the coefficients ξ and α are functions of τ.[14] As noted in Section 2, $\Xi_\tau(z)$ can be treated as a cdf of a random variable on the unit interval. The approximation $\hat{\Xi}_\tau(z)$ equals zero at $z=0$, one at $z=1$, and is increasing in between, so also can be treated as a cdf on the unit interval. In each case, the "mean" is 1/2. We solve for coefficients ξ and α such that the second and fourth moments of these distributions match. (Due to the rotational symmetry of Ξ and $\hat{\Xi}$, the third moments add no new information.)

As Ξ_τ is a compounded beta distribution, its moments are easily obtained. The j^{th} uncentered moment is given by

$$\lambda_j \equiv \frac{1}{(\tau)_j} \int_0^1 (\tau y)_j dy$$

where $(a)_k$ is Pochhammer's notation, i.e., $(a)_0 = 1$, $(a)_1 = a$, $(a)_k = (a)_{k-1}(a+k-1)$. The function $(\tau y)_j$ is merely a j^{th} order polynomial in y, so λ_j has a simple closed form solution for any j.

The corresponding moments for $\hat{\Xi}_\tau$ also have closed-form solution:

$$\hat{\lambda}_j \equiv \frac{1}{j+1} + \xi \frac{j(j-1)}{2} \frac{(\alpha)_{j-1}}{(2\alpha)_j}.$$

We set $\lambda_j = \hat{\lambda}_j$ for $j=2$ and $j=4$, and solve for ξ and α:

$$\alpha = \frac{3(\tau^2 + 6\tau + 6)}{3\tau^2 + 13\tau + 18} \qquad \text{and} \qquad \xi = \frac{2\alpha + 1}{3(\tau+1)}.$$

The approximation is extraordinarily precise over the entire range of τ values. In the four panels of Figure 6, we plot $\Xi_\tau(z) - z$ and $\hat{\Xi}_\tau(z) - z$ for $\tau = (1, 8, 64, 512)$. Subtracting out the linear component serves to heighten the visual differences between Ξ and our approximation, yet in each case the fit is nearly perfect. The approximation also satisfies the desired limiting behavior. When $\tau = 0$, $\alpha = \xi = 1$, so $\hat{\Xi}_\tau(z) = 1/2$. When $\tau = \infty$, $\xi = 0$ and $\alpha = 1$, so $\hat{\Xi}_\tau(z) = z$.

In Section 3, it is suggested that for reasonably large values of τ, we can obtain a highly tractable form for σ^2 if we approximate $\alpha = 1$ and $\xi = 1/\tau$. Figure 7 shows how α and ξ vary with τ. In the upper panel, we see that $\alpha(\tau)$ is nonlinear for low values of τ, but asymptotes to one as τ heads towards infinity (note the log-scale on the τ axis). In the lower panel, we see that $\xi(\tau)$ converges quite closely to $1/\tau$ by $\tau = 100$ (log-scale on both axes).

[14] Weighting $(z(1-z))^{\alpha-1}$ by the beta function is natural as it transforms the last piece of $\hat{\Xi}_\tau$ into a beta pdf.

