Chapter 5
Approximating Economic Capital

It has long been an axiom of mine that the little things are infinitely the most important.

(Arthur Conan Doyle)

As the previous chapter makes abundantly clear, the computation of credit-risk economic capital is, irrespective of the approach taken, computationally intensive and slow. This is simply a fact of life. Clever use of parallel-processing and variance-reduction techniques can improve this situation, but they cannot entirely alleviate it. In some settings, however, speed matters. It would be extremely useful—and, at times, even essential—for certain applications, to perform large numbers of *quick* credit-risk economic-capital estimates. *Two* classic applications include:

- LOAN PRICING: computing the marginal economic-capital impact of adding a new loan, or treasury position, to the current portfolio; and
- STRESS TESTING: assessing the impact of a (typically adverse) global shock to the entire portfolio—or some important subset thereof—upon one's overall economic-capital assessment.

The needs of these applications are slightly different. Loan pricing examines a single, potential addition to the lending book; not much change is involved, but the portfolio perspective is essential.[1] Time is also very much of the essence. One simply cannot wait for 90 minutes (or so) every time one wishes to consider an alternative pricing scenario.[2] Stress-testing involves many, separate portfolio and instrument-level calculations; significant portfolio change and dislocation generally occurs.[3] One might be ready to wait a few days for one's stress-testing results, but if one desires frequent, sufficiently nuanced, and regular stress analysis, such

[1] Chapter 6 considers this application in much more detail.

[2] Even waiting 10 minutes would (eventually) make loan originators crazy and stifle their ability to consider a broad range of potential loan structures for their clients.

[3] Chapter 12 is dedicated to the discussion of stress testing.

D. Bolder, *Modelling Economic Capital*, Contributions to Finance and Accounting,
https://doi.org/10.1007/978-3-030-95096-5_5

slowness is untenable. To be performed effectively—albeit for different reasons—both applications thus require swiftness.

Sadly, speed of execution is one element woefully lacking from our large-scale industrial credit-risk economic capital model. What is to be done? Since it is simply unworkable to use the simulation engine to inform these estimates, an alternative is required. The proposed solution involves the construction of a *fast*, first-order, instrument-level *approximation* of one's economic-capital allocation. Approximation suggests the presence of some error, but in both applications an immediate, but slightly noisy, estimate is vastly superior to a more accurate, but glacially slow value.

This might seem a bit confusing. The economic-capital model is, itself, a simulation-based approximation. We would, in essence, be performing an approximation of an approximation. While there is some truth to this criticism, there is no other reasonable alternative. Failure to find a fast, semi-analytic, approximation would undermine our ability to sensibly employ our credit-risk economic capital model in a few important, and highly informative, applications.

Such an approach to this underling problem is *not* without precedent. Bolder and Rubin [7], for example, investigate a range of flexible high-dimensional approximating functions—in a rather different analytic context—with substantial success. Ribarits et al. [25] present—based on Pykhtin [23]—precisely such an economic-capital approximation for loan pricing. This work is a generalization of the so-called granularity adjustment—initially proposed by Gordy [13]—for Pillar II concentration-risk calculations.[4] This regulatory capital motivated methodology, while powerful, is *not* directly employed in our approach for one simple reason: we wish to tailor our approximation to our specific modelling choice and parametrization. A second reason is that we will require separate, although certainly related, approximations for both default and credit-migration risk. The consequence is that, although motivated by the general literature, much of this chapter is specific to the NIB setting. It will hopefully, however, stimulate ideas and provide a conceptual framework for others facing similar challenges.

5.1 Framing the Problem

Before jumping into specific approximation techniques, let's try to frame the problem. To summarize, we require a *fast* assessment of the marginal economic-capital consumption—at a given point in time—operating at the individual loan and obligor level. We may ultimately apply it at the portfolio level, but it needs to operate at the individual exposure level.[5] The plan is to construct a closed-form—or, at

[4] This area, an important part of the regulatory world, is covered extensively in Chap. 11.

[5] A portfolio analysis would, consequently, involve summing the implications across a range of individual positions.

least, semi-analytical—mathematical approximation of the economic-capital model allocations. There are, fortunately, many possible techniques in the mathematical and statistical literature that might be useful in this regard. Broadly speaking, there are *two* main approaches to such an approximation:

1. a structural description; or
2. a reduced-form—or, empirically motivated—approach.

The proposed economic-capital approximation methodology in this chapter is based, more or less, on the former method. That said, this need not be so. The sole selection criterion—ignoring, for the moment, speed—is predication accuracy. This suggests that we might legitimately consider reduced-form techniques as well.

This brings us to the general approximation problem. Let us begin by defining $y \in \mathbb{R}^{I\times 1}$ as the vector-quantity that we are trying to approximate; this is basically our collection of I economic-capital estimates at a given point in time. We further denote $X \in \mathbb{R}^{I\times\kappa}$ as a set of explanatory, or instrument, variables that can act to describe the current economic-capital outcomes.[6] Good examples of instruments would include the size of the position, its default probability, assumed recovery rate, industrial sector, and geographic region. Our available inputs thus involve I economic-capital allocations along with κ explanatory variables for each observation.

The (slow) simulated-based economic-capital model, $f(X)$, is a complicated function of these explanatory variables. Conceptually, it is a mapping of the following form:

$$y \leftarrow \boxed{\text{Economic-Capital Model: } f(X) + \epsilon} \leftarrow X \tag{5.1}$$

where ϵ is observation or measurement error.[7] In words, therefore, $f(X)$ is basically the economic-capital model. Regrettably, we do not have a simple description of f; if we did, of course, we would not find ourselves in this situation! Practically, our approximation is attempting to bypass the unknown function, f, and replace it with an approximator. Conceptually, this replaces Eq. 5.1 with

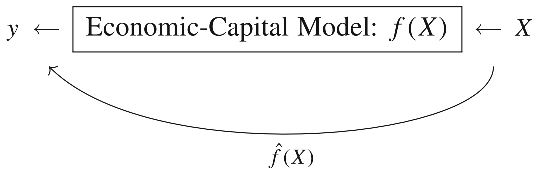

[6] In a machine-learning context, these are typically referred to as features.

[7] In our case, of course, this is more readily conceptualized as simulation error.

where we denote our approximation as $\hat{f}(X)$. The trick, and our principal task in this chapter, is to find a sensible, rapidly computed, and satisfactorily accurate choice of $\hat{f}$. There are no shortage of potential candidates. Function approximation is a heavily studied area of mathematics. Much of the statistics literature relating to parameter estimation touches on this notion and the burgeoning discipline of machine learning is centrally concerned with prediction of (unknown and complicated) functions.[8]

If y is the true model output, then our approximation can be characterized as:

$$\hat{y} = \hat{f}(X). \tag{5.2}$$

We write this value without error, which is a bit misleading. There is, quite naturally, error associated with our new approximation. We prefer, however, to describe this directly in comparison to the observed y. Indeed, a useful choice $\hat{f}$ involves a small distance between y and $\hat{y}$; this includes the approximation error. There are a variety of ways to measure this distance, but one common, and useful, metric is defined as:

$$\mathbb{E}\left[\left(y - \hat{y}\right)^2\right] = \mathbb{E}\left[\left(f(X) + \epsilon - \hat{f}(X)\right)^2\right]. \tag{5.3}$$

This is referred to as the mean-squared error. It is basically the average distance between the sum of the squared approximation errors. Squaring the errors has the desirable property of transforming both over- and underestimates into positive figures; it also helpfully generates a continuous and differentiable error function. The measurement error is also clearly inherited from the original (unknown) function f. The bottom line is that we seek an approximator, $\hat{y}$, that keeps the error definition in Eq. 5.3 at a relatively acceptable level.

With some patience, the mean-squared error can also tell us something useful about our problem. Employing a few algebraic tricks, we may re-write the previous expression as:

$$\mathbb{E}\left[\left(y - \hat{y}\right)^2\right] = \underbrace{\overbrace{\text{var}[\hat{f}(X)]}^{\text{Variance}} + \overbrace{\left(\mathbb{E}[\hat{f}(X)] - f(X)\right)^2}^{\text{Bias}^2}}_{\text{Reducible}} + \underbrace{\text{var}(\epsilon)}_{\text{Irreducible}}. \tag{5.4}$$

Our measure of distance between the observed and approximated economic-capital values can be generally categorized into *two* broad categories: reducible and

[8] There are many good references for the area of machine, or statistical, learning. A natural starting point, however, is Hastie et al. [16]. A more introductory version of this material is found in James et al. [17].

irreducible error.[9] Reducible error, which itself can be broken down into variance and bias—can be managed with intelligent model selection. Irreducible error stems from observation error; in our case, this relates to simulation noise.[10] This is something we need to either live with or, when feasible, actively seek to reduce. Bias describes fundamental differences between our approximation and true model. We might, for example, use a linear model to approximate something that is inherently non-linear. Such a choice would introduce bias. Variance comes from the robustness of our estimates across datasets. It seeks to understand how well, if we reshuffle the cards and generate a new sample, our approximator might perform. A low-variance estimator would be rather robust to changes in one's observed data-set.

Statisticians and data scientists speak frequently of a variance-bias trade-off. This implies that simultaneous reduction of both aspects is difficult (or even impossible). Statisticians, with their focus on inference, tend to accept relatively high bias for low variance. Machine-learning algorithms, due to their flexibility, tend to have lower bias, but the potential for higher bias.[11] The proposed approximation model—leaning towards the classical statistical school of thought—employs a multivariate linear regression with structurally motivated response variables. In principle, therefore, this choice involves a reasonable amount of bias. Our credit-risk economic-capital model is not, as we've seen, particularly linear in its construction. The implicit logic behind this choice, of course, is that it provides a stable, low-variance estimator.

Colour and Commentary 53 (APPROXIMATING ECONOMIC CAPITAL)*: There is a certain irony in expending—in the previous chapters—such a dramatic amount of mental and computational resources for the calculation of economic capital, only to find it immediately necessary to approximate it. This is the dark side associated with the centrality of credit-risk economic capital. Determining the marginal economic-capital consumption associated with adding (or changing) one, or many, loans has a broad range of useful appli-*

(continued)

[9] To arrive at the decomposition in Eq. 5.4, one needs to rewrite the right-hand-side of Eq. 5.3 as,

$$\mathbb{E}\left[\left(f(X) + \epsilon - \hat{f}(X) \underbrace{+\mathbb{E}[\hat{f}(X)] - \mathbb{E}[\hat{f}(X)]}_{=0}\right)^2\right], \tag{5.5}$$

and recall that, through independence and zero expectation, the product of ϵ with any term—save itself—vanishes. The rest is expansion, simplification, and tedium. See again Hastie et al. [16] or James et al. [17] for more detail, and useful colour, on this computation.

[10] In our case, this can be controlled—albeit not without cost—by increasing the total number of simulations. This is an important practical application of the convergence discussion in Chap. 4.

[11] See Breiman [8] and Bolder [6] for more detailed background on the relative differences between statistical and machine-learning approaches to this general problem.

