

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Granularity Adjustment for Mark-to-Market Credit Risk Models

Michael B. Gordy and James Marrone

2010-37

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Granularity Adjustment for Mark-to-Market Credit Risk Models*

Michael B. Gordy and James Marrone
Federal Reserve Board

3 June 2010

Abstract

The impact of undiversified idiosyncratic risk on value-at-risk and expected shortfall can be approximated analytically via a methodology known as granularity adjustment (GA). In principle, the GA methodology can be applied to any risk-factor model of portfolio risk. Thus far, however, analytical results have been derived only for simple models of actuarial loss, i.e., credit loss due to default. We demonstrate that the GA is entirely tractable for single-factor versions of a large class of models that includes all the commonly used mark-to-market approaches. Our approach covers both finite ratings-based models and models with a continuum of obligor states. We apply our methodology to CreditMetrics and KMV Portfolio Manager, as these are benchmark models for the finite and continuous classes, respectively. Comparative statics of the GA with respect to model parameters in CreditMetrics reveal striking and counterintuitive patterns. We explain these relationships with a stylized model of portfolio risk.

Keywords: granularity adjustment, idiosyncratic risk, portfolio credit risk, value-at-risk, expected shortfall

*We thank Chris Finger of RiskMetrics Group and Jing Zhang of Moody's KMV for helpful input on modeling details, and Serge Darolles for helpful discussion. The opinions expressed here are our own, and do not reflect the views of the Board of Governors of the Federal Reserve System. Email: (Michael.Gordy@frb.gov) and (James.V.Marrone@frb.gov).

In the portfolio risk-factor frameworks that underpin both industry models of credit value-at-risk (VaR) and the Internal Ratings-Based (IRB) risk weights of Basel II, credit risk in a portfolio arises from two sources, systematic and idiosyncratic. Systematic factors represent the effect of unexpected changes in macroeconomic and financial market conditions on the performance of borrowers. Borrowers may differ in their degree of sensitivity to systematic risk, but few firms are completely insulated from the wider economic conditions in which they operate. Therefore, the systematic component of portfolio risk is unavoidable and only partly diversifiable. Idiosyncratic factors represent the risks that are particular to individual borrowers. As a portfolio becomes more fine-grained, in the sense that the largest individual exposures account for a vanishing share of total portfolio exposure, idiosyncratic risk is diversified away at the portfolio level.

In some settings, including the IRB approach of Basel II, the computation of VaR is dramatically simplified if it is assumed that bank portfolios are *perfectly* fine-grained, that is, that idiosyncratic risk has been fully diversified away, so that portfolio loss depends only on systematic risk. Real-world portfolios are not, of course, perfectly fine-grained. When there are material name concentrations of exposure, there will be a residual of undiversified idiosyncratic risk in the portfolio. The impact of undiversified idiosyncratic risk on VaR can be approximated analytically via a methodology known as *granularity adjustment*. In principle, the granularity adjustment (GA) can be applied to any risk-factor model of portfolio credit risk. Thus far, however, analytical results have been derived only for simple models of *actuarial* loss, i.e., credit loss due to default. The implicit view appears to be that the GA would be tedious to derive, or perhaps even intractable, for the more complicated models of *mark-to-market* credit loss. Large banks typically model credit loss in market value terms, and even the model underpinning the IRB approach of Basel II is in this advanced class.¹ In this paper, we demonstrate that the GA is in fact entirely tractable for a large class of models that includes single-factor versions of all the commonly used mark-to-market approaches. If notation is chosen judiciously, the resulting derivations and calculations are concise and straightforward.

In Section 1, we review the established results in the literature on granularity adjustment and introduce the basic notation. Our general solution for mark-to-market models is given in Section 2. This solution covers both finite ratings-based models and models with a continuum of obligor states. In Section 3, we apply our methodology to CreditMetrics and KMV Portfolio Manager as these are the benchmark models for the finite and continuous classes, respectively. Comparative statics with respect to model parameters are explored in Section 4. Some of the comparative statics appear counterintuitive at first glance, so in Section 5 we explain these results with a stylized model of portfolio risk.

¹The IRB risk-weight formulae for corporate loans (Basel Committee on Bank Supervision, 2006, ¶272) are organized in a way that visually suggests actuarial concepts, but the maturity adjustment maps to capital charges derived in a mark-to-market setting (see Gordy and Lütkebohmert, 2010, §1).

1 Granularity adjustment

For simplicity in exposition, we first consider risk-measurement for a portfolio of n homogeneous positions. We wish to model the portfolio loss rate, \tilde{L} , at a fixed horizon $t = H$ with current time normalized to $t = 0$. Let L_i denote the loss at the horizon on position i (expressed as a percentage of current value), so that the portfolio loss rate is simply

$$\tilde{L} = \frac{1}{n} \sum_{i=1}^n L_i. \quad (1)$$

For a given target solvency probability $q \in (0, 1)$, value-at-risk is defined as the q^{th} percentile of the distribution of loss. Let $\alpha_q(Y)$ denote the q^{th} percentile of the distribution of random variable Y , i.e.,

$$\alpha_q(Y) \equiv \inf\{y : \Pr(Y \leq y) \geq q\}. \quad (2)$$

In terms of this more general notation, VaR is $\alpha_q(\tilde{L})$.

Let X denote the set of systematic risk factors that are realized at the horizon. A critical assumption of all risk-factor portfolio models is that all dependence in loss across positions is due to common dependence on X , so that L_i is independent of L_j when conditioned on X . As n grows to infinity, all idiosyncratic sources of risk vanish, so $|\tilde{L} - \mathbb{E}[\tilde{L}|X]| \rightarrow 0$, almost surely. This implies that $\alpha_q(\tilde{L}) \rightarrow \alpha_q(\mathbb{E}[\tilde{L}|X])$ as $n \rightarrow \infty$. This result is especially useful when X is univariate and conditional expected loss is increasing in X , and we henceforth impose these assumptions. Subject to mild restrictions, $\alpha_q(\mathbb{E}[\tilde{L}|X])$ is equal to $\mathbb{E}[\tilde{L}|\alpha_q(X)]$, which is easily calculated in analytical form.

The difference $\alpha_q(\tilde{L}) - \alpha_q(\mathbb{E}[\tilde{L}|X])$ represents the effect of undiversified idiosyncratic risk in the portfolio. This difference cannot be obtained in analytical form, but we can construct an asymptotic approximation in orders of $1/n$.

$$\alpha_q(\tilde{L}) - \alpha_q(\mathbb{E}[\tilde{L}|X]) = -\frac{1}{n} \frac{1}{2h(\alpha_q(X))} \frac{d}{dx} \left(\frac{\mathbb{V}[L_1|X=x]h(x)}{\frac{d\mathbb{E}[L_1|X=x]}{dx}} \right) \Bigg|_{x=\alpha_q(X)} + o(1/n), \quad (3)$$

where $h(\cdot)$ is the density of X . The dominant term on the right hand side is the granularity adjustment.

The GA extends naturally to heterogeneous portfolios. Let A_i be the current size of exposure i . This is the face value of the instrument in an actuarial setting, and is the current market value in a mark-to-market setting. Let $a_i = A_i / \sum_{j=1}^n A_j$ be the portfolio weights. Imposing minor restrictions on the sequence A_1, A_2, \dots so that the $\sum_{i=1}^n a_i^2 \rightarrow 0$ as $n \rightarrow \infty$ (see Assumption $(\mathcal{A}-2)$)

in Gordy, 2003), we have

$$\text{GA} = \frac{-1}{2h(\alpha_q(X))} \frac{d}{dx} \left(\frac{V[\tilde{L}|X=x]h(x)}{\frac{dE[\tilde{L}|X=x]}{dx}} \right) \Big|_{x=\alpha_q(X)} \quad (4)$$

This form of the GA was first suggested by Wilde (2001). Martin and Wilde (2002) gave a more rigorous derivation of Wilde’s formula based on theoretical work by Gouriéroux, Laurent, and Scaillet (2000). Gordy (2004) presents a survey of these developments and a primer on the mathematical derivation.²

The GA of equation (4) applies under either accounting paradigm for loss.³ Under an actuarial definition, loss L_i on position i is the product of a default indicator for i and the loss given default (LGD) suffered on that position. LGD is expressed as a percentage of exposure and may itself be stochastic. Heretofore, all applications of the GA to portfolio credit risk have been in an actuarial setting. Wilde (2001) provides analytical solutions to equation (4) for the CreditRisk⁺ model and for an actuarial version of the CreditMetrics model. Analysis of CreditMetrics is developed further by Emmer and Tasche (2005). Even for the special case of a homogeneous portfolio and zero recovery on defaulted loans, the Emmer and Tasche solution suggests some complexity. Expressed in the notation to be introduced below, we have

$$\begin{aligned} \text{GA} = & -\frac{1}{n} \frac{1}{2 \frac{\sqrt{\rho}}{\sqrt{1-\rho}} \phi\left(\frac{C_0 - \alpha_q(X)\sqrt{\rho}}{\sqrt{1-\rho}}\right)} \left[\frac{\sqrt{\rho}}{\sqrt{1-\rho}} \phi\left(\frac{C_0 - \alpha_q(X)\sqrt{\rho}}{\sqrt{1-\rho}}\right) \left(1 - 2\Phi\left(\frac{C_0 - \alpha_q(X)\sqrt{\rho}}{\sqrt{1-\rho}}\right)\right) \right. \\ & \left. + \alpha_q(X) + \frac{\sqrt{\rho}}{\sqrt{1-\rho}} \frac{C_0 - \alpha_q(X)\sqrt{\rho}}{\sqrt{1-\rho}} \Phi\left(\frac{C_0 - \alpha_q(X)\sqrt{\rho}}{\sqrt{1-\rho}}\right) \Phi\left(\frac{\alpha_q(X)\sqrt{\rho} - C_0}{\sqrt{1-\rho}}\right) \right] \quad (5) \end{aligned}$$

The original result, in Emmer and Tasche (2005, Remark 2.3), incorrectly has a minus sign in place of the second plus sign on the second line of equation (5). The same sign error is found in the more general result in Proposition 2.2 of that paper. The obscurity of this error, which we believe has not been noticed until now, perhaps reflects the opacity of the formulae.

