Model Foundations for the Supervisory Formula Approach

Michael B. Gordy*
Board of Governors of the Federal Reserve System

July 2004

In its proposal for a New Basel Accord, the Basel Committee on Bank Supervision (2004, Part 2.IV) offers two methodologies for assigning regulatory capital charges to securitization exposures under the Internal Ratings-Based (IRB) approach. The Ratings-Based Approach sets capital primarily as a function of an external rating, such as might be assigned by Moody's or S&P, and is to be employed whenever an external rating is available. As many securitization exposures are not externally rated, the alternative Supervisory Formula Approach (SFA) determines required capital as a function of the characteristics of the collateral pool and contractual properties of the tranche. The chapter sets forth the theoretical foundation for the SFA provided by the "uncertainty in loss prioritization" (ULP) model of Gordy and Jones (2003).

Regulatory application places important constraints on model design. To avoid the inconsistencies that presently make securitization an instrument for regulatory arbitrage (see Jones 2000), the securitization model must be fully compatible with the model underpinning IRB capital charges on whole loans. Cost-effective application to a wide variety of securitizations and participating institutions dictates that the capital rule 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 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. The ULP model satisfies both of these constraints.

The model foundation for the IRB approach is a special case of the class of credit value-at-risk (VaR) models currently in widespread use. Economic capital is set to cover total mark-to-market credit losses over a one-year horizon with probability q.[1] It is assumed that the bank's credit portfolio is infinitely fine-grained in the sense that any single obligor represents a negligible

*To appear in *Structured Credit Products: Pricing, Rating, Risk Management and Basel II*, William Perraudin, Ed., Risk Books, 2004. The opinions expressed here are those of the author, and do not reflect the views of the Board of Governors or its staff. I thank Peter Carr, Enrico De Giorgi, Hemant Ishwaran, David Jones and Dilip Madan 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.

share of the portfolio's total exposure, and that a single, common systematic risk factor drives all dependence across credit losses in the portfolio. The salient 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 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.

The ULP model imposes the ASRF assumptions on the bank's overall portfolio. Bearing in mind that an obligor may appear concurrently in multiple securitization transactions and directly in the bank's own loan book, we assume that the bank's aggregated exposure to each obligor represented in the collateral pool is small relative to the bank's overall portfolio. The ULP model introduces a new source of risk specific to the securitization transaction, and so we require as well that the bank's total exposure to each transaction (that is, summing across the tranches held by the bank) is small relative to the bank's overall portfolio. To be clear, the model does not require that the bank's total holdings of securitization instruments be a small share of its total portfolio. Most importantly, the collateral pool need not be infinitely fine-grained. Indeed, the ULP model can be applied to securitizations of pools ranging from a single loan to infinitely-many loans.

To avoid dependence on the contractual details of the cashflow waterfall, it is necessary to take a stylized view of the allocation of pool credit losses across tranche investors. Pykhtin and Dev (2002, 2003) propose perhaps the simplest approach, which is to assume that losses in the pool are allocated deterministically according to a strict loss prioritization (SLP) rule. Let ζ denote the credit enhancement level, defined as the sum of the par values of all more-junior tranches, and let T denote the tranche's par value or thickness. In practice, both ζ and T are readily observable to market participants. Under the SLP rule, the tranche absorbs pool losses only in excess of ζ, up to a maximum of T. That is, if L denotes pool losses, then the loss for the tranche investor is $\min\{L, \zeta + T\} - \min\{L, \zeta\}$.

The ULP model generalizes this simple rule by adding a stochastic element to the distribution of loss across tranche investors. The credit enhancement level determines ex-ante *expected* protection against losses in the pool, but the actual protection conferred by the prioritized waterfall is random and realized only at the horizon. This 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

[2] See Gordy (2003) for a derivation of this result under minimal restrictions on the portfolio and very general modeling assumptions.

prioritization of cash flows over the life of a securitization, the nominal 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 the model is *not* positing that there is operational or legal risk in the execution of securitization contracts. The new source of uncertainty 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. The intuition draws from the long vein of econometrics literature on models with hidden parameters. 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 along the lines of Duffie and Garleanu (2001) would naturally be preferred.

We first develop the ULP model under very general assumptions similar to those in Gordy (2003). A general methodology for simplifying computation of capital charges is then set forth. The model specification is completed in the next section, where we impose the specific functional forms and distributional assumptions to which IRB capital charges are calibrated. This ensures full consistency between the IRB approach and our proposed treatment of securitization exposures. We conclude with discussion of regulatory application.

1 Modeling uncertainty in loss prioritization

We seek to assign an economic capital charge to a securitization exposure held in an asymptotically fine-grained portfolio. As in the Basel II framework for the treatment of whole loans, total capital must be sufficient to cover portfolio credit loss up to some percentile q of the loss distribution.

Dependence between obligors in the economy arises due to a single systematic risk-factor X. Conditional on X, losses on each loan in the bank's portfolio and in the securitized pool are independent. Let L be book-value (or "default-mode") loss incurred within the pool as a share of total pool exposure, and let $H_q(\cdot)$ be the cumulative distribution function (cdf) of L conditional on $X = x_q$, where x_q denotes the q^{th} percentile of the distribution of X. 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 (2003).

The collateral pool is securitized into a continuum of prioritized tranches, each of which is

[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.

of infinitesimal thickness.[4] Assume for the moment that losses are allocated deterministically in accordance with strict loss prioritization (SLP), and let $K_{\text{slp}}(\zeta)$ be the cumulative capital charge (as a share of total pool exposure) on the juniormost share ζ of the structure. Under the ASRF assumptions, this is given by

$$K_{\text{slp}}(\zeta) = \text{E}\left[\min\{\zeta, L\} | X = x_q\right] = \int_0^1 \min\{\zeta, \ell\}\, dH_q(\ell) = \zeta - \int_0^\zeta H_q(\ell) d\ell \tag{1}$$

where the last equality follows from integration-by-parts. The marginal contribution to the cumulative capital function, or $K'_{\text{slp}}(\zeta)$, gives the capital charge per dollar of tranche notional value for the infinitesimally thin tranche at credit enhancement level ζ. If the bank holds the set of infinitesimal tranches from ζ to $\zeta + T$, then economic capital on the bank's position is $K_{\text{slp}}(\zeta + T) - K_{\text{slp}}(\zeta)$ per dollar of total pool exposure. Expressed as a percentage of the bank's exposure, capital is $(K_{\text{slp}}(\zeta + T) - K_{\text{slp}}(\zeta))/T$.