Colour and Commentary 53 (continued)
cations. Examples include risk-adjusted return computations, loan pricing, stress testing, sensitivity analysis, and even loan impairments. The slowness of the base simulation-based computation makes—using the simulation engine—such rapid and flexible marginal computations practically infeasible. Their impossibility does not, of course, negate their usefulness. As such, the most logical, and pragmatic, course of action is to construct a fast, accurate, semi-analytic approximation to permit such analysis. Moreover, such effort is not wasted. Not only does it facilitate a number of productive computations, but it also provides welcome, incremental insight into our modelling framework.

5.2 Approximating Default Economic Capital

We've established that we need an approximation technique that, given the existing portfolio, can operate at the instrument level. As established in previous chapters, there are *two* distinct flavours of credit-risk economic capital: default and migration. While we need both, different strategies are possible. We could group the two together and construct a model to approximate the *combined* default and migration economic capital. Alternatively, we could build separate approximators for each element. Neither approach is right or wrong; it depends on the circumstances. Our choice is for distinct treatment of default and migration risk.[12] As a consequence, we'll develop our approximators separately beginning with the dominant source of risk: default.

5.2.1 Exploiting Existing Knowledge

As a general tenet, if you seek to perform an approximation of some object, it is wise to exploit whatever knowledge you have about it. Specific qualities and attributes about the economic-capital model can thus help inform our approximation decisions. We know that our production credit-risk model is a multivariate t-threshold model. We also know that threshold models essentially randomize the conditional probability of default—and thereby induce default correlation—through the introduction of common systemic variables.[13] Extreme realizations of these systemic variables push up the (conditional) likelihood of default and dispropor-

[12] We'll defend this choice later in the discussion, once we've built a bit more understanding of the general approximation approach.

[13] See Bolder [5, Chapter 4] for more colour on the global properties of threshold models.

tionately populate the tail of any portfolio's loss distribution. This link between systemic factors and tail outcomes seems like a sensible starting point.

This brings us directly to the latent creditworthiness state variable introduced in Chap. 2. To repeat it once again for convenience, it has the following form:

$$\Delta X_i = \sqrt{\frac{\nu}{W}}\left(\alpha_i \mathrm{B}_i \Delta z + \sqrt{1-\alpha_i^2}\Delta w_i\right), \tag{5.6}$$

for $i = 1, \ldots, I$. For our purposes, it will be useful to simplify somewhat the notation and structure of the model. These simplifications will enable the construction of a stylized approximator. In particular, the product of systemic factors and their loadings is, by construction, standard normally distributed; that is, $\mathrm{B}_i \Delta z \sim \mathcal{N}(0,1)$. To reduce the dimensionality, we simply replace this quantity with a single standard normal random variable, z. Moreover, to underscore the distance to the true model and make it clear we are operating in the realm of approximation, let us set $\Delta w_i = \upsilon_i$ and $\Delta X_i = y_i$. This yields

$$y_i = \sqrt{\frac{\nu}{W}}\left(\alpha_i z + \sqrt{1-\alpha^2}\upsilon_i\right), \tag{5.7}$$

which amounts to a univariate approximation of our multivariate t-threshold model. Some information is clearly lost with these actions, most particularly relating to regional and sectoral aspects of each credit obligor. This will need to be addressed in the full approximation.

A few additional components bear repeating from Chap. 2. The default event (i.e., $\mathcal{D}_i$) occurs if y_i falls below a pre-defined threshold, $K_i = F_{\mathcal{T}_\nu}^{-1}(p_i)$ or,

$$\mathbb{I}_{\mathcal{D}_i} = \mathbb{I}_{\left\{y_i \leq F_{\mathcal{T}_\nu}^{-1}(p_i)\right\}}. \tag{5.8}$$

which allows to get a step closer to our object of interest. The default loss for the ith obligor, $L_i^{(d)}$, can be written as

$$\begin{aligned} L_i^{(d)} &= \mathbb{I}_{\mathcal{D}_i} \overbrace{(1-\mathcal{R}_i)}^{\gamma_i} c_i, \\ &= \underbrace{\mathbb{I}_{\left\{y_i \leq F_{\mathcal{T}_\nu}^{-1}(p_i)\right\}}}_{\text{Equation 5.8}} \gamma_i c_i, \end{aligned} \tag{5.9}$$

where $\mathcal{R}_i$, γ_i, and c_i denote (as usual) the recovery rate, loss-given-default, and exposure-at-default of the ith obligor, respectively. The d superscript in Eq. 5.9 explicitly denotes default to keep this estimator distinct from the forthcoming migration case; despite the additional notational clutter, it will prove useful throughout the following development. The magnitude of the default loss thus ultimately depends upon the severity of the common systemic-state variable (i.e., z and W) outcomes. Equations 5.6 to 5.9 thus comprise a brief, somewhat stylized representation of our current production model.

The kernel of previous knowledge that we wish to exploit relates to the interaction between systemic variable outcomes and the conditional probability of default. This relationship is captured via the so-called conditional default loss. Conditionality, in this context, involves assuming that our common random variates, z and W, are provided in advance; we will refer to these given outcomes as z^* and w^*. Under this approach to the problem, we can derive a fairly manageable expression for this quantity as

$$\begin{aligned}
\underbrace{\mathbb{E}\left(L_i^{(d)}\middle| z=z^*, W=w^*\right)}_{\text{Conditional default loss}} &= \mathbb{E}\left(\underbrace{\mathbb{I}_{\left\{y_i \leq F_{\mathcal{T}_\nu}^{-1}(p_i)\right\}}\gamma_i c_i}_{\text{Equation 5.9}}\middle| z=z^*, W=w^*\right), \qquad (5.10)\\
&= \mathbb{P}\left(\underbrace{\sqrt{\frac{\nu}{W}}\left(\alpha_i z + \sqrt{1-\alpha_i^2}\upsilon_i\right)}_{y_i} \leq F_{\mathcal{T}_\nu}^{-1}(p_i)\middle|\right.\\
&\qquad \left. z=z^*, W=w^*\right)\mathbb{E}(\gamma_i)c_i,\\
&= \mathbb{P}\left(\upsilon_i \leq \frac{\sqrt{\frac{W}{\nu}}F_{\mathcal{T}_\nu}^{-1}(p_i)-\alpha_i z}{\sqrt{1-\alpha_i^2}}\middle| z=z^*, W=w^*\right)\\
&\qquad \mathbb{E}(\gamma_i)c_i,\\
&= \underbrace{\Phi\left(\frac{\sqrt{\frac{w^*}{\nu}}F_{\mathcal{T}_\nu}^{-1}(p_i)-\alpha_i z^*}{\sqrt{1-\alpha_i^2}}\right)}_{p_i(z^*,w^*)}\mathbb{E}(\gamma_i)c_i,
\end{aligned}$$

given that υ_i follows, by construction, a standard normal distribution. This expression—or, at least, a variation of it—forms the structural basis for much of current regulatory guidance.[14] Equation 5.10 also plays a central role in our *default* economic-capital approximation.

To move further, we first need to be a bit more precise on the actual values associated with our conditioning variables. Equation 5.10 applies to any loss associated with arbitrary systemic-variable realizations. If we may select any choice of z^* and w^*—and we're interested in worst-case default losses—why not pick

[14] See BIS [2, 3, 4] for rather more detail. We will turn to regulatory questions in Chap. 11.

really bad ones? Indeed, why not use the same level of confidence (i.e., 0.9997) embedded in our economic-capital metrics, which we will generically denote as α^*. If we seek an extreme downside observation for the common systemic risk factor, a sensible choice is

$$z^* = \Phi^{-1}(1 - \alpha_z^*), \tag{5.11}$$

where α_z^* denotes the confidence level associated with z. For the common mixing variable, a natural choice would be

$$w^* = \text{Inv} - \chi^2(1 - \alpha_w^*, \nu), \tag{5.12}$$

where again α_w^* is separately defined. In principle, we would prefer that $\alpha_z^* = \alpha_w^*$. As is often the case, the situation is a bit more nuanced. While the choice of systemic-state variable outcome is quite reasonable, unfortunately, the value in Eq. 5.12 does not perform particularly well for estimating worst-case default loss. In some cases, it works satisfactorily, whereas in other settings it is far too extreme and leads to dramatic overestimates of the conditional default loss.

This brings us to an interesting insight into the t-threshold model. The larger the confidence interval used to evaluate Eq. 5.12, the smaller the value of the common mixing variable, w^*. The smaller the w^* outcome, the smaller the ratio $\sqrt{\frac{w^*}{\nu}}$, which in turn modifies the default threshold. Reducing the size of the threshold, essentially amounts to a reduction in the creditworthiness of the obligor; making default more probable.[15] The conditional default probability associated with extreme outcomes of z^* and w^* is thus hit from both ends: on one side there is a nasty systemic state variable outcome, while on the other the threshold goal-line has been moved closer. Setting them both to extreme values can, in many cases, lead to really dramatic default losses.

The best way, perhaps, to understand the nature of this calculation is to consider a practical example. Table 5.1 summarizes all of the key inputs for an illustrative, but arbitrary, credit obligor. Three levels of α_w^* confidence interval are provided for the w^* outcomes, but the z^* outcomes are held constant. Key components of the computation are displayed illustrating the impact of w^* and the sensitivity of overall results to this choice. Innocently moving α_w^* from 0.6667 to 0.9997—all else equal—leads to fairly dramatic increases in the default loss estimate.

Figure 5.1 explores this question in more detail by widening our gaze. It displays, for three different choices of α_w^*, the *normalized* estimated and observed default loss values across the entire portfolio.[16] The clear conclusion is that the estimated value increases steadily—and ultimately rather significantly exceeds the observed

[15] Low levels of confidence have the opposite effect; they increase the ratio $\sqrt{\frac{w^*}{\nu}}$, pushing out the threshold, and effectively increasing the credit quality for all obligors.

[16] Normalized, in this context, means that the the average *observed* values are subtracted and the result is divided by the volatility of the *observed* default losses; this essentially means that the axes

Table 5.1 *A simple example*: The underlying table summarizes, for an arbitrary and illustrative exposure, the various quantities and calculations involved in the construction of a conditional probability of default with the preceding conditioning variables, z^* and w^*.