In a mark-to-market setting, “loss” is an ambiguous concept. One needs to choose a reference point (i.e., the value of the instrument that counts as zero loss) and a convention for discounting to the present. A typical definition is the difference between expected return and realized return, discounted back to today at the riskfree rate. Return is defined as the ratio of market value at the horizon (inclusive of cashflows received during the period $(0, H]$, accrued to the horizon at the

²Gordy and Lütkebohmert (2010) address practical considerations for application to Basel II. Granularity adjustment has also been applied to option pricing (Gagliardini and Gouriéroux, 2009), pricing and risk-measurement of CDOs (Antonov et al., 2005), econometrics (Gouriéroux and Monfort, 2009/10; Gouriéroux and Jasiak, 2008), simulation methods (Gordy and Juneja, forthcoming), and modeling systemic risk contributions in banking systems (Tarashev et al., 2010).

³When applied in a mark-to-market setting, mild additional restrictions are required to bound the conditional second moment of portfolio loss (Gordy, 2003, §3, Assumption $(\mathcal{A} - 1)$).

riskfree rate) to the current market value. We adopt this convention, but note that it is generally trivial to modify our results to accommodate other definitions.

To formalize, let $B_t(T)$ be the money market fund, i.e., $B_t(T)$ is the value at T of a unit of currency invested at date t in a riskless continuously compounded money market fund. We can write this as

$$B_t(T) = \exp\left(\int_t^T r_s ds\right)$$

where r_t is the instantaneous short rate. Portfolio credit risk models generally exclude interest rate risk, so we assume that the path of r_t is deterministic (though not necessarily constant). We multiply intra-horizon cashflows by $B_t(H)$ to accrue to the horizon, and divide by $B_0(H)$ to discount horizon values back to today. Let W_i be the return on position i at the horizon, and define loss as $L_i = (E[W_i] - W_i)/B_0(H)$. Aggregate portfolio return is $\tilde{W} = \sum_{i=1}^n a_i W_i$, and aggregate portfolio loss is

$$\tilde{L} = \frac{1}{B_0(H)}(E[\tilde{W}] - \tilde{W}). \quad (6)$$

Let $\mu_i(x)$ denote the conditional expected return $E[W_i|X = x]$ as a function of x , and similarly define

$$\tilde{\mu}(x) = E[\tilde{W}|X = x] = \sum_{i=1}^n a_i \mu_i(x).$$

With this notation, we can write the asymptotic VaR as

$$\text{VaR}^\infty = E[\tilde{L}|X = \alpha_q(X)] = \frac{1}{B_0(H)}(E[\tilde{\mu}(X)] - \tilde{\mu}(\alpha_q(X))).$$

Let $\sigma_i^2(x)$ be the conditional variance $V[W_i|X = x]$. Due to the conditional independence of the position losses, we can write the conditional variance for portfolio return as

$$\tilde{\sigma}^2(x) = V[\tilde{W}|X = x] = \sum_{i=1}^n a_i^2 \sigma_i^2(x).$$

From equation (6), we have

$$\frac{d}{dx} E[\tilde{L}|X = x] = -\tilde{\mu}'(x)/B_0(H)$$

and

$$V[\tilde{L}|X = x] = \tilde{\sigma}^2(x)/B_0(H)^2$$

so equation (4) can be re-written as

$$\text{GA} = \frac{1}{2} \frac{1}{B_0(H)} \frac{1}{h(\alpha_q(X))} \frac{d}{dx} \left(\frac{\tilde{\sigma}^2(x)h(x)}{\tilde{\mu}'(x)} \right) \Big|_{x=\alpha_q(X)} \quad (7a)$$

This is the form in which we will calculate the GA.

In many commonly-used models, the distribution of X is such that $h'(x)/h(x)$ takes a simple form. This is most notably the case when X is normally distributed (as in the models considered in Section 3), for which we have $h'(x)/h(x) = -x$. Here it can be convenient to apply the product rule to the derivative in the GA formula, and arrive at

$$\text{GA} = \frac{1}{2} \frac{1}{B_0(H)} \left(\frac{\tilde{\sigma}^2(\alpha_q(X)) h'(\alpha_q(X))}{\tilde{\mu}'(\alpha_q(X)) h(\alpha_q(X))} + \frac{d}{dx} \left(\frac{\tilde{\sigma}^2(x)}{\tilde{\mu}'(x)} \right) \Big|_{x=\alpha_q(X)} \right) \quad (7b)$$

We have thus far assumed that value-at-risk is the risk-measure of interest. A popular alternative to VaR is expected shortfall (ES). When portfolio loss has continuous distribution, this is defined as

$$\text{ES}_q[\tilde{L}] = \text{E}[\tilde{L} | \tilde{L} \geq \alpha_q(\tilde{L})] \quad (8)$$

Martin and Tasche (2007) and Gordy (2004) show that the granularity adjustment for ES is

$$\text{GA}^{\text{ES}} = \frac{1}{2} \frac{1}{(1-q)} \frac{\text{V}[\tilde{L} | X = \alpha_q(X)] h(\alpha_q(X))}{\frac{d\text{E}[\tilde{L} | X=x]}{dx} \Big|_{\alpha_q(X)}}$$

which we can re-write as

$$\text{GA}^{\text{ES}} = \frac{-1}{2} \frac{1}{B_0(H)} \frac{h(\alpha_q(X)) \tilde{\sigma}^2(\alpha_q(X))}{(1-q) \tilde{\mu}'(\alpha_q(X))} \quad (9)$$

The computations needed for this expression are a subset of the computations needed for equation (7a), so it is clear that the ES GA can readily be calculated whenever the VaR GA can be calculated.

Finally, in some models conditional expected loss is monotonically *decreasing* in X . The above results continue to hold, but with $\alpha_q(X)$ everywhere replaced by $\alpha_{1-q}(X)$ and the sign on GA^{ES} reversed.

2 Conditional expected return and variance functions

To implement the GA, we require tractable expressions for the conditional expected return and conditional variance of return as functions of the realization of X . We consider first the class of credit risk models in which the condition of obligors at the horizon is represented by a finite state space. This includes the important class of ratings-based models, for which the “state” is the obligor’s rating at the horizon.

Let \mathcal{S} be the set of possible obligor states at the horizon. These states are enumerated as $\mathcal{S} = \{0, 1, 2, \dots, G\}$. Let $S_i \in \mathcal{S}$ be the state for obligor i at the horizon. When the states are S&P rating grades, for example, the obligor has defaulted if $S_i = 0$, migrated to CCC if $S_i = 1$, and so on up to $S_i = G$ for migration to AAA.

In all ratings-based models, the return W_i depends on the horizon rating S_i . More generally, we might expect W_i to be influenced by the systematic factor X , as prevailing spreads for a given rating grade typically increase during a credit market downturn. There may also be idiosyncratic influences on the value at the horizon. Many current models allow for idiosyncratic recovery risk (i.e., random LGD) in the default state, and the models could (in principle) easily be extended to allow for idiosyncratic spread risk in non-default states. In our framework, we allow for all three sources of risk.

We decompose the conditional expected value by further conditioning on horizon state:

$$\mu_i(x) = \sum_{s=0}^G \mathbb{E}[W_i | X = x, S_i = s] \cdot \Pr(S_i = s | X = x).$$

Let $\pi_{is}(x) \equiv \Pr(S_i = s | X = x)$ and $\lambda_{is}(x) \equiv \mathbb{E}[W_i | X = x, S_i = s]$, so we can write

$$\mu_i(x) = \sum_{s=0}^G \pi_{is}(x) \lambda_{is}(x) = \langle \Pi_i(x), \Lambda_i(x) \rangle \quad (10)$$

where $\Pi_i(x)$ is the vector of $\{\pi_{is}(x)\}$, $\Lambda_i(x)$ is the vector of $\{\lambda_{is}(x)\}$, and where $\langle \cdot, \cdot \rangle$ denotes the inner product.

We proceed similarly for the conditional variance:

$$\begin{aligned} \sigma_i^2(x) &= \mathbb{E}[W_i^2 | X = x] - \mathbb{E}[W_i | X = x]^2 = \sum_{s=0}^G \mathbb{E}[W_i^2 | X = x, S_i = s] \cdot \Pr(S_i = s | X = x) - \mu_i(x)^2 \\ &= \sum_{s=0}^G (\mathbb{V}[W_i | X = x, S_i = s] + \mathbb{E}[W_i | X = x, S_i = s]^2) \cdot \Pr(S_i = s | X = x) - \mu_i(x)^2. \end{aligned}$$

Letting $\xi_{is}^2(x) \equiv \mathbb{V}[W_i | X = x, S_i = s]$, we write

$$\begin{aligned} \sigma_i^2(x) &= \sum_{s=0}^G (\xi_{is}^2(x) + \lambda_{is}(x)^2) \pi_{is}(x) - \mu_i(x)^2 \\ &= \langle \Xi_i(x), \Pi_i(x) \rangle + \langle \Lambda_i(x)^2, \Pi_i(x) \rangle - \langle \Lambda_i(x), \Pi_i(x) \rangle^2 \quad (11) \end{aligned}$$

where $\Xi_i(x)$ is the vector of $\{\xi_{is}^2(x)\}$ and $\Lambda_i(x)^2$ is the vector of $\{\lambda_{is}(x)^2\}$.