By construction, the sum of capital charges across all tranches is $K_{\text{slp}}(1)$. At $\zeta = 1$, we have

$$K_{\text{slp}}(1) = \text{E}\left[\min\{1, L\} | X = x_q\right] = \text{E}\left[L | x_q\right] = \mathcal{K}_{\text{pool}},$$

where $\mathcal{K}_{\text{pool}}$ denotes the capital the pool would require if held on balance sheet. Thus, securitization does not alter the total capital that must be held against a pool of loans.

We now generalize this basic result to allow for uncertainty in loss prioritization. Let $Z(\zeta)$ be the effective credit enhancement level associated with the nominal credit enhancement level ζ. Both ζ and Z are expressed as a share of total pool exposure, so $Z(\zeta)$ is a random process mapping $[0, 1]$ to $[0, 1]$. A natural choice for the specification of Z is the Dirichlet process with parameter τ.[5] Under this specification, we have $Z(0) = 0$ and $Z(1) = 1$ (with certainty), and the marginal distribution for $Z(\zeta)$ is $\text{Beta}(\tau\zeta, \tau(1 - \zeta))$. Intuitively, the assumption of a Dirichlet process allows the effective credit enhancement level $Z(\zeta)$ to "wiggle" around its expected value ζ while still pinning down the boundary values at $\zeta = 0$ and $\zeta = 1$. Parameter τ controls the variance of the process around its expected path. As $\tau \to \infty$, the wiggle room disappears, and $|Z(\zeta) - \zeta| \to 0$, almost surely. Sample paths for $\tau = 20$ and $\tau = 400$ are shown in Figures 1 and 2, respectively.

When Z is independent of all other risks in the portfolio, the ASRF assumptions imply (as usual) that the appropriate capital charge for a tranche is its expected loss conditional on $X = x_q$.[6]

[4] Pykhtin and Dev (2002, 2003) appear to have been the first to use infinitesimally thin tranches as an analytical convenience in modeling securitization.

[5] More precisely, Z follows a Dirichlet process with parameter τU, where U is the uniform measure on $[0, 1]$. The Dirichlet process was introduced in the Bayesian literature by Ferguson (1973). It can arise as the limiting process for the Dirichlet distribution, which is a multivariate generalization of the beta distribution (Ishwaran and Zarepour 2002, Theorem 3). It has previously appeared in the finance literature in Chow (1998) and as the "gamma bridge" of Ribeiro and Webber (2003).

[6] If securitization instruments were to account jointly for a significant share of the bank's portfolio, and if (conditional on $X = x_q$) the Z processes were not independent across those securitizations, then the risk in loss prioritization would not be diversified away, and thus would demand capital of its own. Similar reasoning motivates the assumption that the bank's aggregate exposure to each obligor in the collateral pool constitutes a trivial share of the bank's portfolio.

Figure 1: Sample paths of Dirichlet Process

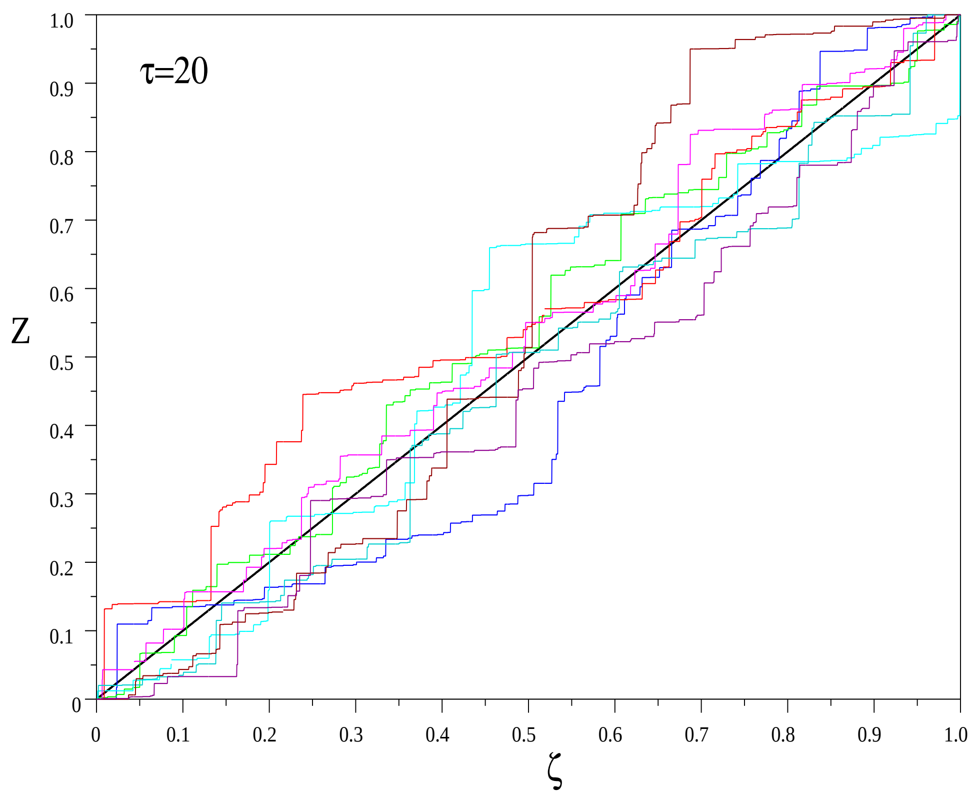

The cumulative capital function is then

$$
\begin{aligned}
K(\zeta) = \mathrm{E}\left[\min\{Z(\zeta), L\} | X{=}x_q\right] & \\
&= \mathrm{E}\left[Z(\zeta)\right] - \mathrm{E}\left[\int_0^{Z(\zeta)} H_q(\ell) d\ell\right] \\
&= \zeta - \int_0^1 \frac{z^{\tau\zeta-1}(1-z)^{\tau(1-\zeta)-1}}{B(\tau\zeta, \tau(1-\zeta))} dz \int_0^z H_q(\ell) d\ell \\
&= \zeta - \int_0^1 (1 - B(z; \tau\zeta, \tau(1-\zeta))) H_q(z) dz \qquad (2)
\end{aligned}
$$

where the function $B(y; a, b)$ is the Beta(a, b) cdf evaluated at y.[7] Because $Z(1) = 1$ with certainty, we have $K(1) = \mathcal{K}_{\text{pool}}$ as under the SLP case. Thus, uncertainty in loss prioritization does not alter the total required capital for the securitization.