Quantity	Definition	Severity of w^*		
		$\alpha_w^* = 0.6667$	$\alpha_w^* = 0.9000$	$\alpha_w^* = 0.9997$
Degrees of freedom	ν	70		
Confidence level	$\alpha^* \equiv \alpha_z^*$	0.9997		
Exposure	c_i	65,000,000		
Credit state	S_i	11		
Loss-given-default	$\mathbb{E}(\gamma_i)$	0.52		
Default probability	p_i	0.55%		
Systemic weight	α_i	0.19		
Systemic shock	$z^* = \Phi^{-1}(1-\alpha_z^*)$	−3.43		
Mixing-variable shock	$w^* = \text{Inv} - \chi^2(1-\alpha_w^*, \nu)$	64.39	55.33	36.39
Threshold value	$F_{\mathcal{T}_\nu}^{-1}(p_i)$	−2.61		
Mixing ratio	$\sqrt{\frac{w^*}{\nu}}$	0.96	0.89	0.72
Key ratio	$\frac{\sqrt{\frac{w^*}{\nu}} F_{\mathcal{T}_\nu}^{-1}(p_i) - \alpha_i z^*}{\sqrt{1-\alpha_i^2}}$	−1.11	−0.42	0.00
Conditional default probability	$p_i(z^*, w^*)$	13.37%	33.85%	50.00%
Conditional default loss	$\mathbb{E}\left(L_i^{(d)} \middle\vert z = z^*, W = w^*\right)$	4,519,809	11,442,358	16,900,000

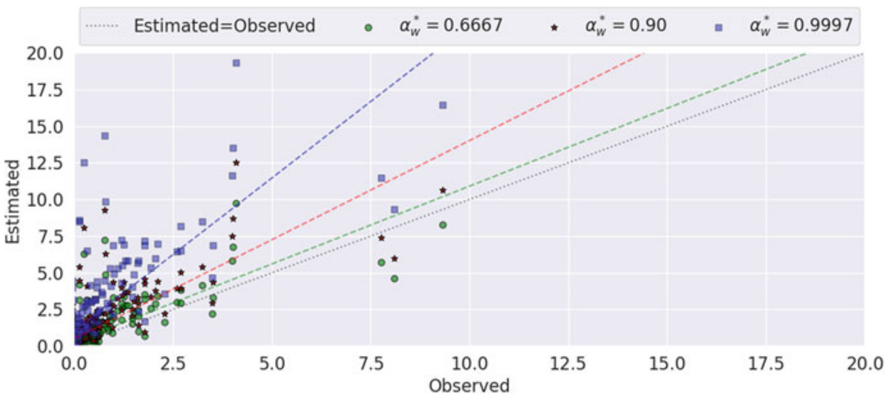

Fig. 5.1 *Choosing* α_w^*: The preceding graphic displays, for three different choices of α_w^* used to identity w^*, the *normalized* estimated and observed default loss values. The larger the choice of α_w^*, the more extreme the default-loss estimates.

outcomes—for larger values of α_w^*. The degree of sensitivity of the conditional default probability to the choice of confidence level for the mixing variable is rather surprising. It seems that this variable's principle role is not to generate extreme loss outcomes, but rather simply to induce a joint t distribution; and, as a result, tail dependence.[17] For the remainder of this analysis, therefore, we will use the more neutral, but slightly conservative, value of $\alpha_w^* = 0.6667$. Different values are naturally possible, and this is by no means an optimized choice, but it performs quite reasonably.

Colour and Commentary 54 (THE t-THRESHOLD MIXING VARIABLE)*: In the Gaussian threshold model, the conditional probability of default depends solely on the provided realization of the systemic risk factor (or factors). Intuition about the randomization of the (conditional) default probability—and the inducement of default correlation—is straightforward: extreme systemic factor realizations lead to correspondingly high probabilities (and magnitudes) of default loss. Moving to the t-threshold setting complicates things. A second conditioning variable—the χ^2-distributed quantity—must also be revealed. One's first reaction is to simultaneously draw an extreme mixing-variable outcome. Our analysis demonstrates that this yields unrealistically large default loss probabilities.*[a] *This appears to suggest that the mixing variable's role is* not *to generate extreme outcomes, but rather to alter the joint distribution and induce tail dependence. For this reason, when using the conditional default probability of the t-threshold model for our approximation analysis, we assume a relatively neutral draw from the χ^2-distributed mixing variable.*

[a] Effectively, extreme mixing variable outcomes pull in the default threshold for all obligors. Combined with a simultaneously bad draw from the systemic state variable, obligors are hit from both sides and the impact is excessive.

5.2.2 Borrowing from Regulatory Guidance

Equation 5.10 actually turns out to be a critical ingredient of our initial closed-form—although, admittedly simple—approximation of the marginal economic-capital consumption associated with a given exposure. A few additional steps are

in Fig. 5.1 denote standard deviations from the observed mean. We'll use this idea throughout the chapter.

[17] Moreover, given their independence, the joint extreme coincidence of systemic and mixing variable outcomes should be exceedingly rare. Forcing them both to the extremes clearly scales up the default-loss outcomes in an unrealistic manner.

nonetheless required. Let us define $\mathcal{A}_i^{(d)}(\alpha^*)$ as the observed *default* economic-capital allocation associated with the ith obligor computed at confidence level, α^*. Recall that economic-capital is, conceptually, written as

$$\mathcal{A}_i^{(d)}(\alpha^*) = \underbrace{\text{Worst-Case Loss}_i(\alpha^*) - \text{Expected Loss}_i}_{\text{Unexpected Loss}_i(\alpha^*)}. \tag{5.13}$$

For the (Pillar I) computation of minimum regulatory capital requirements—see Chap. 11 for much more on these ideas—the internal ratings-based approach offers the following exposure-level formula for the direct computation of this quantity:

$$\mathcal{A}_i^{(d)}(\alpha^*) \approx \underbrace{\overbrace{\mathbb{E}\left(L_i^{(d)} \middle| z = z^*, W = w^*\right)}^{\text{Equation 5.10}} - \mathbb{E}\left(L_i^{(d)}\right)}_{\text{Unexpected Loss}_i(\alpha_z^*)}, \tag{5.14}$$

which basically amounts to the difference between conditional and unconditional default loss. This quantity is, in the classical sense, a VaR-related measure of risk.[18]

Exploitation of Eq. 5.14 leads us to the kernel of the approximation method. The mathematical structure of the approximation is

$$\begin{aligned}
\underbrace{\mathcal{A}_i^{(d)}(\alpha_z^*)}_{\equiv \mathcal{A}_i^{(d)}(\alpha^*)} &\approx \mathbb{E}\left(L_i^{(d)} \middle| z = z^*, W = w^*\right) - \mathbb{E}\left(L_i^{(d)}\right), \qquad (5.15)\\
&\approx \underbrace{\mathbb{E}\left(L_i^{(d)} \middle| \alpha_z^*\right)}_{\text{Equation 5.10}} - p_i \cdot \mathbb{E}(\gamma_i) \cdot c_i,\\
&\approx \underbrace{\Phi\left(\frac{\sqrt{\frac{w^*}{\nu}} F_{\mathcal{T}_\nu}^{-1}(p_i) - \alpha_i z^*}{\sqrt{1-\alpha_i^2}}\right)}_{p_i(\alpha_z^*)} \mathbb{E}(\gamma_i) c_i - p_i \cdot \mathbb{E}(\gamma_i) \cdot c_i,\\
&\approx \Big(p_i(\alpha_z^*) - p_i\Big) \cdot \mathbb{E}(\gamma_i) \cdot c_i,
\end{aligned}$$

for $i = 1, \ldots, I$ risk-owner risk contributions. To be fair, the previously described unexpected loss expression does not precisely coincide with the regulatory definition in Eq. 5.14. The t-threshold model leads to a few conceptual changes. The most

[18] Our interest is in an expected-shortfall metric, but we will address this issue in a moment.

important is the presence of our common mixing variable, w^*. For this computation, as already indicated, we will set w^* at a fixed, reasonably neutral quantity—to induce tail dependence—and treat the α_z^* and z^* as the true variables of interest.[19] Collecting our thoughts, Eq. 5.15 is an analytical representation of a regulatory capital, VaR-based economic-capital allocation. The intuition is that a worst-case outcome for the default loss is inferred from a catastrophically bad outcome of the systemic variable, z^*. Subtracted from this is the expected loss computed by averaging over all possible outcomes of the systemic state variable.[20]

There is a definitional issue associated with the approximation in Eq. 5.15: it is based on the idea of a VaR unexpected-loss estimator. We use, as highlighted in previous chapters, the expected-shortfall metric. Using the definition of expected-shortfall, and applying it to Eq. 5.15, this can be rectified as

$$\underbrace{\frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^1 \mathcal{A}_i^{(d)}(x)dx}_{\mathcal{E}_i^{(d)}(\alpha_z^*)} \approx \frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^1 \underbrace{\left(p_i(x)-p_i\right)\cdot\mathbb{E}(\gamma_i)\cdot c_i\,dx}_{\text{Equation 5.15}}, \tag{5.16}$$

$$\begin{aligned}
\mathcal{E}_i^{(d)}(\alpha_z^*) &\approx \mathbb{E}(\gamma_i)\cdot c_i\cdot \underbrace{\frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^1 p_i(x)dx}_{\tilde{p}_i(\alpha_z^*)} - \frac{p_i\mathbb{E}(\gamma_i)\cdot c_i}{1-\alpha_z^*}\int_{\alpha_z^*}^1 dx,\\
&\approx \mathbb{E}(\gamma_i)\cdot c_i\cdot\tilde{p}_i(\alpha_z^*) - \frac{p_i\mathbb{E}(\gamma_i)\cdot c_i}{1-\alpha_z^*}\Big[x\Big]_{\alpha_z^*}^1,\\
&\approx \mathbb{E}(\gamma_i)\cdot c_i\cdot\tilde{p}_i(\alpha_z^*) - \frac{p_i\mathbb{E}(\gamma_i)\cdot c_i}{\cancel{1-\alpha_z^*}}\cancel{(1-\alpha_z^*)},\\
&\approx \left(\tilde{p}_i(\alpha_z^*)-p_i\right)\cdot\mathbb{E}(\gamma_i)\cdot c_i.
\end{aligned}$$

There is both good and bad news. The good news is that, when moving to the expected-shortfall setting, the basic form of our approximation is preserved. The bad news is that the integral $\tilde{p}_i(\alpha_z^*)$ is not, to the best of our knowledge, available in closed form. Computation of this quantity requires solving a one-dimensional numerical-integration problem. To solve 500+ individual problems, this takes a bit less than 2 seconds. For large-scale approximations—such as stress-testing calculations—this has the potential to slow things down somewhat, but no real damage is involved. For an individual loan, thankfully, this numerical element

[19] There is, therefore, a small mathematical sleight of hand with the confidence levels.

[20] For this reason, each $\mathbb{E}\left(L_i^{(d)}\right)$ is independent of z.

Table 5.2 *Simple economic-capital estimates*: This table illustrates some key results of using Eq. 5.10 and 5.16 to estimate the default-related economic-capital consumption associated with the simple example introduced in Table 5.1.