We now turn to the class of credit risk models in which the obligor state at the horizon can take on a continuum of values. This includes structural approaches based on the Merton (1974) model in which obligor credit risk can be measured by the standardized distance between the obligor's asset value and default threshold. In industry practice, KMV Portfolio Manager is the most widely-used implementation of this approach.

As before, \mathcal{S} is the set of possible obligor states at the horizon. Typically we would have $\mathcal{S} \subseteq \mathfrak{R}$, but that is not strictly necessary for our purposes. Adapting our earlier notation, let $\pi_i(s; x)$ be the conditional probability density function for S_i , and let

$$\begin{aligned}\lambda_i(s; x) &\equiv \text{E}[W_i|X = x, S_i = s] \\ \xi_i^2(s; x) &\equiv \text{V}[W_i|X = x, S_i = s].\end{aligned}$$

When working with this class of models, let $\langle \cdot, \cdot \rangle$ denote the Hermitian inner product, so that

$$\langle \Lambda_i(x), \Pi_i(x) \rangle = \int_{\mathcal{S}} \lambda_i(s; x) \pi_i(s; x) ds.$$

With this notation, the inner product representations of $\mu_i(x)$ and $\sigma_i^2(x)$ in equations (10) and (11) continue to hold.

In both the discrete and continuous state space cases, the derivative of the inner product is given by the usual product rule, e.g.,

$$\frac{d}{dx} \langle \Lambda_i(x), \Pi_i(x) \rangle = \langle \Lambda_i'(x), \Pi_i(x) \rangle + \langle \Lambda_i(x), \Pi_i'(x) \rangle.$$

The derivatives of the $\mu_i(x)$ and $\sigma_i^2(x)$ functions are therefore easily obtained from the derivatives of the constituent $\Pi_i(x)$, $\Lambda_i(x)$ and $\Xi_i(x)$ functions.

3 Application

To apply the results of the previous section to a given model, we need to have tractable and differentiable expressions for the $\Pi_i(x)$, $\Lambda_i(x)$ and $\Xi_i(x)$ functions. For the models most widely-used in practice, these functions are indeed easily obtained and even more easily differentiated. We demonstrate with application to CreditMetrics and KMV Portfolio Manager, as these are the benchmark models for the finite and continuous classes, respectively.

3.1 CreditMetrics

CreditMetrics is perhaps the most widely-known industry model of portfolio credit risk.⁴ The model is loosely styled on the classic structural model of Merton (1974), but is calibrated to credit ratings rather than equity price and balance sheet information. Obligor rating is taken as a sufficient statistic of the term-structure of firm default risk on a single-name basis, and rating transitions are assumed to follow a time-homogeneous Markov chain. This implies that the unconditional

⁴CreditMetrics was first described by Gupton, Finger, and Bhatia (1997), and subsequently has been developed and refined as a vendor model by RiskMetrics Group.

distribution over horizon rating depends only on current rating. We write $\bar{\pi}_{gs}$ for the probability of transition from current grade g to grade s at the horizon.

Associated with obligor i is a latent ‘‘asset return’’ variable Y_i , which is assumed to be distributed $\mathcal{N}(0,1)$. The real line is partitioned into ‘‘bins’’ corresponding to the possible state outcomes in \mathcal{S} . Given current rating g_i , the obligor defaults if $Y_i \leq C_{g(i),0}$, transitions to CCC if $C_{g(i),0} < Y_i \leq C_{g(i),1}$, and so on, for fixed bin threshold values

$$-\infty \equiv C_{g,-1} \leq C_{g,0} \leq \dots \leq C_{g,G} \equiv \infty.$$

For transition probabilities to be consistent with the given $\bar{\pi}_{gs}$, we have

$$C_{g,s} = \Phi^{-1} \left(\sum_{j=0}^s \bar{\pi}_{gj} \right)$$

for $s = 0, \dots, G - 1$.

To induce dependence across obligors, we decompose the asset returns as

$$Y_i = \sqrt{\rho_i}X + \sqrt{1 - \rho_i}\epsilon_i. \quad (12)$$

The systematic factor X is distributed $\mathcal{N}(0,1)$, which implies that $h(x) = \phi(x)$ and $h'(x) = -x\phi(x)$. The idiosyncratic factors ϵ_i are iid $\mathcal{N}(0,1)$ and independent of X . It is straightforward to show (see, e.g., Gordy, 2001) that the conditional probability distribution for S_i is given by

$$\pi_{is}(x) = \Phi \left(\frac{C_{g(i),s} - x\sqrt{\rho_i}}{\sqrt{1 - \rho_i}} \right) - \Phi \left(\frac{C_{g(i),s-1} - X\sqrt{\rho_i}}{\sqrt{1 - \rho_i}} \right). \quad (13)$$

Exploiting the relationship $\phi'(x) = -x\phi(x)$, the derivatives of $\pi_{is}(x)$ are easily obtained.

Relative to the general framework of Section 2, CreditMetrics imposes simplifying assumptions on the distribution of return in each state. Market credit spreads at the horizon are taken as *deterministic* functions of rating grade. In the case of default, there is only *idiosyncratic* risk in recovery. Therefore, for all horizon states,

$$\lambda_{is}(x) = \mathbb{E}[W_i | X = x, S_i = s] = \mathbb{E}[W_i | S_i = s],$$

so we can take $\lambda_{is}(x)$ as a constant λ_{is} . For the conditional variance, we have $\xi_{is}^2(x) = 0$ for all $s \geq 1$. In the default state, $\xi_{i0}^2(x)$ is replaced by the constant $\xi_{i0}^2 = \mathbb{V}[W_i | S_i = 0]$.

The return in the default state is affine in the recovery rate $(1 - \text{LGD})$. For example, say we have a loan with biannual coupons of $c_i/2$, face value of 1, and current value P_{i0} . If we assume that default occurs just before the horizon of $H = 1$ year, then the first coupon is received and the second is accrued into the claim. The return in the default state is therefore $((1 - \text{LGD}_i)(1 + c_i/2) +$

$B_{1/2}(1)c_i/2)/P_{i0}$. In the CreditMetrics model, it is assumed that LGD_i is drawn as an independent beta-distributed random variable with specified mean $ELGD_i$ and variance $VLGD_i^2$, which implies that λ_{i0} is affine in $ELGD_i$ and ξ_{i0}^2 is affine in $VLGD_i^2$. For parsimony in data requirements, it is usually assumed that

$$VLGD_i^2 = \nu \cdot ELGD_i \cdot (1 - ELGD_i) \quad (14)$$

for fixed volatility parameter ν . In this case, ξ_{i0}^2 is affine in ν .

At this point in the analysis, only the state returns λ_{is} remain to be calculated. In the original version of CreditMetrics, as documented by Gupton et al. (1997), pricing at the horizon followed a discounted contractual cashflow approach. For greater internal consistency, later versions of CreditMetrics adopted a modified version of the Hull and White (2000) methodology. In this approach, the term-structures of risk-neutral default probabilities for each grade are backed out from the observed term-structures of ratings-based credit spreads. It is then trivial to obtain prices for each obligor rating at the horizon by summing the discounted expected cashflows. For our purposes in this paper, either methodology (or, indeed, any number of other pricing methodologies) can be used. One must be able to calculate the return in each obligor state in order to implement CreditMetrics, so the calculation of the λ_{is} imposes no burden that is peculiar to granularity adjustment.

Finally, we note that the GA for the actuarial version of CreditMetrics can be obtained as a special case with our general framework. To calculate the actuarial GA, fix $\lambda_{is} = 1$ for all non-default $s \geq 1$, $\lambda_{i0} = 1 - ELGD_i$ and $\xi_{i0}^2 = VLGD_i^2$, and fix both the coupon rate and the riskfree rate to zero. The GA formula of Emmer and Tasche (2005) is a special case of our formula in which ELGD is fixed to 1 and VLGD to zero.

3.2 KMV Portfolio Manager

Like CreditMetrics, Moody's KMV model is based on the classic structural model of Merton (1974). Whereas CreditMetrics takes a stylized approach to the model, KMV is firmly grounded in the substance of the structural relationship between firm asset value and debt performance. The model can be divided into two components. Default prediction (i.e., estimation of the term-structure of firm default probabilities, or "EDFs") is provided by KMV Credit Monitor (Crosbie and Bohn, 2003; Kealhofer, 2003a). Portfolio risk is assessed by KMV Portfolio Manager (Kealhofer and Bohn, 2001). We develop the GA for version 1.4 of KMV Portfolio Manager, and note that the current version of the model may differ in important respects.

The portfolio model takes as input the term-structure of EDFs for each obligor in the portfolio, and we do the same here. Specifically, for each firm i , we take as input parameters the probability of default at or before the horizon ($EDF_{i0,H}$) and the probability of default at or before loan maturity (EDF_{i0,T_i}).