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

Figure 2: Sample paths of Dirichlet Process

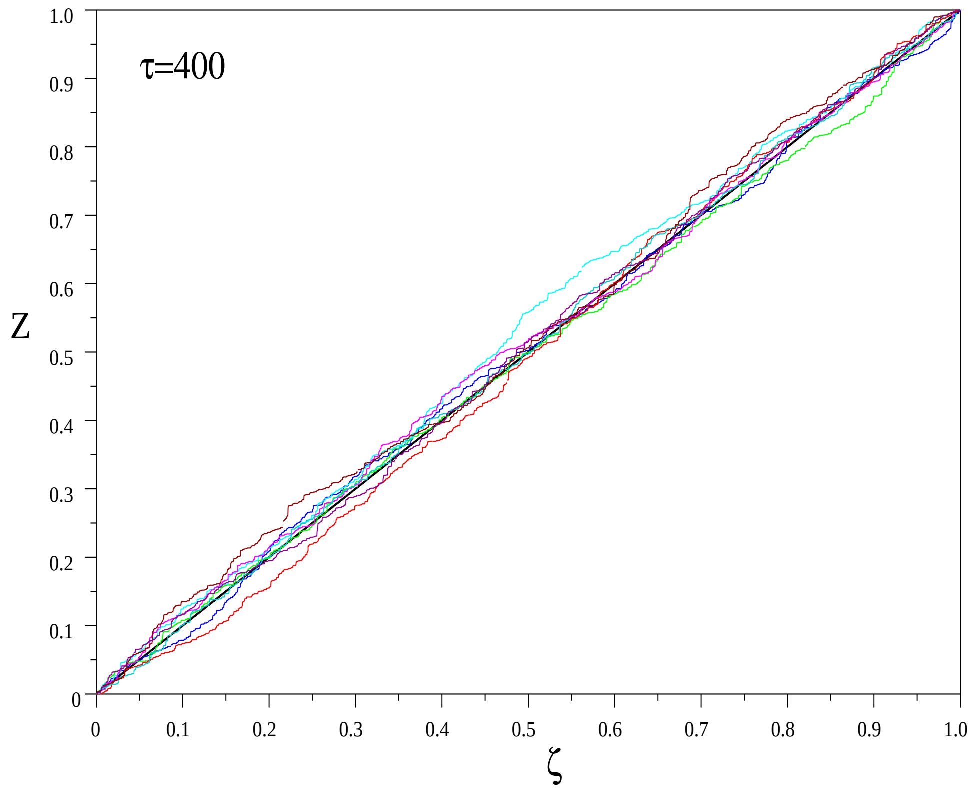

The ULP model embeds the SLP model as a limiting case. As $\tau \to \infty$, uncertainty in the division of risk vanishes, and the beta cdf for $Z(\zeta)$ converges to $\mathbb{1}_{\{z \geq \zeta\}}$. Thus,

$$\lim_{\tau \to \infty} K(\zeta) = \zeta - \int_0^1 (1 - \mathbb{1}_{\{z \geq \zeta\}}) H_q(z) dz = \zeta - \int_0^{\zeta} H_q(z) dz = K_{\text{slp}}(\zeta).$$

As $\tau \to 0$, the contractual exposure shares become less and less informative of the actual division of risk. The limiting distribution for $Z(\zeta)$ is Bernoulli such that $Z(\zeta) = 1$ with probability ζ and $Z(\zeta) = 0$ and probability $1 - \zeta$. Thus,

$$\lim_{\tau \to 0} B(z; \tau\zeta, \tau(1 - \zeta)) = 1 - \zeta \qquad \forall z \in (0, 1).$$

Substitute this result in equation (2) to obtain

$$\lim_{\tau \to 0} K(\zeta) = \zeta - \int_0^1 \zeta H_q(z) dz = \zeta \cdot \mathrm{E}\left[L | x_q\right],$$

which implies a proportional sharing of $\mathcal{K}_{\text{pool}}$ 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.

If the collateral pool is itself asymptotically fine-grained, as is a reasonable characterization of most retail securitizations, then equation (2) has a simple analytic solution. In the asymptotic case, $H_q(z) = \mathbb{1}_{\{z \geq \text{E}[L|x_q]\}}$. Using the integral

$$\int B(y; a, b) dy = y \cdot B(y; a, b) - \frac{a}{a+b} \cdot B(y; a+1, b) \tag{3}$$

we find

$$K(\zeta) = \zeta \cdot B(\text{E}\left[L|x_q\right]; \tau\zeta + 1, \tau(1-\zeta)) + \text{E}\left[L|x_q\right](1 - B(\text{E}\left[L|x_q\right]; \tau\zeta, \tau(1-\zeta))). \tag{4}$$

In the SLP case, any tranche of a securitization of an asymptotically fine-grained pool that is senior to the $\mathcal{K}_{\text{pool}}$ threshold requires zero capital, and any tranche of such a securitization that is strictly junior to the $\mathcal{K}_{\text{pool}}$ threshold requires dollar-for-dollar capital. In our generalized model, senior tranches always require some capital because of the possibility that the "realized" exposure of the senior tranches exceeds $1 - \mathcal{K}_{\text{pool}}$, and junior tranches require less than dollar-for-dollar capital because of the possibility that their realized exposure is less than $\mathcal{K}_{\text{pool}}$. Thus, the extended model unambiguously increases capital for tranches senior to $\mathcal{K}_{\text{pool}}$, and unambiguously reduces capital for tranches junior to $\mathcal{K}_{\text{pool}}$. For tranches that straddle this breakpoint, the effect is ambiguous but typically small.