Quantity	Definition	Value
Model default economic capital	$\mathcal{E}_i^{(d)}(\alpha^*)$	4,314,601
Model default economic-capital ratio	$\frac{1}{c_i}\mathcal{E}_i^{(d)}(\alpha^*)$	6.6%
VaR regulatory-capital approximation	$\mathcal{A}_i^{(d)}(\alpha_z^*) = \left(p_i(\alpha_z^*) - p_i\right) \cdot \mathbb{E}(\gamma_i) \cdot c_i$	4,333,909
VaR regulatory-capital ratio	$\frac{1}{c_i}\mathcal{A}_i^{(d)}(\alpha_z^*)$	6.7%
Expected-shortfall regulatory-capital approximation	$\mathcal{E}_i^{(d)}(\alpha_z^*) = \left(\tilde{p}_i(\alpha_z^*) - p_i\right) \cdot \mathbb{E}(\gamma_i) \cdot c_i$	5,358,698
Expected-shortfall regulatory-capital ratio	$\frac{1}{c_i}\mathcal{E}_i^{(d)}(\alpha_z^*)$	8.2%

requires only a fraction of a second.[21] As a consequence, this semi-analytic twist to our approximation does not delay the associated applications in any appreciable way.

Table 5.2 takes the simple example introduced in Table 5.1 and—with the help of Eqs. 5.10 and 5.16—provides some insight into how well the approximations actually work. The true model-based default economic-capital value is about EUR 4 million amounting to roughly $6\frac{1}{2}\%$ of the position's overall exposure. This figure seems sensible for a position falling, albeit slightly, below investment grade. The VaR-based approximation, from Eq. 5.15, also generates an estimate a bit north of EUR 4 million or around 6.7%. The VaR estimator, in this case, thus generates a marginal over-estimate. Incorporation of the expected-shortfall element—using, in this case, the numerical-integration estimator introduced in Eq. 5.16—yields a rather higher estimate of about EUR 5.4 million or in excess of 8%. This represents a rather significant overstatement of the model-based economic capital.

Our simple example illustrates that, as a generic approximation, the presented results appear to be in the vicinity. That said, the VaR estimator seems to outperform the expected-shortfall calculation. To make a more general statement, of course, it is necessary to look beyond a single instance. We need to systemically examine the relationship between actual economic-capital estimates and our approximating expressions. Figure 5.2 consequently compares our regulatory motivated economic-capital approximations—computed with Eq. 5.16—to the true simulation-model outcomes. Since the actual levels of economic capital are unimportant, the results

[21] The implementation makes use of `scipy`'s `quadrature` function. See Ralston and Rabinowitz [24] for more on Gaussian quadrature and numerical integration.

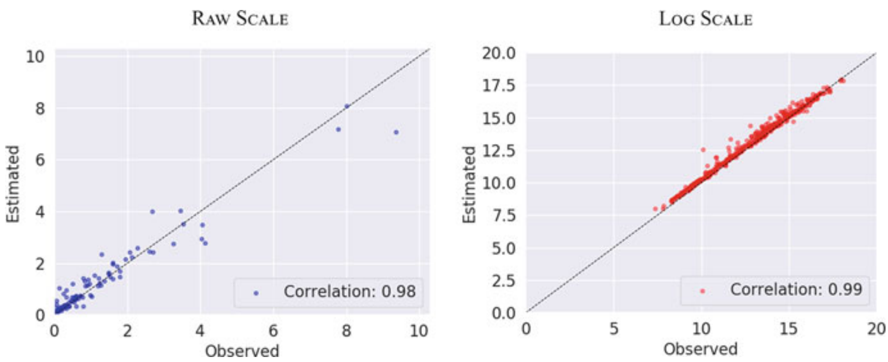

Fig. 5.2 *Degree of agreement*: The preceding graphics, in raw and logarithmic terms, illustrate the degree of agreement between our (normalized) approximation in Eq. 5.16 and the (normalized) observed simulation-based economic-capital contributions. The overall level of linear correlation is fairly respectable.

have again been normalized. The left-hand graphic outlines the comparative results on a normalized basis similar to Fig. 5.1. Visually, the agreement appears sensible, although there are numerous instances where individual estimates deviate from the observed values by a sizable margin. Interestingly, at the overall portfolio level, despite the approximation noise, there is a really quite respectable cross-correlation coefficient of 0.98. This suggests that, although Eq. 5.16 may not always capture the level of default risk, it does catch the basic pattern.

A related issue is the broad range of economic-capital values. For some obligors, the allocation approaches zero, whereas for others it can be one (or more) orders of magnitude larger. This stems from the exponential form of default probabilities as well as concentration in the underlying portfolio. This fact, combined with the the patchy performance of the approximation in Eq. 5.16, can lead to scaling problems. As a consequence, the right-hand graphic illustrates the approximation fit when applying natural logarithms to both (normalized) observed and approximated default-risk values. This simple transformation smooths out the size dimension and increases the correlation coefficient to 0.99. Although far from perfect, we may nonetheless cautiously conclude that there is some merit to this basic approach.

Colour and Commentary 55 (The Approximation Kernel)*: The central component of the default economic capital approximator is pinched directly from the logic employed in the Basel Internal-Rating Based (or IRB) approach. More specifically, worst-case default losses are approximated via the conditional probability of default associated with a rather highly adverse*

(continued)

Colour and Commentary 55 (continued)
realization of the systemic risk factor.[a] *While sensible and intuitive, this is* not *an assumption-free choice. As covered in detail in Chap. 11, Gordy [13] shows us that the IRB model is only consistent with the so-called single-factor asymptotic risk-factor model. The principal implication is that this modelling choice focuses* entirely *on systemic risk; it ignores the idiosyncratic effects associated with portfolio concentrations. The strength of this assumption—as well as the concentrations in our real-life portfolio—implies that this base structure can only take us so far. It will be necessary, as we move forward, to adjust our approximation to capture the ignored idiosyncratic elements.*

[a] We also draw a more neutral-valued χ^2 mixing variable to accommodate the t-threshold structure of our production model.

5.2.3 A First Default Approximation Model

Raw use of Eq. 5.16 is likely to be sub-optimal. Its high correlation with observed values, however, suggests that it makes for a useful starting point. What we need is the ability to allow for somewhat more approximation flexibility in capturing individual differences. In other words, we seek an approximation model. This is the underlying motivation behind transforming Eq. 5.16 into an ordinary least-squares (OLS), or regression, estimator. The presence of regression coefficients allows us to exploit the promising correlation structure and also increases the overall fit to the observed data.

The basic linear structure follows from Eq. 5.16 and we use the logarithmic operator to induce an additive form as

$$\mathcal{E}_i^{(d)}(\alpha_z^*) \approx \underbrace{\left(\tilde{p}_i(\alpha_z^*) - p_i\right)\mathbb{E}(\gamma_i)c_i}_{\text{Equation 5.16}}, \tag{5.17}$$

$$\mathcal{E}_i^{(d)}(\alpha_z^*) \approx \tilde{p}_i(\alpha^*)\mathbb{E}(\gamma_i)c_i - \underbrace{p_i\mathbb{E}(\gamma_i)c_i}_{\mathbb{E}\left(L_i^{(d)}\right)},$$

$$\mathcal{E}_i^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_i^{(d)}\right) \approx \tilde{p}_i(\alpha_z^*)\mathbb{E}(\gamma_i)c_i,$$

$$\ln\left(\mathcal{E}_i^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_i^{(d)}\right)\right) \approx \ln\left(\tilde{p}_i(\alpha_z^*)\mathbb{E}(\gamma_i)c_i\right),$$

$$\ln\left(\mathcal{E}_i^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_i^{(d)}\right)\right) \approx \ln\left(\tilde{p}_i(\alpha_z^*)\right) + \ln\left(\mathbb{E}(\gamma_i)\right) + \ln(c_i).$$

Equation 5.17 is an approximation. If we re-imagine this linear relationship with an intercept and error, we have

$$\underbrace{\ln\left(\mathcal{E}_i^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_i^{(d)}\right)\right)}_{y_i^{(d)}} = \xi_0 + \xi_1 \underbrace{\ln\left(\tilde{p}_i(\alpha_z^*)\right)}_{X_{i,1}^{(d)}} + \xi_2 \underbrace{\ln\left(\mathbb{E}(\gamma_i)\right)}_{X_{i,2}^{(d)}} + \xi_3 \underbrace{\ln(c_i)}_{X_{i,3}^{(d)}} + \epsilon_i^{(d)}, \tag{5.18}$$

for $i = 1, \ldots, I$ risk-owner risk contributions.[22] The left-hand side of Eq. 5.18 is a transformation of the observed default economic-capital estimates, while the elements of our approximation from Eq. 5.16 have become additive explanatory variables.

Equation 5.18 directly translates, of course, into the classical ordinary least squares setting and reduces to the familiar model,

$$y^{(d)} = X^{(d)}\Xi + \epsilon^{(d)}, \tag{5.19}$$

where

$$y^{(d)} = \begin{bmatrix} \ln\left(\mathcal{E}_1^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_1^{(d)}\right)\right) \\ \vdots \\ \ln\left(\mathcal{E}_I^{(d)}(\alpha_z^*) + \mathbb{E}\left(L_I^{(d)}\right)\right) \end{bmatrix}, \tag{5.20}$$

and

$$X^{(d)} = \begin{bmatrix} 1 & \ln\left(\tilde{p}_1(\alpha_z^*)\right) & \ln\left(\mathbb{E}(\gamma_1)\right) & \ln(c_1) \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \ln\left(\tilde{p}_I(\alpha_z^*)\right) & \ln\left(\mathbb{E}(\gamma_I)\right) & \ln(c_I) \end{bmatrix}, \tag{5.21}$$

and $\epsilon^{(d)} \sim \mathcal{N}(0, \sigma^2 I)$ for $\sigma \in \mathbb{R}_+$. The OLS estimator for $\Xi = \begin{bmatrix} \xi_0 & \xi_1 & \xi_2 & \xi_3 \end{bmatrix}^T$ determined, as usual, through the minimization of $(y^{(d)} - X^{(d)}\Xi)^T(y^{(d)} - X^{(d)}\Xi)$

[22] It is customary to use the symbol β for regression coefficients. Since we already employ this character to denote factor loadings, we will use an alternative esoteric Greek letter to avoid confusion.

Table 5.3 *A first default-approximation model*: This table summarizes the results of the base default economic-capital approximation model. It has an intercept and *three* regulatory capital requirement motivated explanatory variables. The amount of variance explained is fairly encouraging, but the proportional errors are depressingly (even frighteningly) large.

Quantity	Description	Estimate	Standard error	t-statistic	p-value
ξ_0	Intercept	−0.91	0.13	−7.06	0.00
ξ_1	(Extreme) conditional default probability	0.92	0.02	53.18	0.00
ξ_2	Loss-given-default	1.08	0.03	35.10	0.00
ξ_3	Exposure	1.04	0.01	158.95	0.00
Root mean *squared* error		45%			
Mean absolute error		18%			
Median absolute error		10%			
R^2		0.923			

is thus simply the well-known least-squares solution:

$$\hat{\Xi} = \left(X^{(d)^T} X^{(d)} \right)^{-1} X^{(d)^T} y^{(d)}. \tag{5.22}$$

Standard errors and test-statistic formulae also follow from similarly well-known results.[23] As discussed previously, this does not imply that our model is inherently linear, but rather that we are knowingly accepting a certain degree of model bias for low variance and the ability to use statistical-inference techniques.