In the structural approach, default occurs when asset value falls short of the fixed liabilities of

the firm at $t = H$. We take as the obligor state variable the log return on firm assets. The asset value is assumed to follow a geometric Brownian motion with drift under the physical measure, so that

$$S_i = (\zeta_i - \eta_i^2/2)H + \eta_i\sqrt{H}Z_i, \quad (15)$$

where ζ_i is the drift, η is the volatility, and Z_i is a shock distributed $\mathcal{N}(0, 1)$. As in CreditMetrics, the shock is decomposed into systematic and idiosyncratic components:

$$Z_i = \sqrt{\rho_i}X + \sqrt{1 - \rho_i^2}\epsilon_i \quad (16)$$

where the systematic factor X and idiosyncratic factors ϵ_i are independent and distributed $\mathcal{N}(0, 1)$.

The domain of possible states at the horizon is the real line. Because the shocks are normally distributed, it is easily seen that the cdf of S_i is

$$\Pr(S_i \leq s) = \Phi\left(\frac{s - (\zeta_i - \eta_i^2/2)H}{\eta_i\sqrt{H}}\right) \quad (17)$$

The default threshold $C_{i,H}$ must therefore satisfy the relationship

$$\Phi\left(\frac{C_{i,H} - (\zeta_i - \eta_i^2/2)H}{\eta_i\sqrt{H}}\right) = EDF_{i0,H}.$$

Loosely speaking, the default threshold represents the log of the firm's fixed liabilities. However, because KMV employs a proprietary mapping from the firm's normalized distance-to-default to EDF, the threshold cannot be determined directly from balance sheet information.

The conditional distribution of S_i is also Gaussian. It is easily seen that the conditional density is

$$\pi_i(s; x) = \phi\left(\frac{s - (\zeta_i - \eta_i^2/2)H - x\eta_i\sqrt{H}\sqrt{\rho_i}}{\eta_i\sqrt{H}\sqrt{1 - \rho_i}}\right). \quad (18)$$

The KMV pricing algorithm differs from that of CreditMetrics, but shares the important assumptions that market value at the horizon is a deterministic function of obligor state for surviving obligors, and that, in the case of default, there is only idiosyncratic risk in recovery. Therefore, for all horizon states,

$$\lambda_i(s; x) = \mathbb{E}[W_i|X = x, S_i = s] = \mathbb{E}[W_i|S_i = s],$$

so we can write $\lambda_i(s; x)$ as $\lambda_i(s)$. For the conditional variance, we have $\xi_i^2(s; x) = 0$ for all $s > C_{i,H}$. In the default states, $\xi_i^2(s; x)$ is replaced by a constant ξ_{i0}^2 . This implies that the inner product $\langle \Xi_i(x), \Pi_i(x) \rangle$ in equation (11) takes the simple form

$$\langle \Xi_i(x), \Pi_i(x) \rangle = \xi_{i0}^2 \cdot \Phi\left(\frac{C_{i,H} - (\zeta_i - \eta_i^2/2)H - x\eta_i\sqrt{H}\sqrt{\rho_i}}{\eta_i\sqrt{H}\sqrt{1 - \rho_i}}\right).$$

The KMV pricing methodology separates future contractual cashflows into riskless and risky components. If recovery were deterministic, we would write for the price P_t of a loan at time t

$$P_t = (1 - \text{LGD}) \cdot \text{RFV}_t + \text{LGD} \cdot \text{RYV}_t(s) \quad (19)$$

where RFV_t is the time- t price on a riskless bond of the same contractual terms and $\text{RYV}_t(s)$ is the time- t price of a zero-recovery risky bond of the same contractual terms for an obligor in state s . When recovery is stochastic but idiosyncratic, then recovery risk is not priced in equilibrium, so equation (19) should continue to hold, but with stochastic LGD replaced by its expectation ELGD.

In the event of default at or before the horizon, the value at the horizon is $(1 - \text{LGD}_i) \cdot \text{RFV}_{iH}$. While this recovery value is stochastic, it is invariant with respect to S_i , so we can write

$$\begin{aligned} \lambda_i(s; x) &= \frac{1}{P_{i0}} (1 - \text{ELGD}_i) \text{RFV}_{iH} \equiv \lambda_{i0} \\ \xi_i^2(s; x) &= \left(\frac{\text{RFV}_{iH}}{P_{i0}} \right)^2 \text{VLGD}_i^2 \equiv \xi_{i0}^2 \end{aligned}$$

for all $s \leq C_{i,H}$ and all x . The recovery variance VLGD_i^2 is specified as in equation (14).

In the event of survival, the return is $P_{iH}(s)/P_{i0}$. We write $P_{iH}(s)$ as a function of horizon state because of the dependence of RYV_{iH} on S_i . The calculation of $\text{RYV}_{iH}(s)$ is detailed in Gordy, Heitfield, and Jones (in progress). As we have noted in the context of the CreditMetrics model, one must be able to calculate the return in each obligor state in order to implement the portfolio model, so the calculation of the $\lambda_i(s)$ imposes no burden that is peculiar to granularity adjustment.

4 Comparative Statics

In this section, we explore the comparative statics of the granularity adjustment with special emphasis on the parameters that do not appear under the actuarial paradigm. For the sake of clarity, we adopt the CreditMetrics model in a stylized setting with two non-default rating grades ($G = 2$). We consider a portfolio that is homogeneous in all respects other than initial credit rating, i.e., all loans are of equal size and have the same ELGD and VLGD^2 , and all obligors have the same asset correlation ρ . If all obligors were of the same initial rating as well, then we know from equation (3) that the GA can be written as β/n for β that depends on model parameters but not on n . When obligors are not of the same initial rating, the GA can still be written as β/n if we fix the share of each rating grade in the portfolio. We present the comparative statics in terms of β to avoid dependence on the choice of n .

We parameterize the matrix of unconditional transition probabilities (under the physical measure) as in Table 1. Default probabilities are $\bar{\pi}_{A0}$ and $\bar{\pi}_{B0}$. In our baseline parameterization, we set $\bar{\pi}_{A0}$ to 15 basis points (bp) and $\bar{\pi}_{B0}$ to 300bp, so that grades A and B represent the investment

and speculative grades, respectively. Conditional on survival, the probability of remaining in the initial grade is θ_g . Agency ratings are known to be “sticky” (Altman and Rijken, 2004; Löffler, 2004), so we set $\theta_A = \theta_B = 0.9$.

Table 1: Transition probabilities

	A	B	D
A	$(1 - \bar{\pi}_{A0})\theta_A$	$(1 - \bar{\pi}_{A0})(1 - \theta_A)$	$\bar{\pi}_{A0}$
B	$(1 - \bar{\pi}_{B0})(1 - \theta_B)$	$(1 - \bar{\pi}_{B0})\theta_B$	$\bar{\pi}_{B0}$

Our baseline portfolio is composed of equal numbers of grade A and grade B loans. Each loan has face value 1, ELGD of 50%, and maturity of 3 years. Coupons are paid biannually. In the event of default, it is assumed that the first coupon is received in full, and the second coupon is accrued into the legal claim in bankruptcy. Asset correlation is fixed at $\rho = 0.2$. The riskfree rate is a constant $r = 5\%$ and the variance parameter for LGD is $\nu = 0.25$. The horizon is $H = 1$ year and the target solvency probability is $q = 99.9\%$.

For consistency with the current generation of CreditMetrics, we use the pricing approach of Hull and White (2000). This requires that we have for each obligor the term-structure of risk-neutral cumulative default probabilities, which in practical application is extracted from the observed term-structure of credit spreads. For our purposes, a parametric approach is preferable, so we obtain the risk-neutral term structure of default probabilities by adding a parameterized risk-premium to the term structure of default probabilities under the physical measure. Since the CreditMetrics model assumes that ratings follow a time-homogeneous Markov process, we can take powers of the transition matrix in Table 1 to obtain the physical cumulative default probability at any horizon. Let $\bar{\pi}_g(t, T)$ denote the physical probability of default between time t and time T for an obligor in grade g at time t . The Markovian structure of the model implies that $\bar{\pi}_g(t, T) = \bar{\pi}_g(0, T - t)$. We convert to risk-neutral probabilities $\bar{\pi}_g^*(t, T)$ as in the KMV model:

$$\bar{\pi}_g^*(t, T) = \Phi \left(\Phi^{-1}(\bar{\pi}_g(t, T)) + \psi \sqrt{T - t} \sqrt{\rho} \right),$$

where ψ is the “market Sharpe ratio” that determines risk premia. The economic rationale for this specification is put forth by Kealhofer (2003b) and Agrawal et al. (2004). We follow KMV in setting a baseline value of $\psi = 0.4$.

Our comparative statics are *total* derivatives. As parameter values change, the par coupon for the loans will change as well. To take the total derivative of β with respect to, say, $\bar{\pi}_{A0}$, we maintain the initial par value of each loan by changing the coupon to its par value for each value of $\bar{\pi}_{A0}$. This approach is most consistent with economic intuition.