For less fine-grained pools, uncertainty in loss prioritization has a smaller effect. Figure 3 shows the effect of τ and n using the model specification imposed in the SFA. The case of $n = \infty$ is shown in the bottom panels. When $\tau = \infty$, marginal capital is dollar-for-dollar up to $\mathcal{K}_{\text{pool}}$ 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 collateral pool of $n = 128$ (middle row of panels), which might be representative of the most fine-grained CDO pools, undiversified idiosyncratic risk is sufficient to smooth away the cliff effect at $\mathcal{K}_{\text{pool}}$.[8] In this case, the capital and marginal capital curves for $\tau = 1000$ are indistinguishable from those of $\tau = \infty$. Uncertainty in loss prioritization 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 $Z(\zeta)$ to have any additional smoothing effect.

[8] Conversations with market participants suggest that few CDOs contain over 200 names. When corrected for concentration in exposure sizes as described later in this chapter, we would expect that "effective n" would typically be under 100.

Figure 3: Effect of τ on capital and marginal capital

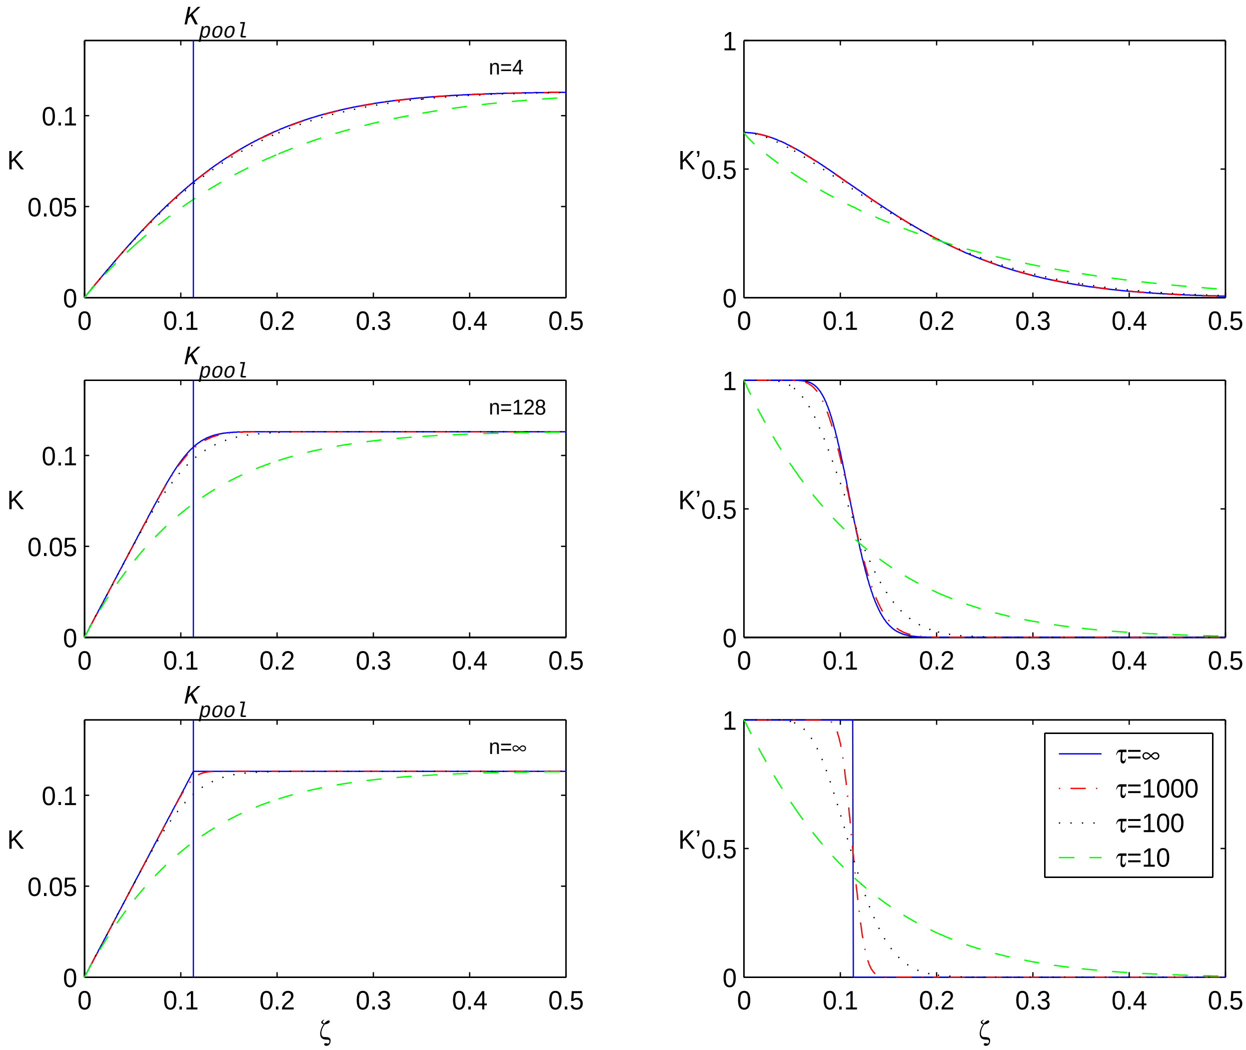

Note: Homogeneous collateral pools with default probability PD $= 0.02$, expected LGD $= 0.5$ and asset-correlation $\rho = 0.2$. We set VaR target quantile $q = 0.999$ and recovery risk parameter $\gamma = 0.25$.

2 Fitting a simple functional form

Unless we restrict ourselves to analysis of securitizations of asymptotically fine-grained pools, the pool's conditional loss distribution $H_q(z)$ is likely to be analytically intractable, which would require that we use numerical integration or simulation to solve for $K(\zeta)$ in equation (2).[9] For regulatory purposes, a simpler and more transparent functional solution is required, even if it comes at slight expense in precision.

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

$$F(\zeta) = 1 - \frac{K'(\zeta)}{K'(0)}. \tag{5}$$

This definition is useful because it exploits three known properties of the first derivative of $K(\zeta)$: $K'(\zeta)$ is nonincreasing on the unit interval, and we have $K'(0) \approx 1 - H_q(0)$ and $K'(1) = 0$.[10] 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 approximated by the cdf of a simple distribution such as the beta.