Table 5.3 summarizes the results of fitting the previously described approximation model to observed economic-capital allocation data associated with an arbitrary date in 2020. All *four* of the parameter estimates are strongly statistically significant. The actual parameter values are also positive for the conditional default probability, loss-given-default, and exposures, which seems plausible. The R^2 statistic of roughly 0.9 suggests that the amount of total variance explained by a simple linear model is encouraging.

The actual calculation of approximated (i.e., predicted) default-risk values actually involves a small bit of algebraic gymnastics. In particular,

$$\begin{aligned} \hat{y}^{(d)} &= X^{(d)} \hat{\Xi}, \\ e^{\hat{y}^{(d)}} &= e^{X^{(d)} \hat{\Xi}}, \\ e^{\ln\left(\widehat{\mathcal{E}^{(d)}(\alpha_z^*)} + \mathbb{E}\left(L^{(d)}\right)\right)} &= e^{X^{(d)} \hat{\Xi}}, \\ \widehat{\mathcal{E}^{(d)}(\alpha_z^*)} &= e^{X^{(d)} \hat{\Xi}} - \mathbb{E}\left(L^{(d)}\right), \end{aligned} \tag{5.23}$$

[23] For much more detailed background on the rudiments of OLS, please see Judge et al. [19].

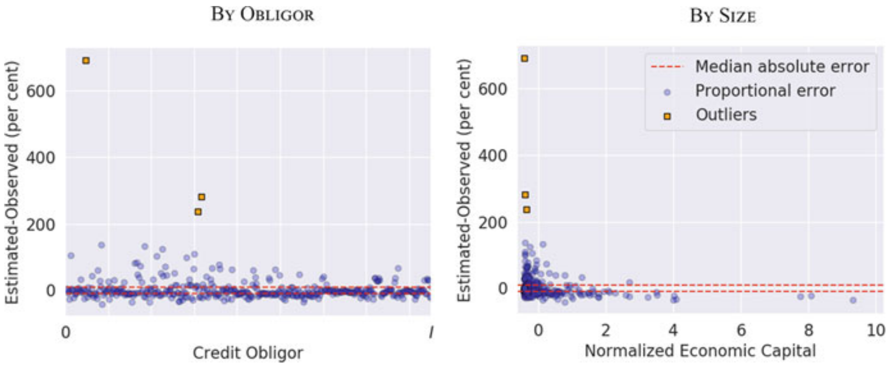

Fig. 5.3 *Initial error analysis*: The preceding graphics display the model errors—estimated less observed values—organized by both individual credit obligor and normalized default economic capital. The handful of extreme outliers are presented separately.

where $\widehat{\mathcal{E}^{(d)}(\alpha_z^*)}$ and $\mathbb{E}\left(L^{(d)}\right)$ denote the I-dimensional vectors of approximated default economic-capital values and expected losses, respectively. Equation 5.23 permits computation of root-mean-squared, mean-absolute and median-absolute error measures. The results are provided, in Table 5.3, in (proportional) percentage terms rather than logarithmic or currency space.[24] These goodness-of-fit measures—assessing the (in-sample) correspondence between observed and predicated values—unfortunately paint a less rosy picture. The mean-absolute error is roughly 20%; implying that, on average, the approximation misses its mark by about one fifth. If we turn to the median-absolute error, it improves to one tenth; this strongly suggests the presence of a few large outliers. Our suspicion is confirmed with the—notoriously sensitive to outliers—root-mean-squared error taking a dreadful value of 45%.[25]

Figure 5.3 looks into the outlier question more carefully by illustrating the collection of proportional approximation errors.[26] The left hand graphic indicates

[24] The logarithmic-transformation is helpful to manage scale discrepancies between economic-capital allocations, but it is essentially a statistical trick to improve the overall fit. Since we ultimately require currency estimates, model diagnostics in logarithmic space are *not* terribly helpful and thus excluded.

[25] Outlier sensitivity is easily seen from the definition. Given $\epsilon^{(d)} \in \mathbb{R}^{I \times 1}$, then

$$\text{RMSE} = \sqrt{\frac{{\epsilon^{(d)}}^T \epsilon^{(d)}}{I}}. \tag{5.24}$$

The squared term in the numerator places rather heavy weight on any individual and extreme error (i.e., outlier).

[26] The errors are defined as estimated less observed default economic values divided by their true observed outcome.

a general trend towards over-estimation and a handful of extreme outliers.[27] The right-hand graphic provides some insight into the outliers; they appear to all relate to small, below average economic-capital allocations. This is one of the dangers of using proportional errors; a large proportional error can, when considered in currency terms, be economically immaterial.

Despite an R^2 figure exceeding 0.9, there is work to be done; the proportional errors—even when excluding the outliers—are simply too large to be practically useful. Thankfully, there are good reasons to suspect that the model is incomplete. The approximation is founded on a one-dimensional model, which (by construction) completely ignores concentration, provides no information on random recovery, and fails to incorporate a broader default-correlation perspective. Our approximation model can thus be further improved.

Colour and Commentary 56 (A First Default Approximation Model)*: Our regulatory capital motivated first-order approximation of default economic capital is readily generalized into a linear regression model. This permits calibration to daily data and provides diagnostics on the strength of specific parts of the approximation. The initial results are mixed. On the positive side, the percentage of variance explained exceeds 90% and the individual regression variables are highly statistical significant. More negatively, proportional goodness-of-fit measures suggest an unacceptable level of inconsistency. Part of the problem is explained by a handful of outliers, but the model also has difficulty with a sizable subset of individual obligors. In many ways, this is not a surprise. The explanatory variables in this initial model capture, literally by their definition, only the systemic aspect of economic capital. This suggests identification and incorporation of idiosyncratic and concentration related explanatory variables as a sensible avenue towards improvement of our initial model.*

5.2.4 Incorporating Concentration

Our first naive default approximation model is unfortunately, in its current form, not quite up to the task. This is not just due to a deficiency of the underlying regulatory capital computation, but it is rather a structural issue. Gordy [13]—in his excellent first-principle analysis of the Basel-IRB methodology—demonstrates that this regulatory formula is equivalent to using Vasicek [28]'s single-factor asymptotic risk-factor model (ASRF). This has two immediate consequences. First, it requires

[27] Outliers, in this context, were (somewhat arbitrarily) identified as a proportional error exceeding 200%.

that one's portfolio is—in practical terms—extremely well diversified across many credit obligors. Second, as the name suggests, there is only a single source of risk. As such, there is simply no mechanism to capture regional, sectoral, or firm-size effects. To be blunt, this implies that our initial model in Eq. 5.18 is simply not equipped to capture concentration effects.

Most real-world financial institutions—and NIB is no exception in this regard—have regional and sectoral concentrations. Ignoring this dimension leads, as became clear in the previous section, to structural bias in our economic-capital approximations. The most natural route, remaining in the regulatory area, would involve the so-called granularity adjustment. This quantity, introduced by Gordy [13], is an add-on to the base regulatory capital computation to account for concentration.[28] After significant experimentation, we rejected this path. The principal issue is that it does not (easily, at least) have the requisite flexibility to handle either multiple risk factors or readily incorporate regional and sectoral information.[29]

We opt to follow an alternative, and perhaps more intuitive, strategy. The idea is to assign a concentration score to each individual set of lending exposures sharing similar regional and structural characteristics. We can then include this score as an explanatory variable in our regression model in Eq. 5.18. If done correctly, this would help to incorporate important idiosyncratic features of one's exposures into our approximation.

There is a rich literature on concentration. Concentration and Lorenz curves, Herfindahl-Hirschman indices and Gini coefficients are popular tools in this area.[30] The challenges with these methods, for our purposes, are twofold. First, they require full knowledge of the economic-capital weights, which will not always be (logically) available to us. Second, and perhaps more importantly, we seek a concentration measure that incorporates aspects of our current model with a particular focus on sectoral and regional dimensions.

The proposed score, or index, is a function of *three* pieces of information about the obligor. These include:

1. geographic region;
2. industrial sector; and
3. public-sector or corporate status.

The measure is specialized to the current credit-risk economic-capital model's factor structure. It does not stem from any current literature, but should rather be viewed as something of a heuristic concentration metric.

[28] This interesting and important area is addressed in Chap. 11. If you cannot wait to dig into this fascinating area, you are immediately referred to Lütkebohmert [20, Chapter 11], Gordy and Lütkebohmert [14, 15], Martin and Wilde [21], Emmer and Tasche [9], and Torell [27].

[29] Pykhtin [23] does offer a multivariate approach, which is explored in Chap. 11.

[30] For more background on this fascinating field of study, please see Yitzhaki and Olkin [29], Figini and Uberti [10], Bellalah et al. [1], Milanovic [22], and particularly the extremely useful Lütkebohmert [20].

We take a slow, methodical, and constructive approach to development of our concentration metric. The first step in the fabrication of our index begins by assigning, at a given point in time, total portfolio exposure (i.e., exposure-at-default) to our J systemic state variables. Since our I credit obligors may simultaneously be members of both a region and an industrial sector, managing this seems a bit challenging. How, for example, do we allocate our exposure to these risk factors? The solution is to use the factor loadings—these quantities determine the relative importance of the region and industry systemic factors. Recall, from Chap. 3, that one half of an obligor's weight is allocated to its geographic region, while the remaining half is assigned to its industrial sector. At most, therefore, an individual obligor's exposure can be assigned to two distinct factors. A public-sector entity, which has no obvious industrial classification, receives a full weight to its region. An advantage of this approach is the incorporation of the obligor's factor loadings (indirectly) into our index.

Mathematically, this jth risk factor's exposure assignment—which we will denote as $\bar{\mathcal{V}}_j$—is thus simply,

$$\bar{\mathcal{V}}_j = \sum_{i=1}^{I} \mathrm{B}_{ij} c_i, \tag{5.25}$$

for $j = 1, \ldots, J$ and where, as usual, c_i denotes the ith obligor's exposure at default. In other words, we basically weight the individual exposure estimates by their factor loadings. In practice, this is a very straightforward matrix multiplication,

$$\bar{\mathcal{V}} = \mathrm{B}^T c, \tag{5.26}$$

where $\bar{\mathcal{V}} \in \mathbb{R}^{J\times 1}$, $c \in \mathbb{R}^{I\times 1}$ are $\mathrm{B} \in \mathbb{R}^{I\times J}$.[31] One might be tempted to replace exposure with economic capital (i.e., $\mathcal{E}_i(\alpha^*)$) in Eq. 5.26; it is, in fact, economic-capital concentration that we seek to describe. This turns out to be a bad idea. Were we to use economic-capital in the construction—even indirectly—in the establishment of our concentration metric, we would find ourselves awkwardly having economic-capital embedded in both sides of our regression relationship. Lesser indiscretions have been classified as a criminal act in statistical circles.