We first explore sensitivity of asymptotic VaR and the GA to parameters that appear in both

the MtM and actuarial models, starting with default probabilities $(\bar{\pi}_{A0}, \bar{\pi}_{B0})$. In the upper panel of Figure 1a, we vary $\bar{\pi}_{A0}$ from zero to 300bp while holding $\bar{\pi}_{B0}$ fixed to its baseline value of 300bp. In the lower panel, we vary $\bar{\pi}_{B0}$ from 15bp to 1500bp while holding $\bar{\pi}_{A0}$ fixed to its baseline value of 15bp. As we should expect, asymptotic VaR, denoted VaR^∞ and defined as $E[\tilde{L}|X = \alpha_q(X)]$, is increasing monotonically with both default probabilities under both actuarial and MtM paradigms.⁵ VaR is larger in the MtM setting because it captures migration risk and loss of coupon income in the event of default. When $\bar{\pi}_{A0} = \bar{\pi}_{B0}$, migration risk is eliminated, so the two views of risk are nearly equivalent.

Comparative statics for the GA are displayed in Figure 1b. Except at low values of $\bar{\pi}_{A0}$, the GA increases in the default probabilities. The intuition, which we will develop in greater detail in Section 5, is that extreme losses in the finite portfolio case are most likely to be generated by a combination of unfavorable systematic and idiosyncratic draws, rather than by systematic risk alone. Default events induce larger loss than downgrades, so the idiosyncratic effect will manifest as a higher than conditionally expected default rate. This implies that VaR is more sensitive to default risk than asymptotic VaR^∞ , and therefore that the gap between them should increase with $\bar{\pi}$.

This intuition breaks down when the grade A default probability is very low. For small $\bar{\pi}_{A0}$, it takes a large (and therefore unlikely) idiosyncratic shock to cause a grade A firm to default even when $X = \alpha_{1-q}(X)$. Consequently, for any given realization of the set of idiosyncratic shocks (i.e., the $\{\epsilon_i\}$ of equation (12)), portfolio loss is largest for permutations that disproportionately assign negative shocks to grade B firms. As the scenarios associated with VaR in the finite portfolio case will be disproportionately driven by grade B defaults, VaR must be less sensitive to $\bar{\pi}_{A0}$ than asymptotic VaR^∞ . For this counterintuitive effect to be observed, we must have positive portfolio shares in at least two non-default grades. If we have only a single non-default grade, then β is strictly increasing with the unconditional default probability.

The relative impact of the GA is plotted in Figure 1c. For a portfolio of $n = 1000$, define the relative impact as the ratio of the GA to VaR, where the portfolio VaR is approximated as the sum of asymptotic VaR^∞ and the GA, i.e.,

$$\text{relative impact} = 100 \cdot \frac{\text{GA}_n}{\text{VaR}^\infty + \text{GA}_n} \quad (20)$$

The lower the default probabilities, the greater the relative importance of sampling variation on portfolio risk, so the larger is the GA as a share of VaR.

The effect of portfolio quality on asymptotic VaR^∞ and the GA is consistent with the comparative statics for default probabilities. We find that VaR^∞ and the GA both fall with the share

⁵For the actuarial paradigm, we plot asymptotic “unexpected loss,” defined as $E[\tilde{L}] - \alpha_q(\tilde{L})$, rather than a quantile of the actuarial loss distribution. This is more consistent with the MtM notion of VaR as a quantile of the de-meaned return.

of grade A in the portfolio. The MtM VaR^∞ exceeds the actuarial VaR^∞ at all values of the investment-grade share, but the actuarial GA exceeds the MtM GA. On a relative impact basis, we find that the size of the GA is increasing with the share of grade A.

Comparative statics for recovery rates, shown in Figure 2, are also similar to those for default probabilities. Both asymptotic VaR^∞ and the GA increase with ELGD. The relationships are exactly linear in the actuarial setting, and nearly non-linear in the MtM case (i.e., there is a slight non-linearity due to the effect of ELGD on par spread). The GA increases with ELGD because ELGD controls the magnitude of loss in the default state and, as we have just observed, the finite-portfolio VaR is more sensitive than the asymptotic VaR^∞ to default risk.

Comparative statics for asset correlation are relatively straightforward. As shown in the upper panel of Figure 3, asymptotic VaR is strictly increasing in ρ . At $\rho = 0$, all risk is diversifiable, so VaR^∞ is zero. At $\rho = 1$, asset returns are comonotonic. So long as $\bar{\pi}_{A0} > 1 - q$, all borrowers default in the state $X = \alpha_q(X)$, and VaR^∞ is determined primarily by ELGD. Between these two extremes, VaR^∞ increases monotonically.

The effect of asset correlation on the GA runs in the opposite direction, as seen in the bottom panel of Figure 3. The greater is ρ , the smaller the impact of idiosyncratic risk on asset returns, so the smaller the contribution of idiosyncratic risk to VaR. As ρ falls to zero, one can show analytically that β tends to infinity. As ρ increases to one, β can tend to negative infinity (when $\bar{\pi}_{B0} \geq \bar{\pi}_{A0} > 1 - q$ and $\text{VLGD}^2 > 0$) or to zero (when $1 - q > \bar{\pi}_{B0} \geq \bar{\pi}_{A0}$ or when $\text{VLGD}^2 = 0$), or even to positive infinity (in a subset of the remaining cases). At these endpoints, the asymptotic series underpinning equation (3) diverges, so the first-order GA becomes an unreliable measure of the gap between VaR and asymptotic VaR^∞ .⁶ Nonetheless, negative values for β are not just an artifact. When ρ is near one, the density of the loss distribution becomes multimodal, and in this circumstance asymptotic VaR^∞ can exceed VaR. Martin and Tasche (2007) explain this phenomenon as a concomitant of the failure of sub-additivity in VaR and prove that the granularity adjustment for expected shortfall is always positive.

We can assess the impact of recovery risk through the comparative static with respect to ν . Asymptotic VaR is invariant with respect to idiosyncratic recovery risk, so is invariant with respect to ν . The GA is linear and increasing in VLGD^2 , so also is linear and increasing in ν . The effect is generally large. In our baseline example, the slope $d\beta/d\nu$ is 1.004 for the MtM model and 1.092 for the actuarial model.

We now turn to the parameters that influence risk under the mark-to-market paradigm, but not the actuarial model. The parameter θ controls the degree of stickiness in ratings, conditional on survival to the horizon. In an actuarial setting, asymptotic VaR^∞ and the GA are invariant with respect to non-default transition likelihood. The comparative statics in the MtM setting are somewhat complicated and perhaps surprising. Consider first the effect of varying θ_B on asymptotic

⁶Gordy (2004) explores this convergence issue with a stylized example.

VaR^∞ . As θ_B increases, the conditional and unconditional probabilities of transition from B to A both fall towards zero. Because the conditional probability is already low (i.e., one expects few upgrades when conditioning on a bad systematic draw), the unconditional probability falls by more than the conditional probability, as seen in Figure 4. This closes the gap between $E[\tilde{W}]$ and $\alpha_q(\tilde{W})$, so reduces VaR^∞ (lower panel of Figure 5a).

For θ_A , the story is somewhat more complicated. As seen in Figure 4, the gap between the conditional and unconditional probability of transition from A to B is fairly constant over $0.5 < \theta_A < 0.8$ and then converges rapidly to zero as θ_A increases towards one. In the lower range, the effect on VaR^∞ is dominated by the indirect effect on par coupon rates: as θ_A increases, the par coupon for grade A falls, so the return λ_{AB} associated with downgrade to B is reduced.⁷ The probability weight on λ_{AB} is greater under the conditional distribution than the unconditional, so the gap between $E[\tilde{W}]$ and $\alpha_q(\tilde{W})$ widens. However, in the upper range of θ_A values, the rapid convergence of $\pi_{BA}(\alpha_q(X))$ towards $\bar{\pi}_{BA}$ dominates, and this causes VaR^∞ to decrease with θ_A . This non-monotonic behavior is observed in the upper panel of Figure 5a.

The comparative statics for the GA as a function of θ , displayed in Figure 5b, are the mirror image of the comparative statics for VaR. The intuition is similar to the explanation for the comparative statics with respect to the $\bar{\pi}$. Relative to asymptotic VaR^∞ , finite portfolio VaR is more sensitive to default risk and less sensitive to migration risk. This implies that the GA will increase (decrease) with θ whenever VaR^∞ decreases (increases) with θ .

Loan maturity increases the sensitivity of returns to rating migration, so asymptotic VaR^∞ increases with maturity. At long maturities, the return distribution reflects the long-run steady-state of the rating process, so the relationship becomes flat. This is seen in Figure 6, where we plot VaR against maturity (log-scale) in the upper panel. Similar to the phenomenon observed in the comparative statics for θ , and indeed for the very same reason, the comparative statics for GA (lower panel) with respect to maturity are the mirror image of the comparative statics for asymptotic VaR^∞ .

Comparative statics with respect to the market Sharpe ratio parameter follow the same logic. The higher the risk premium ψ , the larger the loss associated with downward migration, so the higher the asymptotic VaR^∞ (upper panel of Figure 7). Parallel to the pattern observed for θ and maturity, the comparative statics for the GA with respect to ψ are the mirror image of the comparative statics for VaR^∞ (lower panel).

Finally, comparative statics for the riskfree rate are quite straightforward. Both VaR and the GA decline in near linear fashion with the money market return $B_0(H)$, because the riskfree rate has a minimal effect on valuation at the horizon.

⁷The par coupon for grade B also falls, but by a smaller amount. Of all the transition returns, λ_{AB} is most sensitive to θ_A .