Figure 6: Cyclical components of $\Xi_\tau(z)$ and $\hat{\Xi}_\tau(z)$

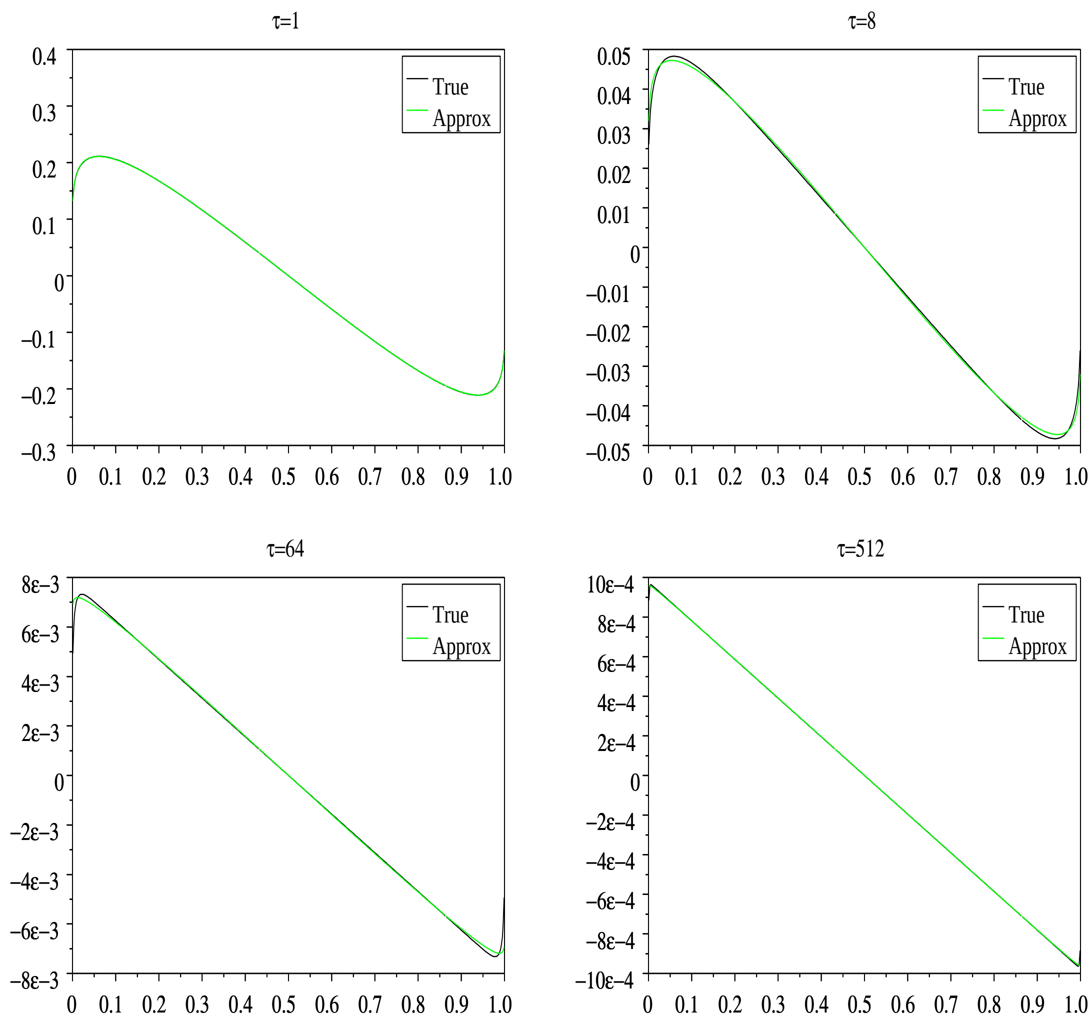

Note: The panels show $\Xi_\tau(z) - z$ and $\hat{\Xi}_\tau(z) - z$ for different values of τ.