Figure 5.4 summarizes, for an arbitrary portfolio during 2020, the percentage allocation of exposure—following from Eq. 5.26—to our $J = 24$ regional and sectoral categorizations. Since the actual amounts and systemic-risk factor identities are relatively unimportant—but we wish to be able to compare to further transformations—the values are illustrated in numbered percentage terms. The top four categories consume almost half of the total exposure. While concentration is typical in most real-life commercial lending and investment portfolios, we should

[31] Many, or even most, of the B factor-loading values are zero. This is a consequence of the overidentifying factor-loading constraints introduced in Chap. 3. Incidentally, it makes no practical difference if one uses the raw factor loadings, β, or their normalized equivalents, B.

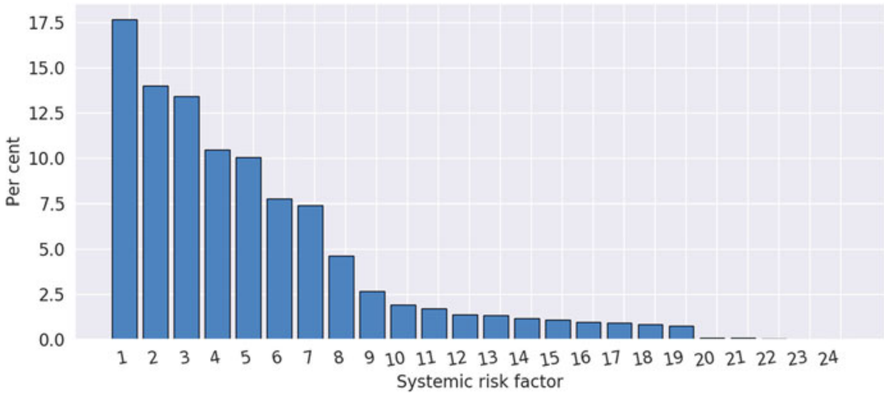

Fig. 5.4 *Exposure by risk factor*: The preceding graphic displays, for an arbitrarily selected portfolio from 2020, the percentage factor-loaded exposure allocations by systemic risk factor computed using Eq. 5.26.

be cautious to read too much into Fig. 5.4. In this intermediate step, we use only the relatively uninformative factor loadings.

Having used factor loadings to aggregate portfolio exposure, the next step involves the incorporation of the systemic-factor dependence structure. Technically, this takes the form of a projection. In particular, the next step is

$$\begin{aligned}\tilde{\mathcal{V}} &= \underbrace{\begin{bmatrix} 1 & \rho_{12} & \cdots & \rho_{1J} \\ \rho_{21} & 1 & \cdots & \rho_{2J} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{J1} & \rho_{J2} & \cdots & 1 \end{bmatrix}}_{\Omega} \begin{bmatrix} \mathcal{V}_1 \\ \vdots \\ \mathcal{V}_J \end{bmatrix}, \\ &= \Omega \bar{\mathcal{V}}. \end{aligned} \tag{5.27}$$

where $\Omega \in \mathbb{R}^{J \times J}$ is the factor-correlation matrix estimated in Chap. 3 and $\tilde{\mathcal{V}} \in \mathbb{R}^{J \times 1}$ is a projected vector of exposure amounts. Technically, $\tilde{\mathcal{V}}$ is a linear projection of $\bar{\mathcal{V}}$ using the projection matrix, Ω—it is essentially a linear (dimension-preserving) mapping from $\mathbb{R}^{J \times 1}$ to $\mathbb{R}^{J \times 1}$. If all of the systemic factors were to be orthonormal, then Ω would be an identity matrix, and the original exposure figures would be preserved. If there is strong positive correlation between the systemic risk factors, conversely, we would expect to see a relative increase in the relative allocation to each risk-factor category.

To see the implications of projecting correlation information onto our factor-arranged exposure outcomes, let's directly compare the allocations $\bar{\mathcal{V}}$ and $\tilde{\mathcal{V}}$ in percentage terms. These results are summarized in Fig. 5.5. Projecting the correlation matrix onto the exposure amounts clearly, and dramatically, changes

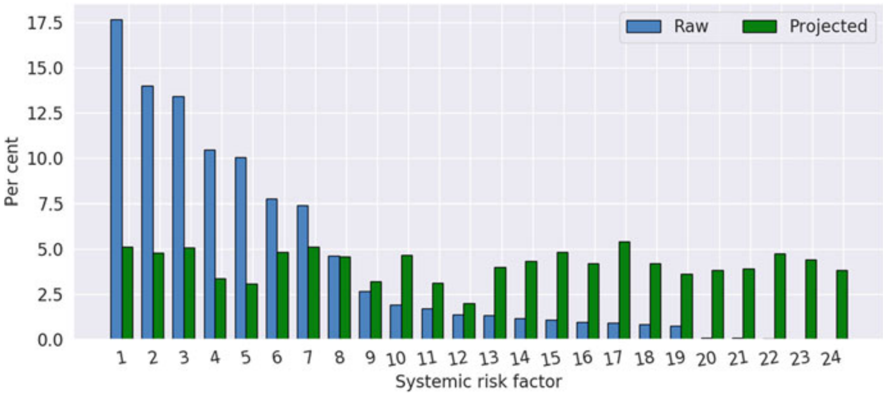

Fig. 5.5 *Incorporating systemic dependence*: The preceding graphic reproduces the percentage factor-loaded exposure allocations from Fig. 5.4, but also adds in the percentage correlation-matrix projected outcomes from Eq. 5.27. The positively correlated risk-factor system flattens out the concentration profile.

the concentration profile. Specifically, projecting $\bar{V}$ flattens out the profile implying, when taking into account factor correlations, that the deviations in concentration are rather different than those suggested by the factor-loading-driven result in Fig. 5.4.

What exactly is going on? The answer, of course, begins with the structure of the systemic risk-factor correlation matrix itself. The risk factors are, as indicated and discussed in previous chapters, a positively correlated system. We can make this a bit more precise. Using an eigenvalue decomposition, it is possible to orthogonalize a correlation matrix into a set of so-called principal components.[32] A useful byproduct of this computation is that one can estimate the amount of variance, in the overall system, explained by each individual orthogonal factor. This provides useful insight into the degree of linear correlation between the raw variables in one's system. Table 5.4 reviews the results of this analysis—applied to our factor correlation matrix, Ω—for the five most important orthogonal factors. The single, most important, orthogonal factor explains in excess of 50% of the overall variance, while the top five (of 24) principal components cover about three quarters of total system variance.

The results in Table 5.4 strongly underscore the rather dependent nature of our risk-factor system and help explain Fig. 5.5. A few systemic risk factors appear to demonstrate a high degree of correlation when we simply sum over the exposures. Controlling for the (generally high) level of linear dependence between these individual systemic risk factors, however, is essential to capturing the model's picture of portfolio concentration.

The final step involves the transformation of the individual values in Eq. 5.27 into a readable and interpretable index. This basically amounts to normalization; it

[32] A fantastic reference for this practical statistical technique is Jollliffe [18]. See Golub and Loan [12] for more on the eigenvalue decomposition.

Table 5.4 *Systemic correlation*: The underlying table illustrates, using the principal-components technique, the amount of variance explained by the five most important orthogonalized factors derived from our systemic correlation matrix, Ω. These results strongly underscore the rather dependent nature of our risk-factor system.

Principal component	% variance explained
1	55.6%
2	8.8%
3	4.6%
4	4.3%
5	3.3%
Total	76.7%

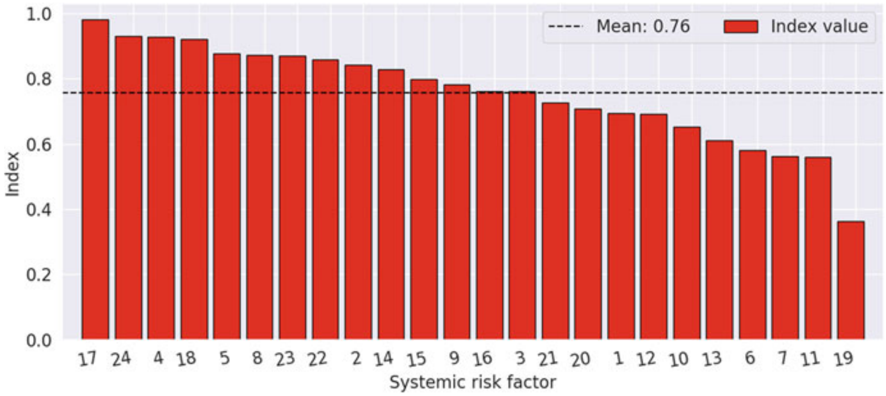

Fig. 5.6 *The concentration index*: The preceding graphic displays—conditional on our arbitrarily selected portfolio from 2020—the set of J distinct concentration index values arranged in descending order.

is consequently transformed into a value in the unit interval as,

$$\hat{\mathcal{V}} = \frac{\tilde{\mathcal{V}}}{\max(\tilde{\mathcal{V}})}. \tag{5.28}$$

Other normalizations are certainly possible. The largest concentration index value takes, of course, a value of unity. $\hat{\mathcal{V}}_j$ thus represents the final concentration-index value for the jth risk factor.

Figure 5.6 illustrates, using Eq. 5.28, the set of J distinct concentration index values displayed in descending order. These values are, once again, conditional on our arbitrary 2020 portfolio, although the results do not appear to vary considerably across time. This is certainly due to high levels of persistence in the portfolio and infrequently updated through-the-cycle parameters.[33] Actual concentration values

[33] It takes time to change the composition of any buy-and-hold portfolio.

for a given risk-owner, or exposure, fall in the range [0.36,1] with a mean of about 0.75. The ordering is broadly similar to that observed in Fig. 5.4, but there are a few surprises.

A separate concentration index is required for each individual credit obligor, but there are only J concentration-index values. The factor-loading matrix comes, once again, to the rescue by offering a logical approach to allocation of the sectoral and regional concentration values to the obligor level. Practically, this amounts to (yet) another matrix multiplication,

$$\mathcal{V} = \mathrm{B}\hat{\mathcal{V}}, \tag{5.29}$$

where $\mathcal{V} \in \mathbb{R}^{I\times 1}$ is the final concentration index. Each element of $\mathcal{V}$ thus provides some assessment of the relative concentration of this individual position. The larger the value, of course, the more a given position leans towards the portfolio's inherent concentrations.

Having constructed the concentration index outcomes, it may still not be immediately obvious what is going on. To provide a bit more transparency, and hopefully insight, we will begin from Eq. 5.29, and gradually work forward as follows, while keeping all the individual ingredients

$$\begin{aligned}
\mathcal{V} &= \underbrace{\mathrm{B}\hat{\mathcal{V}}}_{\substack{\text{Equation}\\ 5.29}}, \\
&= \underbrace{\mathrm{B}\left(\frac{\tilde{\mathcal{V}}}{\max(\tilde{\mathcal{V}})}\right)}_{\substack{\text{Equation}\\ 5.28}}, \\
&= \frac{\mathrm{B}}{\max(\Omega\bar{\mathcal{V}})}\underbrace{\Omega\bar{\mathcal{V}}}_{\substack{\text{Equation}\\ 5.27}}, \\
&= \frac{\mathrm{B}\Omega}{\max\left(\Omega\mathrm{B}^T c\right)}\underbrace{\mathrm{B}^T c}_{\substack{\text{Equation}\\ 5.26}}, \\
&= \left(\frac{\overbrace{\mathrm{B}\Omega\mathrm{B}^T}^{I\times I}}{\underbrace{\max\left(\Omega\mathrm{B}^T c\right)}_{\text{Scalar: } 1\times 1}}\right)c.
\end{aligned} \tag{5.30}$$

An element in the preceding decomposition of the $\mathcal{V}$ computation should look familiar. $B\Omega B^T \in \mathbb{R}^{I \times I}$ is the normalized, factor-loaded, systemic risk-factor correlation matrix. A standardized factor-loaded correlation matrix is thus projected onto the vector of exposure contributions, c, to yield our proposed concentration index. While not perfect, it does appealingly capture a number of key elements of concentration: current portfolio weights, systemic factor loadings, and the correlation structure of these common factors.