5 A beta-trinomial model of portfolio risk

For parameters governing migration risk in CreditMetrics, we have observed a curious “mirror image” pattern, whereby the comparative static for the GA is of opposite sign to the comparative static for asymptotic VaR[∞]. For parameters governing default risk, by contrast, both VaR[∞] and the GA are increasing (except in one corner of the parameter space). In this section, we shed light on these phenomena using a stylized model of portfolio risk. The comparative statics of this simple model lack the complexity and nuance of the patterns in CreditMetrics, but the most salient characteristics are preserved.

We assume a homogeneous portfolio of n positions and specify the return on position i as

$$W_i = c + Y_i + Z_i$$

where c is a constant, Y_i is a return associated with a discrete state variable, and Z_i is distributed $\mathcal{N}(0, \xi^2)$. The three discrete states represent default ($S_i = 0$), downgrade ($S_i = 1$), and unchanged rating ($S_i = 2$). The value of Y_i in the three states is

$$Y_i = \begin{cases} -\lambda_0 & \text{if } S_i = 0, \\ -\lambda_1 & \text{if } S_i = 1, \\ 0 & \text{if } S_i = 2, \end{cases}$$

so that the vector of state-contingent expected returns is $\Lambda = \{c - \lambda_0, c - \lambda_1, c\}$. We include the Z_i shock merely to ensure a continuous loss distribution. The interest rate is fixed to zero, so that $B_0(H) = 1$. We fix the target solvency probability to $q = 99.9\%$.

Let the risk factor X be distributed Beta(p_1, p_2) on the unit interval. Conditional on $X = x$, the state probabilities for Y_i are $\Pi(x) = \{(1-x)^2, x(1-x), x\}$. The Z_i are assumed to be mutually independent and independent of X and all other risks. Conditional on (X, S_i) , Z_i is the only source of uncertainty, so $\Xi = \{\xi^2, \xi^2, \xi^2\}$.

If we assume for parsimony that investors are risk-neutral, then in equilibrium the constant c is chosen so that $\mathbb{E}[W_i] = 0$. This is solved analytically as

$$c = -\mathbb{E}[Y_i] = \lambda_0 \mathbb{E}[\pi_0(X)] + \lambda_1 \mathbb{E}[\pi_1(X)] = \frac{\lambda_0 p_2 (p_2 + 1) + \lambda_1 p_1 p_2}{(p_1 + p_2)(p_1 + p_2 + 1)} \quad (21)$$

where the last equality follows from the moments of the beta distribution. When we impose this choice of c , portfolio loss is simply $\tilde{L} = -\tilde{W}$.

Besides the portfolio size n , the model has only five parameters: $\lambda_0, \lambda_1, p_1, p_2, \xi$. The parameter of greatest interest is λ_1 , because it determines the magnitude of loss in the downgrade state ($S_i = 1$). Parameter λ_0 determines the magnitude of loss in the default state ($S_i = 0$), and corresponds

most directly to ELGD in CreditMetrics. We restrict $0 \leq \lambda_1 \leq \lambda_0$, so that default induces larger loss than downgrade. Parameters (p_1, p_2) jointly control the distribution of systematic risk. The higher is p_1 relative to p_2 , the higher the probability of the “good state,” $S_i = 2$. Increasing both parameters in proportion leaves the expected value of X unchanged but shrinks its variance. In our examples below, we fix $p_2 = 1$ and take $p_1 = 5$ as our baseline value. The final parameter has the narrow purpose of smoothing the return distribution. We aim to choose the smallest value that will be sufficient to eliminate humps in the density function. In our examples below, we fix $\xi = 0.03$.

The model yields analytical solutions for the GA and for the moments, density and cdf of the return distribution. It is already clear that we have simple expressions for the conditional mean $\tilde{\mu}(x)$ and variance $\tilde{\sigma}^2(x)$ and for the derivatives of these functions. For the beta density for X , we have

$$\frac{h'(x)}{h(x)} = \frac{p_1 - 1}{x} - \frac{p_2 - 1}{1 - x},$$

so it is convenient to use equation (7b) for the GA. To obtain the loss distribution, observe that

$$\tilde{W} = c + \tilde{Y} + \tilde{Z}$$

where $n\tilde{Y}$ is a weighted sum of a conditionally trinomial vector with n trials and conditional probabilities given by $\Pi(x)$, and where $\tilde{Z} \sim \mathcal{N}(0, \xi^2/n)$. Let N_s be the number of positions in state s at the horizon, so that $N_0 + N_1 + N_2 = n$, and let $\chi(n_0, n_1)$ denote the unconditional joint probability of $N_0 = n_0, N_1 = n_1, N_2 = n - n_0 - n_1$. For $n_0 \geq 0, n_1 \geq 0$ and $n_0 + n_1 \leq n$, this is given by

$$\begin{aligned} \chi(n_0, n_1) &= \Pr(N_0 = n_0, N_1 = n_1) \\ &= \binom{n}{n_0, n_1, n - n_0 - n_1} \mathbb{E}[(1 - X)^{n_0} (X(1 - X))^{n_1} X^{n - n_0 - n_1}] \\ &= \frac{n!}{n_0! n_1! (n - n_0 - n_1)!} \frac{B(n - n_0 + p_1, 2n_0 + n_1 + p_2)}{B(p_1, p_2)} \end{aligned}$$

where $B(\cdot, \cdot)$ is the beta function. Because \tilde{Y} and \tilde{Z} are independent, the return distribution is the convolution

$$F(w) = \sum_{n_0 + n_1 \leq n} \chi(n_0, n_1) \Phi \left(\frac{n(w - c) + \lambda_0 n_0 + \lambda_1 n_1}{\xi \sqrt{n}} \right) \quad (22)$$

The density of \tilde{W} and the moment generating function follow trivially.

In Figure 8, we plot the density of the return distribution under our baseline parameter assumptions. The coefficients of skewness and kurtosis are -2.3 and 10.1, respectively, which is qualitatively suitable for the distribution of log-return in a credit portfolio.⁸ If we decrease p_1 or increase p_2 ,

⁸In this model, \tilde{W} is comparable to a log-return, rather than an absolute return as assumed up to now. From a

the distribution would flatten out and become less asymmetric.

Asymptotic VaR for this model is easily obtained:

$$\begin{aligned} \text{VaR}^\infty &= \mathbb{E}[\tilde{L}|X = \alpha_{1-q}(X)] = \lambda_0 \pi_0(\alpha_{1-q}(X)) + \lambda_1 \pi_1(\alpha_{1-q}(X)) - c \\ &= \lambda_0 \left((1 - \alpha_{1-q}(X))^2 - \frac{p_2(p_2 + 1)}{(p_1 + p_2)(p_1 + p_2 + 1)} \right) \\ &\quad + \lambda_1 \left(\alpha_{1-q}(X)(1 - \alpha_{1-q}(X)) - \frac{p_1 p_2}{(p_1 + p_2)(p_1 + p_2 + 1)} \right) \end{aligned} \quad (23)$$

where in the last equality we substitute the equilibrium value for c . VaR^∞ increases with λ_j ($j = 0, 1$) if and only if $\pi_j(\alpha_{1-q}(X)) > \mathbb{E}[\pi_j(X)]$. This condition is easily satisfied for the default state ($j = 0$), but holds only within a range of q values for the downgrade state ($j = 1$). This is because downgrade is the intermediate state between extreme outcomes. For x large, the “unchanged rating” state $S_i = 2$ dominates at $X = x$. As x falls, probability mass is shifted both to $S_i = 1$ and $S_i = 0$. However, at still lower values of x , there is too little mass left on $S_i = 2$, so further increases in the probability of $S_i = 0$ come at the expense of the state $S_i = 1$. If the distribution of X is roughly symmetric, $\pi_1(x) = x(1 - x)$ peaks near $x = \mathbb{E}[X]$, and so the unconditional likelihood of the downgrade state $S_i = 1$ is greater than the conditional likelihood given $X = \alpha_{1-q}(X)$. Negative skew in X increases $\mathbb{E}[X]$ and $\alpha_{1-q}(X)$, which in turn reduces $\mathbb{E}[\pi_1(X)]$ and increases $\pi_1(\alpha_{1-q}(X))$. Fixing $p_2 = 1$, we need $p_1 > 3.8$ to guarantee that VaR^∞ increases with λ_1 at the tail probability $q = 0.999$. The restrictions are discussed in more detail in Appendix A.

Comparative statics for the GA are depicted in Figure 9.⁹ As shown in the upper panel, β is increasing with λ_0 . The relationship for λ_1 , shown in the lower panel, is non-monotonic. The slope β is decreasing with λ_1 at low values of λ_1 and increasing at higher values of λ_1 . In practical application, λ_1 takes on low values (less than the baseline value of 0.2) because loss due to downgrade is generally much smaller than loss due to default.¹⁰ Thus, as observed for CreditMetrics, the GA in this model can increase with default loss and decrease with migration loss while asymptotic VaR^∞ is increasing in both parameters.

The GA is decreasing in λ_1 because migration risk has a smaller impact on VaR for the finite portfolio than it does for the asymptotic portfolio. To see this, consider first the probability of downgrade, $S_i = 1$, conditional on a given level of portfolio loss, i.e., $\tilde{L} = \ell$. For the finite portfolio

risk-management perspective, it makes no difference whether we define loss in terms of the absolute or log return.

⁹We have checked that the β of the first-order GA is indeed a close approximation of the gap between VaR and asymptotic VaR^∞ at the modest portfolio size of $n = 500$. At the baseline parameter values, the relative error of the GA is under 0.2%.