We parameterize F in terms of its mean (μ) and variance (σ^2). To get the mean parameter, we rearrange equation (5) 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

$$\mu = \frac{K(1)}{K'(0)} \approx \frac{\mathrm{E}\left[L|x_q\right]}{1 - H_q(0)}. \tag{6}$$

The variance of F is more challenging. By definition,

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

[9]Monte Carlo simulation of $K(\zeta)$ is straightforward but computationally intensive. Since $K(\zeta) = \mathrm{E}\left[\min\{Z(\zeta), L\}|x_q\right]$, we need only draw a sample of paths $\{z_1(\zeta), \ldots, z_T(\zeta)\}$ for $Z(\zeta)$ and a sample of $\{\ell_1, \ldots, \ell_T\}$ for L (from the H_q distribution). $K(\zeta)$ is estimated by $(1/T)\sum \min\{z_i(\zeta), \ell_i\}$.

[10]The exact derivative of $K(\zeta)$ at $\zeta = 0$ is intractable. For the limiting case of $\tau = \infty$, we have $K'_{\mathrm{slp}}(\zeta) = 1 - H_q(\zeta)$, so $K'_{\mathrm{slp}}(0) = 1 - H_q(0)$. One can show that $K'(0)$ converges to this limit at an exponential rate in τ, so the approximation works well for large τ. Moreover, as a practical matter, one finds that the fitted cumulative capital function $\hat{K}(\zeta)$ is relatively insensitive to errors in the $K'(0)$ input.

This integral does not have an analytical solution, but asymptotic expansion in $1/\tau$ provides an approximation for large τ. Gordy and Jones (2003) show that

$$\sigma^2 \approx \frac{1}{K'(0)}\left(\mathrm{V}[L|x_q] + \mathrm{E}\left[L|x_q\right]^2\right) - \mu^2 + \frac{1}{\tau}\frac{1}{K'(0)}\left(\mathrm{E}\left[L|x_q\right]\left(1-\mathrm{E}\left[L|x_q\right]\right) - \mathrm{V}[L|x_q]\right). \quad (7)$$

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 collateral pool (i.e., the impact of pool granularity), and equals the exact formula for σ^2 in the SLP case. The second is the contribution of uncertainty in loss prioritization and is inversely proportional to τ.

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 $\sigma^2 = (1/\tau)\mathrm{E}\left[L|x_q\right]\left(1-\mathrm{E}\left[L|x_q\right]\right)$.

3 A complete specification

Thus far, we have not needed to specify a model for portfolio loss or to choose a cumulative distribution function to assign to the fitting function. We complete the specification in order to arrive at the formulation seen in the SFA.

We assume that the collateral pool is homogeneous and that the conditional loss distribution H_q comes from a single-factor default-mode model with conditionally idiosyncratic recovery risk (that is, LGDs are independent conditional on X). 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 conditional 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

$$\mathrm{V}[L|x_q] = \frac{1}{n}\left(\mathrm{ELGD}^2 p_q(1-p_q) + p_q \mathrm{VLGD}^2\right) \quad (8)$$

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 and that the conditional probability of default is given by

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

where Φ is the standard normal cdf and ρ is the correlation in asset returns.[11] 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

$$\text{VLGD}^2 = \gamma \cdot \text{ELGD} \cdot (1 - \text{ELGD}) \tag{9}$$

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., zero recovery with probability ELGD, full recovery 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 is a natural choice given its unit interval support, and thus has been adopted in the SFA.[12] 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 $Z(\zeta)$.

To distinguish between the "true" fitting function, which is of intractable form, and the 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

$$\begin{aligned} \hat{K}(\zeta) &= \int_0^\zeta \hat{K}'(y)dy = K'(0)\int_0^\zeta (1 - \hat{F}(y))dy \\ &= (1 - H_q(0)) \cdot (\zeta \cdot (1 - B(\zeta; \theta\mu, \theta(1-\mu))) + \mu \cdot B(\zeta; \theta\mu + 1, \theta(1-\mu))) \end{aligned} \tag{10}$$

where the final equality follows using equation (3). Note that when $n = \infty$, we have $\theta = \tau - 1$ and $\hat{K}(\zeta)$ simplifies to

$$\begin{aligned} \hat{K}(\zeta) = \zeta \cdot (1 - B(\zeta; (\tau-1)\mathcal{K}_{\text{pool}}, (\tau-1)(1-\mathcal{K}_{\text{pool}}))) \\ + \mathcal{K}_{\text{pool}} \cdot B(\zeta; (\tau-1)\mathcal{K}_{\text{pool}} + 1, (\tau-1)(1-\mathcal{K}_{\text{pool}})). \end{aligned} \tag{11}$$

Gordy and Jones (2003) examine the robustness of the fitted $\hat{K}(\zeta)$ to the "true" $K(\zeta)$ given by equation (2), and 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 very low expected LGD and low asset-correlation. At least in the

[11] The CreditMetrics model is documented by Gupton, Finger and Bhatia (1997). For a derivation of p_q from this model, see Koyluoglu and Hickman (1998) or Gordy (2000). This same expression appeared earlier in Vasicek's (1991) analysis of the asymptotic loss distribution in KMV Portfolio Manager.

[12] Note that this application of the beta cdf has nothing to do with the use of the beta as the distribution for $Z(\zeta)$ and the distribution for LGD.

context of the proposed New Basel Accord, this circumstance cannot arise in practice.[13]

Consistency with the IRB treatment of whole loans implies that pool capital, $\mathcal{K}_{\text{pool}}$, is equal to the IRB capital charge for the pool. In the one-year default-mode setting of the model, we have $\mathcal{K}_{\text{pool}} = \mathcal{K}_{\text{irb}} = \text{ELGD} \cdot p_q$. Some simplification in implementing the ULP model can be obtained by noting that parameters PD, ρ and q enter the calculations only via p_q, and that $p_q = \mathcal{K}_{\text{irb}}/\text{ELGD}$. Thus, a sufficient set of collateral pool parameters is n, $\mathcal{K}_{\text{irb}}$ and ELGD. As a practical matter, parameterization of the ULP in terms of $\mathcal{K}_{\text{irb}}$ compensates for some of the limitations of the model's default-mode notion of credit loss.[14] 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 equal the sum of economic losses to the tranches.[15] Under the Advanced IRB approach, capital charges for the underlying pool incorporate maturity effects, so parameterization in terms of $\mathcal{K}_{\text{irb}}$ may be more robust than parameterization in terms of PD and ρ.