Figure 7: Coefficients α and ξ as functions of τ

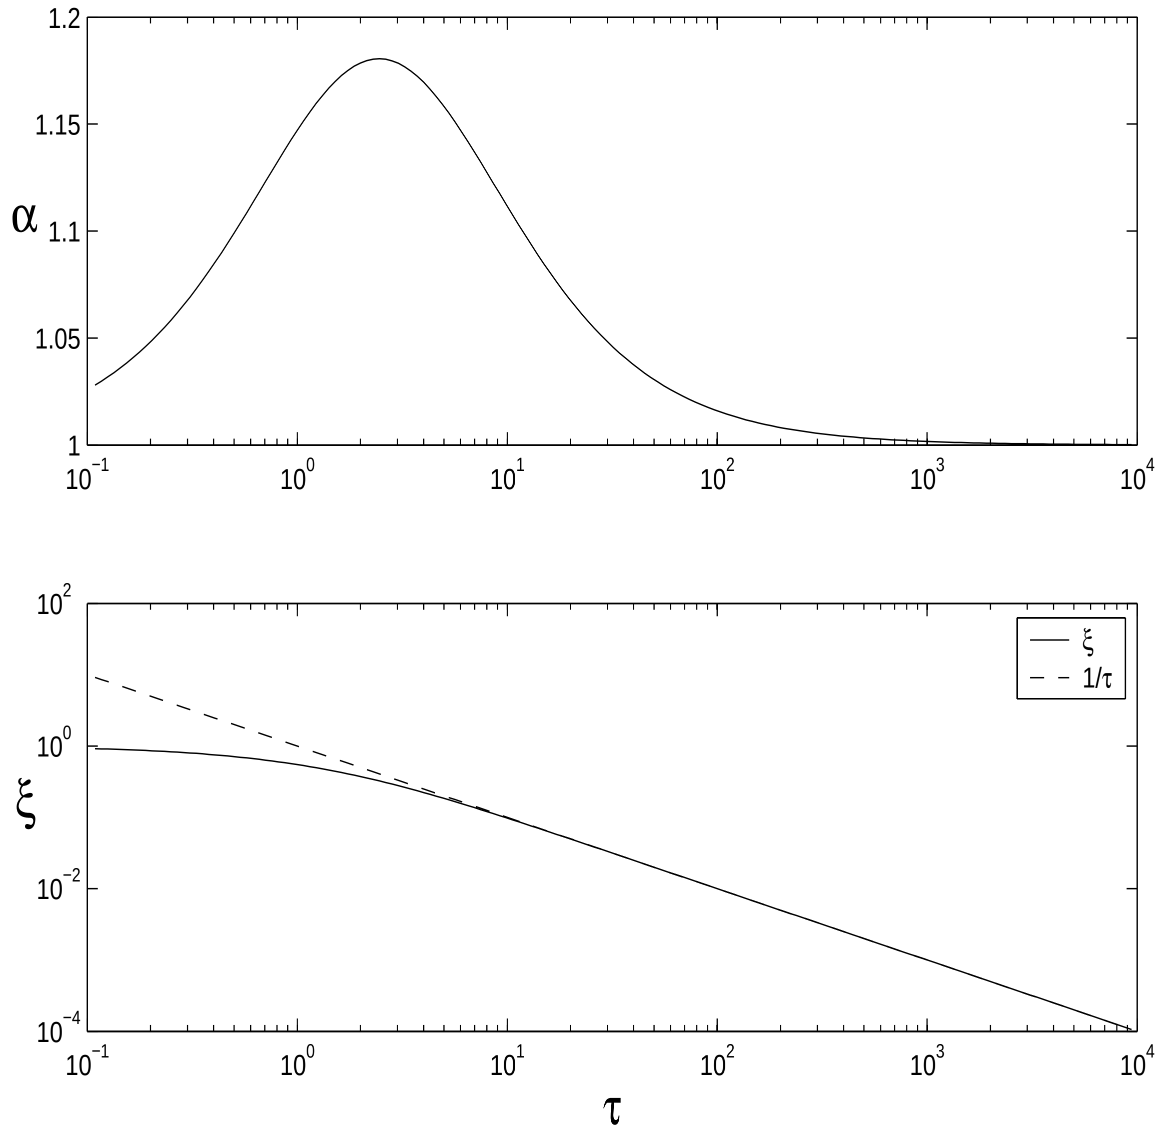

C The function $K[L]$ in the Supervisory Formula Approach

The proposed Supervisory Formula Approach ("SFA") specifies capital for a tranche of credit enhancement level L and thickness T as $(S[L+T]-S[L])$ times the notional size of the underlying pool. The function $S[L]$ is based on another function $K[L]$ plus certain supervisory overrides (i.e., dollar-for-dollar capital up to K_{irb} and a floor level of marginal capital). In this appendix, we show how the function $K[L]$ is grounded in the ULP model.