¹⁰For grade A borrowers in the exercises of Section 4, the implied ratio of λ_1 to λ_0 is less than 0.06 under the baseline parameter values.

case,

$$\begin{aligned} \mathbb{E}[\pi_1(X)|\tilde{L} = \ell] &= \mathbb{E}[\pi_1(X)|\tilde{W} = -\ell] = \frac{1}{n}\mathbb{E}[N_1|\tilde{W} = -\ell] \\ &= \sum_{i=0}^n \frac{i}{n} \Pr(N_1 = i|\tilde{W} = -\ell) = \frac{\sum_{i=0}^n \frac{i}{n} \sum_{n_0=0}^{n-i} \chi(n_0, i) \phi\left(\frac{n(-\ell-c)+\lambda_0 n_0+\lambda_1 i}{\xi\sqrt{n}}\right)}{\sum_{i=0}^n \sum_{n_0=0}^{n-i} \chi(n_0, i) \phi\left(\frac{n(-\ell-c)+\lambda_0 n_0+\lambda_1 i}{\xi\sqrt{n}}\right)} \end{aligned}$$

where the last equality is easily derived using Bayes' rule. A similar expression can be derived for the conditional default probability, $\mathbb{E}[\pi_0(X)|\tilde{L} = \ell]$. For the asymptotic case, there is a one-to-one mapping between a given level of loss and a given quantile of X . Recalling the notation $\tilde{\mu}(x) = -\mathbb{E}[\tilde{L}|X = x]$, we can write the conditional state probability as $\pi_s(\tilde{\mu}^{-1}(-\ell))$.

Under baseline parameter values, we have that

$$\begin{aligned} \mathbb{E}[\pi_0(X)|\tilde{L} = \ell] &> \pi_0(\tilde{\mu}^{-1}(-\ell)) \\ \mathbb{E}[\pi_1(X)|\tilde{L} = \ell] &< \pi_1(\tilde{\mu}^{-1}(-\ell)) \end{aligned}$$

in the tail of the loss distribution. The gaps between finite portfolio and asymptotic conditional probabilities of default and migration are plotted in Figure 10 in the neighborhood of the asymptotic VaR^∞ . The intuition is that the presence of idiosyncratic risk in the finite portfolio implies that a tail loss event need not be ascribed exclusively to an unfavorable tail realization of X . Rather, it is most likely that the residual idiosyncratic risk also contributed to loss, i.e., that the realization of X was less unfavorable than $\tilde{\mu}^{-1}(-\ell)$ and that there were more defaults and fewer non-defaults than conditionally expected.

Now consider that the conditional downgrade probability is decreasing in ℓ in the tail of the loss distribution, as we see in the upper panel of Figure 11. As noted earlier, this is because the conditional default probability (plotted in the lower panel) “steals” mass from the conditional downgrade probability at extreme levels of loss. Thus, we have

$$\mathbb{E}[\pi_1(X)|\tilde{L} = \alpha_q(\tilde{L})] < \mathbb{E}[\pi_1(X)|\tilde{L} = \text{VaR}^\infty] < \pi_1(\alpha_{1-q}(X))$$

and the opposite inequality for conditional default probabilities. This implies that VaR is less sensitive to the return on the downgrade state than is asymptotic VaR^∞ . Since the GA is an approximation to the difference between VaR and VaR^∞ , the GA must be decreasing in λ_1 . By the same logic, VaR is more sensitive to the return on the default state than is asymptotic VaR^∞ , so the GA is increasing in λ_0 .

Conclusion

Granularity adjustment is useful as a gauge of how well a bank has diversified idiosyncratic risk. The results of this paper ease the way for application of the GA methodology to the mark-to-market models that are favored by more sophisticated financial institutions. We have demonstrated that the GA is analytically tractable for a large class of mark-to-market models of portfolio credit risk. This class is restrictive in imposing a single systematic factor, but in other respects is much more general than the models observed in practical application. In particular, we allow in our analysis for spreads at the horizon to depend on the realization of the systematic factor.

We derive explicit expressions for the GA for CreditMetrics and KMV Portfolio Manager. As an application, we explore the comparative statics of the GA in CreditMetrics, and find relationships that are sometimes non-monotonic and sometimes counterintuitive. In particular, we observe that the comparative statics for the GA with respect to transition probabilities, maturity, and the market risk premium are essentially mirror images of the corresponding comparative statics for asymptotic VaR^∞ . We have argued that this phenomenon has a single explanation: The presence of idiosyncratic risk in the finite portfolio implies that a tail loss event need not be ascribed exclusively to an unfavorable tail realization of X . Rather, extreme losses in the finite portfolio case are most likely to be generated by a combination of unfavorable systematic and idiosyncratic draws. Default events induce larger loss than downgrades, so the idiosyncratic effect will manifest as a higher than conditionally expected default rate. This implies that VaR is more sensitive to default risk and less sensitive to migration risk than asymptotic VaR^∞ . As the GA is the difference between VaR and VaR^∞ , the comparative statics for the GA with respect to migration risk parameters must be opposite in sign to the corresponding comparative statics for asymptotic VaR^∞ .

In the absence of an analytical expression for the GA, this phenomenon would have been difficult to uncover. Estimation of the GA by simulation is difficult enough, because simulation noise tends to swamp the small gap between VaR and asymptotic VaR^∞ . Clean simulation-based estimates of the comparative statics would have been even more challenging.

A Parameter restrictions in the beta-trinomial model

In this appendix, we obtain parameter restrictions sufficient to guarantee that asymptotic VaR^∞ will increase with the state return parameters. The necessary and sufficient condition for VaR^∞ to increase with λ_j is

$$\pi_j(\alpha_{1-q}(X)) > E[\pi_j(X)]. \quad (24)$$

For $j = 0$, this is satisfied iff

$$(1 - \alpha_{1-q}(X))^2 > E[(1 - X)^2] = \frac{p_2(p_2 + 1)}{(p_1 + p_2)(p_1 + p_2 + 1)},$$

This holds for $q > q_0^*(p_1, p_2)$, where q_0^* is a threshold depending on the distributional parameters. We find that q_0^* increases in p_1 and decreasing in p_2 , i.e., that large negative skew in X implies high q_0^* . For the case of $p_2 = 1$, the beta cdf simplifies as $H(x) = x^{p_1}$, so $\alpha_{1-q}(X) = (1 - q)^{1/p_1}$. It is then easily shown that

$$q_0^*(p_1, 1) = 1 - \left(1 - \sqrt{\frac{2}{(p_1 + 1)(p_1 + 2)}}\right)^{p_1} \leq 1 - \exp(-\sqrt{2}) \approx 0.757.$$

For the downgrade state ($j = 1$), the left and right hand sides of condition (24) are non-monotonic, i.e.,

$$\alpha_{1-q}(X)(1 - \alpha_{1-q}(X)) > \mathbb{E}[X(1 - X)] = \frac{p_1 p_2}{(p_1 + p_2)(p_1 + p_2 + 1)}.$$

This holds only for q bounded between thresholds (q_1^-, q_1^+) . For the case of $p_2 = 1$, it is easily shown that

$$q_1^\pm(p_1, 1) = 1 - \left(\frac{1}{2} \mp \sqrt{\frac{1}{4} - \frac{p_1}{(p_1 + 1)(p_1 + 2)}}\right)^{p_1}$$

Both q_1^- and q_1^+ increase as p_1 increases. For tail probabilities near $q = 0.999$, the binding constraint is the upper bound q_1^+ , which requires that we push p_1 well above p_2 .

References

- Deepak Agrawal, Navneet Arora, and Jeffrey Bohn. Parsimony in practice: An EDF-based model of credit spreads. *Modeling methodology*, Moody's KMV, April 2004.
- Edward I. Altman and Herbert A. Rijken. How rating agencies achieve rating stability. *Journal of Banking and Finance*, 28(11):2679–2714, November 2004.
- A. Antonov, S. Mechkov, and T. Misirpashaev. Analytical techniques for synthetic CDOs and credit default risk measures. Technical report, NumeriX, May 2005.
- Basel Committee on Bank Supervision. Basel II: International convergence of capital measurement and capital standards: A revised framework. Publication No. 128, Bank for International Settlements, June 2006.
- Peter J. Crosbie and Jeff Bohn. Modeling default risk. Technical report, KMV, January 2003.
- Susanne Emmer and Dirk Tasche. Calculating credit risk capital charges with the one-factor model. *Journal of Risk*, 7(2):85–101, Winter 2005.
- Patrick Gagliardini and Christian Gouriéroux. Approximate derivative pricing for large class of homogeneous assets with systematic risk. February 2009.
- Michael B. Gordy. What wags the tail? Identifying the key assumptions in models of portfolio credit risk. In Carol Alexander, editor, *Mastering Risk*, volume 2: Applications. Financial Times Prentice Hall, 2001.
- Michael B. Gordy. A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12(3):199–232, July 2003.
- Michael B. Gordy. Granularity adjustment in portfolio credit risk measurement. In Giorgio P. Szegö, editor, *Risk Measures for the 21st Century*. John Wiley & Sons, 2004.
- Michael B. Gordy and Sandeep Juneja. Nested simulation in portfolio risk measurement. *Management Science*, forthcoming.
- Michael B. Gordy and Eva Lütkebohmert. Granularity adjustment for Basel II. January 2010.
- Michael B. Gordy, Erik Heitfield, and David Jones. Maturity adjustment in Basel II. in progress.
- Christian Gouriéroux and Joann Jasiak. Granularity adjustment for default risk factor model with cohorts. September 2008.
- Christian Gouriéroux and Alain Monfort. Granularity in a qualitative factor model. *Journal of Credit Risk*, 5(4):29–61, Winter 2009/10.
- Christian Gouriéroux, Jean-Paul Laurent, and Olivier Scaillet. Sensitivity analysis of values at risk. *Journal of Empirical Finance*, 7:225–245, 2000.
- Greg M. Gup-ton, Christopher C. Finger, and Mickey Bhatia. *CreditMetrics–Technical Document*. J.P. Morgan & Co., New York, April 1997.