4 Application to regulatory capital treatment

The Supervisory Formula Approach specifies capital for a tranche of credit enhancement level ζ and thickness T as $(S[\zeta+T]-S[\zeta])$ times the notional size of the collateral pool. The function $S[\zeta]$ has at its baseline the ULP $\hat{K}(\zeta)$ function of equation (10), but imposes certain supervisory overrides on top. In addition to tranche parameters ζ and T, the bank is required to supply collateral pool parameters n, $\mathcal{K}_{\text{irb}}$ and ELGD. Supervisory values of model parameters $\tau = 1000$ and $\gamma = 0.25$ are imposed.[16]

From an operational perspective, the only significant challenge in implementing the ULP model lies in obtaining the pool characteristics n, $\mathcal{K}_{\text{irb}}$ and ELGD. As collateral pools are never truly homogeneous, the number of loans must be adjusted for concentration. For example, a pool with one \$100 million exposure and 99 \$1 million exposures has a loss distribution more akin to that of a homogeneous pool of three or four loans than one of 100 loans. The SFA specifies that the "effective" n be calculated as an inverse Herfindahl index of the exposure sizes in the pool. Gordy (2003, §4) shows that this specification can arise naturally in mapping from a heterogeneous portfolio to a comparable homogeneous portfolio. In some situations, investors may receive incomplete information on the distribution of exposure sizes in a pool, such as the share of the largest exposure

[13]For corporate borrowers, investment grade status implies high asset-correlation. Retail borrowers can have a lower asset-correlation, but securitization of a single retail loan is prohibitively expensive, at least under current market practice.

[14]Pykhtin and Dev's (2002, 2003) model suffers from the same limitations.

[15]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.

[16]Translation from the notation of this chapter to that of Basel Committee on Bank Supervision (2004) is straightforward. The quantities $H_q(0)$, μ, σ^2 and θ map to the SFA quantities h, c, f and g, respectively. The expression for $\text{V}[L|x_q]$ in equation (8) can be rearranged to equal the SFA ν parameter.

and the cumulative share of the four largest exposures. The SFA provides a rule for estimating n from such information that is based on upper bounds for the Herfindahl index given one or two concentration ratios.

For an originator, obtaining the pool average $\mathcal{K}_{\text{irb}}$ and ELGD poses no special burden because these would need to be estimated for regulatory capital if retained on balance sheet. Outside investors sometimes have more limited access to information on the collateral pool. However, it is reasonable to expect investors in unrated tranches to have sufficient information to estimate $\mathcal{K}_{\text{irb}}$. Otherwise, in the absence of a rating, it is difficult to see how an investor could perform due diligence. In most situations, $\mathcal{K}_{\text{irb}}$ for the pool would be calculated on a "bottom-up" basis from instrument-level PD and ELGD estimates, so the pool average ELGD would be straightforward to estimate at the same time. For certain asset classes such as purchased receivables, bottom-up estimation can pose a significant operational burden, so banks are permitted to use a "top-down" approach in which PD and ELGD are estimated at the pool level.[17] If the bank could obtain only historical loss rates for comparable pools, then the most conservative treatment consistent with the estimated expected loss (EL) for the pool is to set ELGD to 100% and PD to EL. $\mathcal{K}_{\text{irb}}$ is then calculated according to the corporate risk-weight function.

At least for the seniormost tranches, setting ELGD to 100% results in at least as high a capital charge as under any other ELGD assumption, so is "conservative" in the supervisory sense.[18] Holding fixed $\mathcal{K}_{\text{irb}}$, n and the supervisory ULP parameters (τ, γ), denote by $K(\zeta;\lambda)$ the $K(\zeta)$ function for ELGD $= \lambda$. In the Appendix, I prove that

Proposition 1 *For all* $\lambda < 1$, $K(\zeta;1) \leq K(\zeta;\lambda)$.

Consider a senior tranche beginning at nominal credit enhancement level ζ and of thickness $1-\zeta$. In the ULP model, capital for such a tranche is given by $\mathcal{K}_{\text{irb}} - K(\zeta;\lambda)$ per dollar of total pool collateral. Proposition 1 implies that this capital charge is maximized when $\lambda = 1$.

The result applies to a collection of tranches as well. Say, for example, that supervisors were to impose ELGD $= 1$ on all tranches above some credit enhancement level ζ (e.g., $\zeta = \mathcal{K}_{\text{irb}}$). Then Proposition 1 implies that total capital on tranches senior to ζ would be greater than under any alternative ELGD assumption. However, if we consider on its own a tranche other than the most senior, it is not possible to say whether the proposed treatment is the most conservative. The intuition is that changing ELGD redistributes capital across tranches, but cannot change the total capital for the securitization as a whole. If capital on the seniormost tranche increases, then capital on the juniormost tranche must be reduced. The effect on the intermediate tranches is ambiguous. As a practical matter, experience with the model suggests that assuming ELGD of 100% should result in conservative treatment for any highly-rated tranche with credit enhancement level above

[17]See Basel Committee on Bank Supervision (2004, §365–7).

[18]For retail pools, estimation of ELGD is unnecessary as it drops out of the capital formula. Given $\mathcal{K}_{\text{irb}}$, we need ELGD only to calculate the contribution of undiversified idiosyncratic risk to the variance of pool loss. As $n \to \infty$, such risk vanishes.

$\mathcal{K}_{\text{irb}}$. For lower mezzanine positions, however, assuming ELGD of 100% may not be desirable from a supervisory viewpoint.

As noted in many industry comments on the Basel proposal, supervisory overrides in the SFA can add a substantial buffer of regulatory capital onto the baseline ULP capital function. The most important of these are a dollar-for-dollar capital requirement for tranches with $\zeta \leq \mathcal{K}_{\text{irb}}$ and a floor on capital of 56 basis points per dollar of notional tranche value. A full explanation for these overrides must consider supervisory concerns and objectives that are outside the context of the ULP model and thus beyond the scope of this chapter. Clearly, however, the overrides serve to limit exploitation of inadequacies in the model's stylized representation of a securitization.

The ULP model assumes that the collateral pool contains only simple whole loans or bonds. In practice, it is not uncommon for securitized pools to include other asset types. In particular, the collateral pool may consist of tranches of other securitizations. Such resecuritizations are often termed "CDOs of CDOs," but the collateral may include a variety of structured finance instruments, such as asset-backed securities (ABS) and commercial mortgage-backed securities.[19] Pools of CDO tranches behave differently from pools of ordinary loans, and so formulaic implementation of the ULP model in a regulatory rule may be subject to manipulation.