Colour and Commentary 57 (INDEXING CONCENTRATION)*: A central task of the credit-risk economic-capital simulation engine is to capture the interplay between diversification and concentration in one's portfolio. This aspect needs to be properly reflected in any approximation method. It takes on particular importance given our previous decision to lean heavily upon a regulatory motivated approximation that entirely assumes away idiosyncratic risk. A possible solution to this question involves the construction of a* concentration *index, which can act as an additional explanatory variable in our regression model. Our proposed choice involves a linear projection of the individual credit obligor exposures at default. The projection matrix is a function of the (normalized) factor loadings and the factor correlation. In this way, the resulting index directly incorporates three important drivers of concentration risk: the factor loadings, systemic correlations, and the current portfolio composition.*

5.2.5 The Full Default Model

We've clearly established the need to incorporate additional idiosyncratic explanatory variables into our base regression model. The previously defined concentration index is a natural candidate; there are many others. Indeed, there is no shortage of possibilities. This means that model selection is actually hard work. The consequence is a very large set of possible extensions and variations to consider.[34] This situation can only be resolved with a significant amount reflection, examination of error graphics and model-selection criteria as well as trial-and-error. Ultimately, we have opted for six *new* explanatory variables to add to our original Eq. 5.18. Despite our best efforts, the revised and extended model should not be considered as an optimal choice. It is better to view it as a sensible element of the set of possible

[34] Even with a few dozen possible explanatory variables, the power set (i.e., set of all possible subsets) is depressingly large.

approximation models. With this in mind, it is described as

$$\overbrace{\ln\left(\mathcal{E}_i^{(d)}(\alpha^*)+\mathbb{E}\left(L_i^{(d)}\right)\right)}^{y_i^{(d)}}$$
$$=\overbrace{\underbrace{\xi_0+\xi_1\ln\left(\tilde{p}_i(\alpha_z^*)\right)+\xi_2\ln\left(\mathbb{E}(\gamma_i)\right)+\xi_3\ln(c_i)}_{\text{Base default model}}+\underbrace{\sum_{k=4}^{\kappa}\xi_k X_k^{(d)}}_{\substack{\text{Additional}\\ \text{explanatory}\\ \text{variables}}}}^{\sum_{k=0}^{\kappa}\xi_k X_i^{(d)}}+\epsilon_i^{(d)}, \tag{5.31}$$

for $i = 1, \ldots, I$ risk-owner risk contributions.

Table 5.5 summarizes the parameter and goodness-of-fit statistics associated with the revised model. Comparing the results to the base approximation summary found in Table 5.3 on page 298, we observe a number of differences. First of all, there is significant improvement in the (in-sample) goodness-of-fit measures. The proportional root-mean squared error is reduced by a factor of five, while our two proportional *absolute* errors are cut by more than half. The median absolute error

Table 5.5 *Full default-approximation model*: This table summarizes the results of the full-blown default economic-capital approximation model including a number of additional regressors as described in Eq. 5.31. The overall fit is substantially improved relative to the base model displayed in Table 5.3.

Quantity	Description	Estimate	Standard error	t-statistic	p-value
ξ_0	Intercept	−4.02	0.09	−45.20	0.00
ξ_1	(Extreme) conditional default probability	0.80	0.01	81.18	0.00
ξ_2	Loss-given-default	1.04	0.01	87.51	0.00
ξ_3	Exposure	1.02	0.00	312.38	0.00
ξ_4	Obligor credit rating	0.03	0.00	13.53	0.00
ξ_5	Concentration index	3.03	0.06	52.87	0.00
ξ_6	Public-sector indicator variable	0.15	0.04	4.31	0.00
ξ_7	Large exposure indicator variable	0.16	0.02	9.25	0.00
ξ_8	Small exposure indicator variable	0.05	0.02	2.69	0.01
Root mean *squared* error		9%			
Mean absolute error		6%			
Median absolute error		4%			
R^2		0.988			

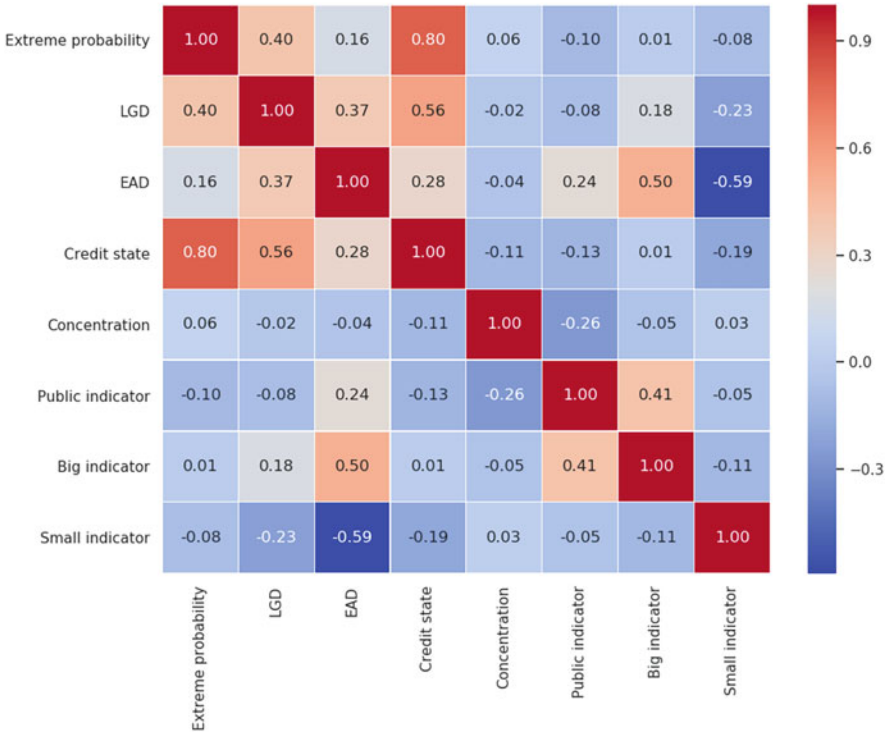

Fig. 5.7 *Default instrument correlation*: This heat-map illustrates the cross-correlation between the key explanatory variables included in the extended regression model summarized in Table 5.5.

is now only about 4% of the base estimate. The amount of variance explained also increases sharply to 0.99. Finally, all of the new parameters are statistically significant and there is no dramatic (economic) change in the ξ_1 to ξ_3 parameters.[35] By almost any metric, the extended model represents a serious upgrade.

The new explanatory variables warrant some additional explanation beyond the brief descriptions provided in Table 5.5. The concentration index was already defined in the previous section and, as we had hoped, turns out to be extremely statistically significant. It clearly provides additional information to the model not encapsulated in the systemic-risk-focused extreme conditional default probability. The obligor credit rating, entered as an integer from 1 to 20, somewhat surprisingly improves the overall fit. We would expect—and, in fact, find in Fig. 5.7—positive correlation between the credit rating and the extreme conditional default probability.

[35] The intercept variable does, however, change significantly.

Despite a slight danger of collinearity, the addition of this piece of information appears to improve the overall model.[36]

We also found it extremely useful to add *three* additional indicator (i.e., dummy) variables. The first identifies public-sector entities in the portfolio as

$$\mathbb{I}_{P_i} = \begin{cases} 1 : \text{the } i\text{th credit obligor is a public-sector entity} \\ 0 : \text{the } i\text{th credit obligor is a corporation} \end{cases}. \tag{5.32}$$

The second isolates particularly large firms with the following logic:

$$\mathbb{I}_{B_i} = \begin{cases} 1 : c_i \geq \text{percentile}(c, 0.90) \\ 0 : c_i < \text{percentile}(c, 0.90) \end{cases}, \tag{5.33}$$

or, in words, the largest 10% of credit exposures are flagged. The small exposure flag works in the same manner on the smallest 10% of exposures. For completeness, it is defined as

$$\mathbb{I}_{S_i} = \begin{cases} 1 : c_i \leq \text{percentile}(c, 0.10) \\ 0 : c_i > \text{percentile}(c, 0.10) \end{cases}. \tag{5.34}$$

While these three choices are clearly important, statistically significant explanatory variables, it is difficult to directly understand their role. Examination of the initial model results suggests that the regulatory capital computation has particular difficulty with public-sector, unusually large, and quite small exposures. This probably relates to underlying concentration and parameter-selection questions. We can conceptualize these final three explanatory variables as allowing the model to better handle exceptions. These indicator variables also do not, as seen in Fig. 5.7, exhibit much correlation with the other explanatory variables. The only exception is, unsurprisingly, the large-exposure indicator and the exposure-at-default variable.

One area of investigation is conspicuous by its absence: random recovery. Nothing in our instrument-variable selection provides any insight into this important dimension. A number of avenues was nonetheless investigated. Key parameters—such as the shape parameters of the underlying beta distribution or the recovery volatility—were added to the model without success. Modified, more severe, loss-given-default values were also explored—again using the characteristics of the underlying beta distributions—without any improvement in model results. On the contrary, these efforts generally led to a deterioration of the overall fit.

Figure 5.8 revisits the error analysis performed in Fig. 5.3 using our extended regression model summarized in Eq. 5.31. To aid in our interpretation, the original outliers—from the base default model—are also displayed. Although the general pattern is qualitatively similar, the range of estimation error is reduced by an order of magnitude. While much of the impact relates to improved handling of the outliers,

[36] This high level of cross correlation should nonetheless be flagged as a potential weak link in the extended regression model.