- John Hull and Alan White. Valuing credit default swaps I: No counterparty default risk. *Journal of Derivatives*, 8(1):29–40, Fall 2000.
- Stephen Kealhofer. Quantifying credit risk I: Default prediction. *Financial Analysts Journal*, 59(1):30–44, Jan–Feb 2003a.
- Stephen Kealhofer. Quantifying credit risk II: Debt valuation. *Financial Analysts Journal*, 59(3):78–92, May–Jun 2003b.
- Stephen Kealhofer and Jeffrey R. Bohn. Portfolio management of default risk. Technical report, KMV Corporation, May 2001.
- Gunter Löffler. An anatomy of rating through the cycle. *Journal of Banking and Finance*, 28(3):695–720, March 2004.
- Richard Martin and Dirk Tasche. Shortfall: a tail of two parts. *Risk*, 20(2):84–89, February 2007.
- Richard Martin and Tom Wilde. Unsystematic credit risk. *Risk*, 15(11):123–128, November 2002.
- Robert C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470, May 1974.
- Nikola Tarashev, Claudio Borio, and Kostas Tsatsaronis. Attributing systemic risk to individual institutions. Working Paper 308, BIS, May 2010.
- Tom Wilde. Probing granularity. *Risk*, 14(8):103–106, August 2001.

Figure 1a: Asymptotic VaR as function of default probabilities





Figure 1b: GA as function of default probabilities





Figure 1c: Relative impact as function of default probabilities





Figure 2: Asymptotic VaR and GA as functions of ELGD





Figure 3: Asymptotic VaR and GA as functions of asset correlation





Figure 4: Conditional transition probabilities



Upper line plots the conditional transition probability $\bar{\pi}_{AB}(\alpha_q(X))$ as a function of θ_A , and lower line plots the conditional transition probability $\bar{\pi}_{BA}(\alpha_q(X))$ as a function of θ_B . Solid line is the -45° line representing the unconditional transition probability.

Figure 5a: Asymptotic VaR as function of stickiness





Figure 5b: GA as function of stickiness





Figure 6: Asymptotic VaR and GA as functions of maturity





Figure 7: Asymptotic VaR and GA as functions of risk premium

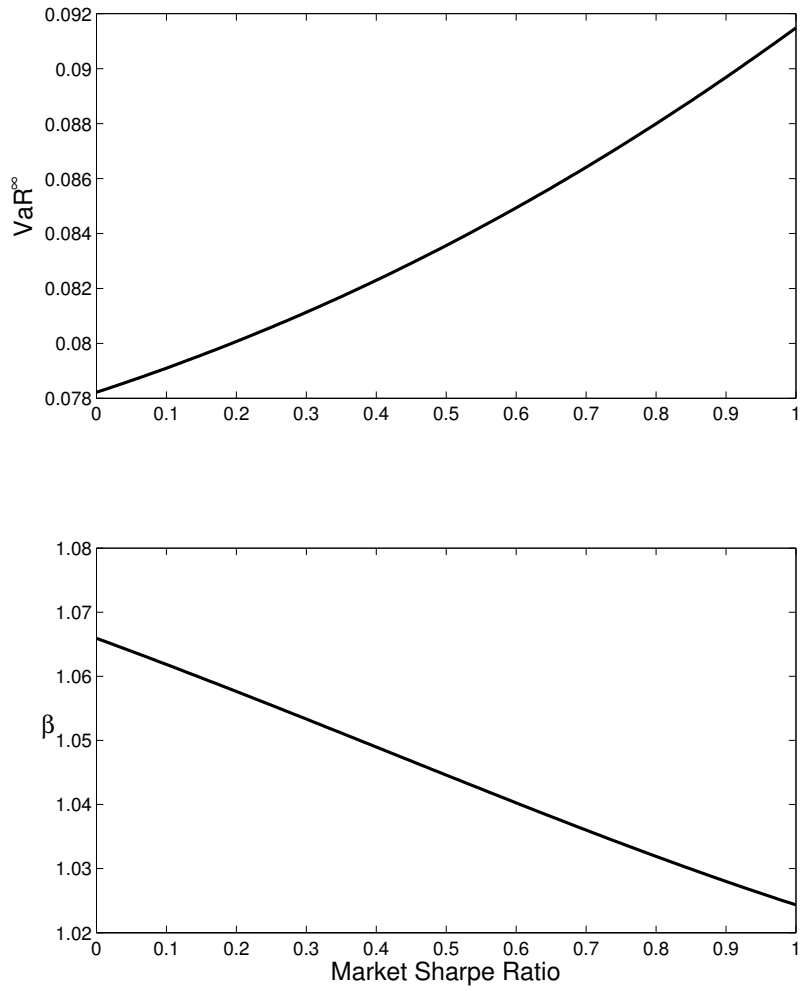





Figure 8: Density of return distribution




Probability density function for \tilde{W} in beta-trinomial model. Parameters are $n = 500$, $\lambda_0 = 1$, $\lambda_1 = 0.2$, $p_1 = 5$, $p_2 = 1$, and $\xi = 0.03$. The constant c is 0.071.

Figure 9: GA as function of state returns



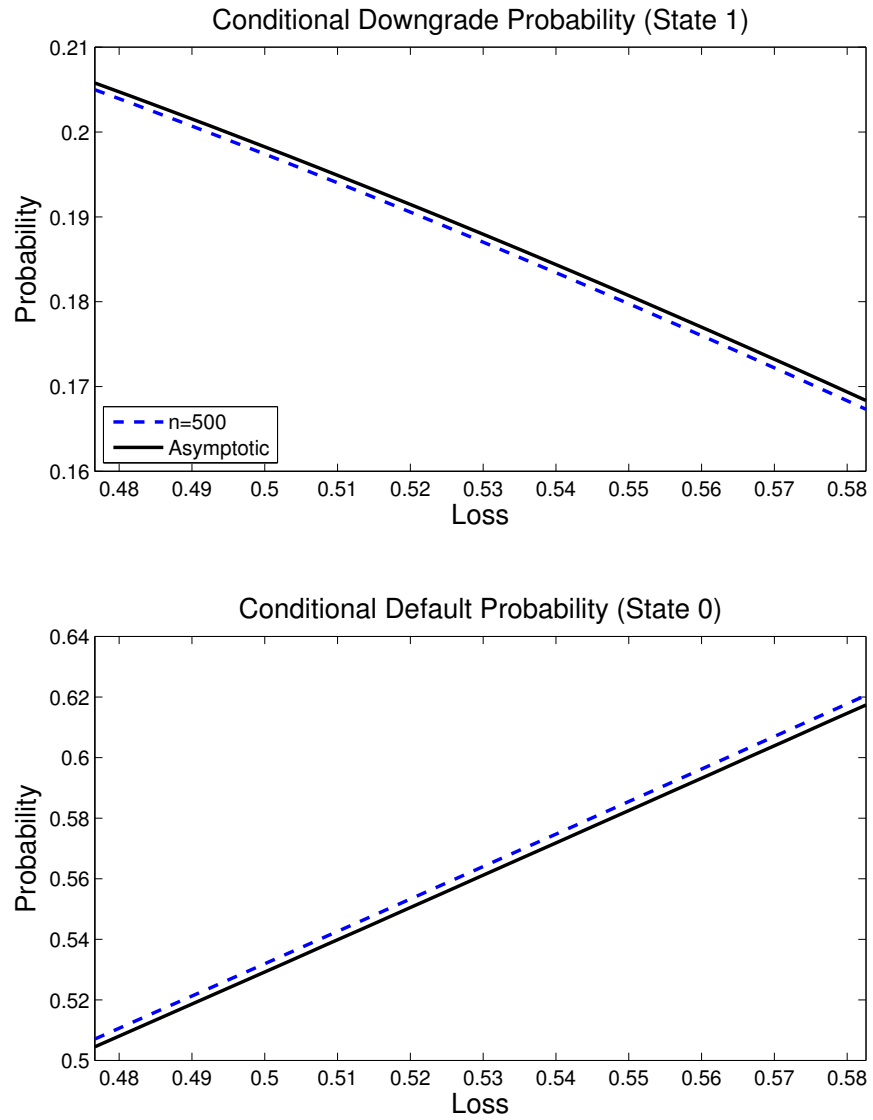
GA in beta-trinomial model. Parameters are $\lambda_1 = 0.2$ (upper panel), $\lambda_0 = 1$ (lower panel), $p_1 = 5$, $p_2 = 1$, and $\xi = 0.03$.

Figure 10: Difference in conditional state probabilities



Difference in conditional state probabilities between finite ($n = 500$) and asymptotic portfolios. Parameters are $\lambda_0 = 1$, $\lambda_1 = 0.2$, $p_1 = 5$, $p_2 = 1$, and $\xi = 0.03$.

Figure 11: Conditional state probabilities



Parameters are $\lambda_0 = 1$, $\lambda_1 = 0.2$, $p_1 = 5$, $p_2 = 1$, and $\xi = 0.03$.