To take a highly stylized example, consider a bank that wishes to securitize a large retail pool of loans. It could package these loans in two ways:

(a) an ordinary ABS in which tranches of the securitized pool are sold directly to investors; or

(b) a CDO of an ABS in which the securitized pool is sold off (as a single 100% tranche) to a CDO conduit, which in turn is broken into tranches for sale to investors.

These structures differ in legal form, but not in risk. That is, the nominal tranche $(\zeta, \zeta + T]$ of security (a) has the same return distribution as the comparable tranche of security (b). However, their treatment under the SFA would differ. For security (a), we would assign $n_a = \infty$ and $\mu_a = \mathcal{K}_{\text{irb}}$. For security (b), we would assign $n_b = 1$. With probability one, there will be at least one default in the underlying ABS retail pool, so the ABS security comprising the CDO collateral pool should be assigned a PD of 100% and ELGD equal to $\mathcal{K}_{\text{irb}}$.[20] This gives us $H_q(0) = 0$, so $\mu_b = \mathcal{K}_{\text{irb}}$. Because the ABS pool is infinitely fine-grained, the "true" conditional variance of loss ($\text{V}[L|x_q]$) for the security is zero. If this were recognized in the SFA, then $\sigma_b^2 = \sigma_a^2$, so $K_a(\zeta) = K_b(\zeta)$ for all ζ. However, for the sake of parsimony, the SFA hardwires the specification for conditional variance in equation (9). This may be appropriate for a collateral pool of whole loans, but in this case implies that $\sigma_b^2 \gg \sigma_a^2$, which is turn implies that $K_a(\zeta) \geq K_b(\zeta)$ for all ζ. Thus, in the absence of the "dollar-for-dollar up to $\mathcal{K}_{\text{irb}}$" override, an originator could restructure (a) into (b) in order to reduce the burden of SFA capital for the junior tranches.

[19] According to Ganapati and Tejwani (2002), the resecuritization market surpassed $45 billion in size by mid-2002 and constituted over one-fifth of the overall new issue volume of cash flow CDOs over the prior two years.

[20] In order to avoid the use of internal models in estimating ELGD, Basel Committee on Bank Supervision (2004, ¶634) requires that securitization exposures in the collateral pool be assigned ELGD of 100%. While this changes the calculations in this example, the qualitative result is the same.

Restructuring can also be used to shift the burden of regulatory capital in the opposite direction. To take another stylized example, consider an originating bank with a collateral pool of ten loans. Say the bank were to divide each loan into 100 identical shares, and then gather together one share of each loan to form 100 identical collateral pools. Each pool is securitized, and in each case the originating bank retains an equity position large enough to absorb all material risk. If the 100 equity positions were gathered together for resecuritization, the new pool would appear to have $n = 100$, but would in fact behave like a pool of ten loans.[21] The problem here is that the restructuring violates the ULP model's assumption of conditional independence in the collateral pool. In the absence of a capital floor in the SFA, the regulatory capital charge on senior tranches of the resecuritization would be too low.

Appendix: Proof of Proposition 1

The proof of Proposition 1 is an application of the theory of stochastic dominance. In order to make clear their dependence on the ELGD parameter, we write the pool loss as L_λ and its conditional cdf as $H_q(\ell;\lambda)$. Equation (2) implies that Proposition 1 holds if and only if

$$\int_0^1 (1 - B(\ell;\tau\zeta,\tau(1-\zeta)))H_q(\ell;\lambda)\,d\ell \leq \int_0^1 (1 - B(\ell;\tau\zeta,\tau(1-\zeta)))H_q(\ell;1)\,d\ell \qquad (12)$$

for all $\zeta \in [0,1]$ and $\lambda < 1$. Define the function $w_\zeta(\ell)$ by

$$w_\zeta(\ell) = \int_0^\ell (1 - B(t;\tau\zeta,\tau(1-\zeta)))\,dt$$

Using integration by parts, we find that equation (12) holds if and only if

$$\int_0^1 w_\zeta(\ell)h_q(\ell;\lambda)\,d\ell \geq \int_0^1 w_\zeta(\ell)h_q(\ell;1)\,d\ell$$

for all $\zeta \in [0,1]$ and $\lambda < 1$. The function $h_q(\cdot;\lambda)$ is the conditional probability density function for L_λ, which allows us to write the necessary and sufficient condition as

$$\mathrm{E}\left[w_\zeta(L_\lambda)|X{=}x_q\right] \geq \mathrm{E}\left[w_\zeta(L_1)|X{=}x_q\right] \qquad (13)$$

for all $\zeta \in [0,1]$ and $\lambda < 1$.

Consider two variables Y_1 and Y_2 with cdfs F_1 and F_2, respectively. We say that F_1 is a mean-preserving spread of F_2 if $\mathrm{E}[Y_1] = \mathrm{E}[Y_2]$ and Y_2 second-order stochastic dominates Y_1 (written as $Y_2 \succ_2 Y_1$). The seminal result of Rothschild and Stiglitz (1970) on mean-preserving spreads is

[21] Again, I am abstracting from the special rule for ELGD on securitization exposure collateral noted in footnote 20.

Lemma 1

For positive random variables Y_1 and Y_2 with $\mathrm{E}[Y_1] = \mathrm{E}[Y_2]$, the following statements are equivalent:

- $Y_2 \succ_2 Y_1$;
- $Y_1 =^d Y_2 + \epsilon$ *for some* ϵ *such that* $\mathrm{E}[\epsilon | Y_2] = 0$;
- $\mathrm{E}[u(Y_2)] \geq \mathrm{E}[u(Y_1)]$ *for all* u *increasing and concave; and*
- $\int_0^y (F_2(t) - F_1(t))dt \leq 0$ *for all* $y \geq 0$.

Because the beta cdf is increasing (at least weakly), the w function is increasing and concave (at least weakly). Therefore, if $H_q(\cdot; 1)$ is a mean-preserving spread of $H_q(\cdot; \lambda)$, then the condition in equation (13) follows immediately.