In the Second Working Paper on Securitisation (Basel Committee on Bank Supervision 2002, Annex 3, ¶574), the function $K[L]$ is specified as

$$K[L] = (1-h)\cdot((1-Beta[L;a,b])L + Beta[L;a,b]c) \tag{22}$$

where

$$\begin{aligned}
h &= (1-K_{\text{irb}}/LGD)^N \\
c &= K_{\text{irb}}/(1-h) \\
\nu &= \frac{1}{N}\left((LGD-K_{\text{irb}})K_{\text{irb}} + 0.25(1-LGD)K_{\text{irb}}\right) \\
f &= \left(\frac{\nu+K_{\text{irb}}^2}{1-h} - c^2\right) + \frac{(1-K_{\text{irb}})K_{\text{irb}}-\nu}{(1-h)\tau} \\
g &= \frac{(1-c)c}{f} - 1 \\
a &= g\cdot c \\
b &= g\cdot(1-c).
\end{aligned}$$

Some translation of notation is needed. The SFA credit enhancement level L is represented as ζ in this paper. The SFA effective number of loans N is here denoted n, and the SFA expected loss given default LGD is denoted ELGD in this paper. The beta cumulative distribution function ($Beta$ in equation 22) is the function $B(\zeta;\cdot,\cdot)$.

The SFA parameters h and c correspond to $H_q(0)$ and μ, respectively. When recovery risk is assumed to be idiosyncratic to the obligor, we have $K_{\text{irb}} = \text{ELGD}\cdot p_q$. Therefore, if K_{irb} and ELGD are observed, we can set $p_q = K_{\text{irb}}/\text{ELGD}$. In a default-mode setting, $H_q(0) = (1-p_q)^n$. Thus, in terms of the SFA notation, we have $h = (1-K_{\text{irb}}/LGD)^N$. By equation (7), $\mu = \text{E}\left[L|x_q\right]/(1-H_q(0)$. As $K_{\text{irb}} = \text{E}\left[L|x_q\right]$, we have $c = K_{\text{irb}}/(1-h)$ in the SFA notation.

The SFA parameter ν is equal to $\text{V}[L|x_q]$ as given in equation (15). If we apply the assumed functional form in equation (17) for VLGD, we can reorder terms to get

$$\text{V}[L|x_q] = \frac{1}{n}\text{ELGD}\cdot p_q\cdot\left((\text{ELGD}-\text{ELGD}\cdot p_q) + \gamma(1-\text{ELGD})\right).$$

Substitute K_{irb} for $\text{ELGD}\cdot p_q$ and set $\gamma = 0.25$ to get the SFA ν. By a similar sequence of substitutions, we can show that the SFA parameter f is our σ as given by equation (14). Finally, SFA parameters g, a and b correspond in a straightforward manner to θ, $\theta\mu$ and $\theta(1-\mu)$, respectively.

References

Basel Committee on Bank Supervision, "The Internal Ratings-Based Approach: Supporting Document to the New Basel Capital Accord," Technical Report, Bank for International Settlements January 2001.

____ , "Second Working Paper on Securitisation," Technical Report, Bank for International Settlements October 2002.

Duffie, Darrell and Nicolae Garleanu, "Risk and Valuation of Collateralized Debt Obligations," *Financial Analysts Journal*, January-February 2001, *57* (1), 41–59.

Fisher, Franklin M. and John J. McGowan, "On the Misuse of Accounting Rates of Return to Infer Monopoly Profits," *American Economic Review*, March 1983, *73* (1), 82–97.

Gordy, Michael B., "A Comparative Anatomy of Credit Risk Models," *Journal of Banking and Finance*, January 2000, *24* (1-2), 119–149.

____ , "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules," Technical Report FEDS 2002-55, Board of Governors of the Federal Reserve System October 2002.

Gupton, Greg M., Christopher C. Finger, and Mickey Bhatia, *CreditMetrics–Technical Document*, New York: J.P. Morgan & Co. Incorporated, April 1997.

Jones, David, "Emerging problems with the Basel Capital Accord: Regulatory capital arbitrage and related issues," *Journal of Banking and Finance*, January 2000, *24* (1-2), 35–58.

Kealhofer, Stephen and Jeffrey R. Bohn, "Portfolio Management of Default Risk," Technical Report, KMV Corporation May 2001.

Pykhtin, Michael and Ashish Dev, "Credit Risk in Asset Securitizations: Analytical Model," *Risk*, May 2002.

____ **and** ____ , "Credit Risk in Asset Securitizations: The Case of CDOs," *Risk*, forthcoming.