Let $U_\lambda(i)$ be the loss on loan i ($i = 1, \ldots, n$) in the pool for ELGD $= \lambda$. Conditional on $X = x_q$, we have $U_\lambda(i) \equiv D_i \cdot \mathrm{LGD}_i$ where D_i is a Bernoulli default indicator with $\Pr(D_i = 1 | x_q) = \mathcal{K}_{\mathrm{irb}}/\lambda$ and LGD_i is a beta-distributed loss given default with mean λ and variance $\gamma\lambda(1 - \lambda)$. When ELGD equals 1, LGD_i is equal to 1 with certainty, so each U_1 is distributed Bernoulli with $\Pr(U_1(i) = 1 | x_q) = \mathcal{K}_{\mathrm{irb}}$. For $\lambda < 1$, the U_λ have a mixed distribution. However, regardless of λ, we have $\mathrm{E}[U_\lambda(i) | x_q] = \mathcal{K}_{\mathrm{irb}}$. To show second-order stochastic dominance, we need the following "single-crossing" result:

Lemma 2

Let Y_1 and Y_2 be random variables on $[0, 1]$ with cdfs F_1 and F_2 and equal means. If there exists $y_0 > 0$ such that $F_1(y) \geq F_2(y)$ for $0 \leq y < y_0$ and $F_1(y) \leq F_2(y)$ for $y_0 \leq y \leq 1$ (and we do not have $F_1(y) = F_2(y)$ everywhere on at least one of these intervals), then $Y_2 \succ_2 Y_1$.

The lemma is a corollary of Fishburn and Vickson (1978, §2.16).

In our application, $\Pr(U_1 \leq y | x_q) = 1 - \mathcal{K}_{\mathrm{irb}}$ for all $y < 1$. For any $\lambda < 1$,

$$\Pr(U_\lambda = 0 | x_q) = 1 - \mathcal{K}_{\mathrm{irb}}/\lambda < 1 - \mathcal{K}_{\mathrm{irb}} = \Pr(U_1 = 0 | x_q).$$

Because the beta distribution is continuous, there must exist $\delta > 0$ such that for all $y \in (1 - \delta, 1)$,

$$\Pr(U_\lambda \leq y | x_q) > 1 - \mathcal{K}_{\mathrm{irb}} = \Pr(U_1 \leq y | x_q).$$

Thus, the two distributions must cross at least once. They cannot cross more than once, because the conditional cdf for U_1 is flat and the conditional cdf for U_λ is strictly increasing. Therefore, the single-crossing condition of Lemma 2 is satisfied, so $U_\lambda(i) \succ_2 U_1(i)$ conditional on x_q for all loans i and for all $\lambda < 1$.

Our final lemma is

Lemma 3

Let $Y_j(i)$ for $i = 1, \ldots, n$ be iid draws from a distribution F_j on the positive real line. Let $M_j = \frac{1}{n}\sum_{i=1}^{n} Y_j(i)$ for $j = \{1, 2\}$, and let G_j denote the distribution of M_j. If F_1 is a mean-preserving spread of F_2, then G_1 is a mean-preserving spread of G_2.

Proof:[22] By Lemma 1, we have a sequence of variables ϵ_i such that $Y_1(i) =^d Y_2(i) + \epsilon_i$ and $\mathrm{E}[\epsilon_i | Y_2(i)] = 0$. Therefore, we have $M_1 =^d M_2 + \epsilon^*$ where $\epsilon^* \equiv \frac{1}{n}\sum_{i=1}^{n} \epsilon_i$. Because the $Y_j(i)$ are iid across i, the ϵ_i must be iid and independent of $Y_2(k)$ for all pairs (i, k). Therefore,

$$\mathrm{E}[\epsilon^* | M_2] = \frac{1}{n}\sum_{i=1}^{n} \mathrm{E}[\epsilon_i | Y_2] = 0.$$

By Lemma 1, G_1 is a mean-preserving spread of G_2.

Conditional on $X = x_q$, the $U_\lambda(i)$ are independently and identically distributed. The pool loss is given as the mean of the $U_\lambda(i)$. Therefore, Lemma 3 implies directly that $H_q(\cdot; 1)$ is a mean-preserving spread of $H_q(\cdot; \lambda)$, which in turn implies that Proposition 1 holds.

References

Basel Committee on Bank Supervision, "International Convergence of Capital Measurement and Capital Standards: A Revised Framework," Publication No. 107, Bank for International Settlements June 2004.

Chow, William, "A Dirichlet Process Mixture Model of Overnight HIBOR," December 1998.

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

Ferguson, Thomas S., "A Bayesian Analysis of Some Nonparametric Problems," *Annals of Statistics*, March 1973, *1* (2), 209–230.

Fishburn, Peter C. and Raymond G. Vickson, "Theoretical Foundations of Stochastic Dominance," in "Stochastic Dominance," Lexington Books, 1978.

Ganapati, Sunita and Gaurav Tejwani, "Structured Credit Strategies: The Case for Structured Finance CDOs," Structured Credit Research, Lehman Brothers, July 2002.

Gordy, Michael and David Jones, "Random Tranches," *Risk*, March 2003, *16* (3), 78–83.

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," *Journal of Financial Intermediation*, July 2003, *12* (3), 199–232.

[22] I am grateful to E. De Giorgi for suggesting this proof.

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

Ishwaran, Hemant and Mahmoud Zarepour, "Dirichlet prior sieves in finite normal mixtures," *Statistica Sinica*, 2002, *12*, 941–963.

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.

Koyluoglu, H. Ugur and Andrew Hickman, "Reconciliable Differences," *Risk*, October 1998, *11* (10), 56–62.

Pykhtin, Michael and Ashish Dev, "Credit Risk in Asset Securitizations: Analytical Model," *Risk*, May 2002, *15* (5), S16–S20.

___ **and** ___ , "Coarse-grained CDOs," *Risk*, January 2003, *16* (1), 113–116.

Ribeiro, Claudia and Nick Webber, "Valuing Path Dependent Options in the Variance-Gamma Model by Monte Carlo with a Gamma Bridge," *Journal of Computational Finance*, Winter 2003, *7* (2).

Rothschild, Michael and Joseph E. Stiglitz, "Increasing Risk I: A Definition," *Journal of Economic Theory*, 1970, *2* (3), 225–243.

Vasicek, Oldrich A., "Limiting Loan Loss Probability Distribution," KMV Corporation, August 1991.