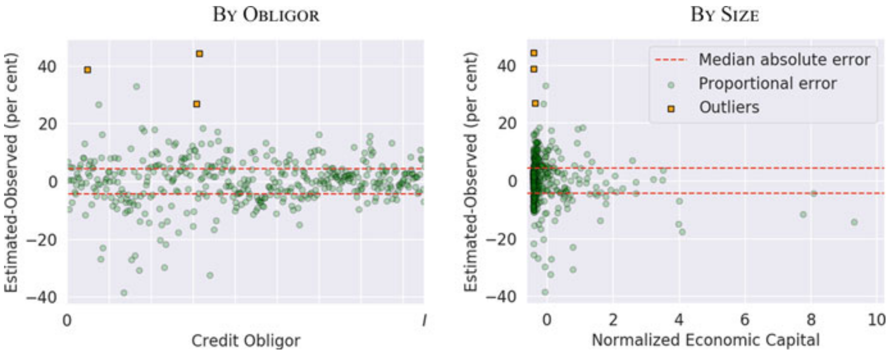

Fig. 5.8 *Revised default-error analysis*: The preceding graphics replicate the analysis displayed in Figure 5.3 using the extended regression model summarized in Eq. 5.31. The original outliers—from the initial model—are displayed to illustrate the overall improvement.

there is also a more general effect. The difficulty of approximating small, below-average economic-capital allocations is not entirely resolved; part of the issue relates to the proportional nature of the error measure, but it remains a weakness.

While imperfect, we conclude that this is a satisfactory approximation model. It exploits some basic relationships involved in the computation of economic capital, is relatively parsimonious, and generates a reasonable degree of fit and statistical robustness. Some weaknesses remain; most particularly relatively high cross correlation between a few explanatory variables as well as a tendency to significantly overestimate a subset of below-average exposures.

Colour and Commentary 58 (A WORKABLE DEFAULT APPROXIMATION MODEL)*: Given the complexity of the credit-risk economic-capital model, we should not expect to be able to construct a perfect approximator.*[a] *Exploiting our knowledge of the problem nonetheless allows us to build a regression model with fairly informative systemic and idiosyncratic explanatory variables. The final result is a satisfactory description of the underlying model. Satisfactory, in this context, implies a median proportional approximation error in the neighbourhood of 4%. While a lower figure would naturally be preferable, we can live with this result. Construction of such models, it should be stressed, is* not *a one-off exercise. The approximation model is re-estimated daily and the error structure is examined within our internal diagnostic dashboards. On an annual basis, a more formal model-selection (or rather verification) procedure is performed. Ultimately, we are always on the lookout for ways to improve this approximation.*

[a] If we could, then we might wish to dramatically rethink our model-estimation approach.

5.3 Approximating Migration Economic Capital

The preceding approximation is focused entirely on the notion of credit-risk economic-capital associated with default events. While it might be tempting to fold in the migration element into the previous approximation model, there, at least, *two* reasonably compelling reasons to separately approximate the default and migration effects. First, although default is a special case of migration, the loss mechanics differ importantly. This makes constructing a joint approximation more difficult. Although presumably still possible, separate structurally motivated approximation models are more natural to construct and easier to follow.

The second reason is more practical. Distinct default and migration models permit autonomous prediction of these two effects. This may be interesting for analytic purposes; understanding the breakdown between these two sources of risk can provide, for example, incremental insight into the model outcomes. It may also be necessary from a portfolio perspective. In certain stress testing or loan-impairment calculations, for example, we may not be directly interested in the migration allocation. In other cases, we may be interested *only* in migration effects. A joint model would make it difficult to manage such situations.

5.3.1 Conditional Migration Loss

Having justified this choice, let us turn our attention to a possible migration-approximation model. We begin, as in the default setting, by trying to exploit our knowledge of the basic calculation and adding some structure to the migration problem. Denoting $L_i^{(m)}$ as the credit-migration loss associated with the ith credit obligor at time $t+1$, we recall that it has the following form,

$$L_i^{(m)} = \underbrace{\Big(\mathbb{S}\left(S_{i,t+1}\right) - \mathbb{S}\left(S_{i,t}\right)\Big)}_{\Delta\mathbb{S}_i} \tau_i c_i, \tag{5.35}$$

where $\mathbb{S}_i$ is the credit spread, $S_{i,t}$ is the credit state at time t, τ_i is the modified spread duration, and c_i is the exposure of the ith position. The time-homogeneous credit spread, of course, is determined directly by the credit state.[37] The spread change (i.e., $\Delta\mathbb{S}_i$) is thus determined by any movement in an obligor's credit state from one time step to the next. The spread movement is then multiplied by the modified spread duration and position exposure to yield an approximation of the migration loss.

Can we construct, using similar logic to the default case, a conditional migration loss estimate? The short answer is yes, there is clearly some conditionality in the

[37] The actual parametrization of these spreads is described in Chap. 3.

migration loss. We will denote, to ease the notation, the current credit state, at time t, as S_{i0}. This is known. The actual loss outcome depends upon the credit-state value at time $t+1$; we'll call this S_{i+}. To formalize this idea, let us rewrite Eq. 5.35 in a similar form and applying the expectation operator to both sides,

$$L_i^{(m)} = \underbrace{\left(\mathbb{S}\left(S_{i+}\right) - \mathbb{S}\left(S_{i0}\right)\right)\tau_i c_i}_{\text{Equation 5.35}}, \tag{5.36}$$

$$\begin{aligned}
\mathbb{E}\left(L_i^{(m)} \middle| S_{i+} = s\right) &= \mathbb{E}\left(\left(\mathbb{S}\left(S_{i+}\right) - \mathbb{S}\left(S_{i0}\right)\right)\tau_i c_i \middle| S_{i+} = s\right), \\
&= \left(\underbrace{\mathbb{E}\left(\mathbb{S}\left(S_{i+}\right) \middle| S_{i+} = s\right)}_{\text{This is a random variable}} - \underbrace{\mathbb{S}\left(S_{i0} \middle| S_{i+} = s\right)}_{\text{This is a constant}}\right)\tau_i c_i, \\
&\approx \left(\mathbb{S}(s) - \mathbb{S}\left(S_{i0}\right)\right)\tau_i c_i.
\end{aligned}$$

While perhaps not earth-shattering, this expression underscores the fact that the only really uncertain element in the credit-migration calculation is the future credit-state value, S_{i+}.[38] The million dollar question relates to our choice of conditioning variable, s. One option would be to set $s = \mathbb{E}(S_{i+})$.

For most credit counterparties—save those already at the extremes of the credit scale—credit migration can involve either an improvement or deterioration of their credit quality. This, in turn, implies either profit or loss associated with the associated credit-migration effects described in Eq. 5.36. Naturally, when looking to the extreme tail of the loss distribution, we would expect the vast majority of credit-migration outcomes to involve downgrade, credit-spread widening, and ultimately loss. We will thus, for our purposes, focus our attention predominately on extreme downside outcomes.

Let's begin with the basics. S_{i+} is a random variable and its expectation will depend upon an entity's appropriate transition probabilities. Its unconditional expectation is

$$\begin{aligned}
\mathbb{E}(S_{i+}) &= \sum_{k=1}^{q} \underbrace{\mathbb{P}\left(S_{i,t+1} = k \middle| S_{i,t} = S_{i0}\right)}_{\text{Transition probability}} \cdot k, \\
&= \sum_{k=1}^{q} p_{kS_{i,t}} \cdot k,
\end{aligned} \tag{5.37}$$

[38] The final step is something of an approximation. Although we developed an expression for the credit spreads in Chap. 3, it does not turn out to be helpful in this case.

recalling that $q = 21$ is the number of distinct credit states (including default) and where the individual transition probabilities (i.e., the p's) are simply elements from the $S_{i,t}$th row of the transition matrix, P. The expected one-step forward credit state is simply the transition-probability-weighted sum of the possible future credit states. Transition matrices nonetheless exhibit a significant amount of inertia.[39] The actual expected credit state, in one period's time, is unlikely to stray very far away from the current credit state. Consequently, Eq. 5.37—in its current form—will generally *not* be terribly informative about credit-migration economic-capital outcomes.

How might we resolve this issue? The worst-case migration loss should logically occur when the (unknown) future credit state, S_{i+}, takes its worst possible level. Under our t-threshold credit risk model, the actual credit-state outcome depends on draws from the systemic-risk and mixing-variable distributions. We might, therefore, employ the same trick as we used in the default setting; select an extreme systemic risk-factor variable outcome, z^*. The conditional expectation of S_{i+} for the ith credit obligor, given z^* is a better candidate for the as-yet-undetermined s in Eq. 5.36. It can be correspondingly written as

$$\begin{aligned}
\mathbb{E}(S_{i+}|z^*) &= \sum_{k=1}^{q} \mathbb{P}\left(S_{i,t+1} = k \middle| S_{i,t} = S_{i0}, z^*\right) \cdot k, \qquad (5.38)\\
&= \underbrace{\sum_{k=1}^{q-1} \mathbb{P}\left(K_{S_{i,t}}(k) \leq \Delta X_i \leq K_{S_{i,t}}(k+1) \middle| z^*\right) \cdot k}_{\text{Migration}}\\
&\quad \underbrace{+\mathbb{P}\left(\Delta X_i \leq K_{S_{i,t}}(q) \middle| z^*\right) \cdot q}_{\text{Default}},\\
&= \sum_{k=1}^{q-1} \left(\mathbb{P}\left(\Delta X_i \leq K_{S_{i,t}}(k+1) \middle| z^*\right) - \mathbb{P}\left(\Delta X_i \leq K_{S_{i,t}}(k) \middle| z^*\right)\right)\\
&\quad \cdot k + \mathbb{P}\left(\Delta X_i \leq K_{S_{i,t}}(q) \middle| z^*\right) \cdot q,\\
&= \sum_{k=1}^{q-1} \underbrace{\left(\Phi\left(\frac{\sqrt{\frac{w^*}{\nu}} K_{S_{i,t}}(k+1) - \alpha_i z^*}{\sqrt{1-\alpha_i^2}}\right) - \Phi\left(\frac{\sqrt{\frac{w^*}{\nu}} K_{S_{i,t}}(k) - \alpha_i z^*}{\sqrt{1-\alpha_i^2}}\right)\right)}_{p_i(z^*,k)}
\end{aligned}$$

[39] We saw this behaviour in Chap. 3 and will discuss it again in Chap. 7.

$$
\cdot k + \underbrace{\Phi\left(\frac{\sqrt{\frac{w^*}{\nu}}F_{\mathcal{T}_\nu}^{-1}(p_i) - \alpha_i z^*}{\sqrt{1-\alpha_i^2}}\right)}_{p_i(z^*,q)\equiv p_i(\alpha^*)} \cdot q,
$$

$$
= \underbrace{\sum_{k=1}^{q-1} p_i(z^*, k)\cdot k}_{\text{Migration}} + \underbrace{p_i(z^*, q)\cdot q}_{\text{Default}},
$$

where

$$
K_{S_{i,t}}(j) = F_{\mathcal{T}_\nu}^{-1}\left(\sum_{w=j}^{q} p_{w,S_{i,t}}\right), \tag{5.39}
$$

is the ith obligor credit-migration threshold for movement into state j.[40] As in the default setting, we use a fixed, relatively neutral realization of the mixing variable, w^*. Equation 5.38 is a rather unwieldy expression, but we already have experience with these conditional-default probability expressions and, perhaps most importantly, they are readily and quickly computed.

We are not quite ready for prime-time usage. The approximation in Eq. 5.38 is a bit too conservative; the simple reason is that it includes the default probability. This component, by construction, will be carved out of the credit-migration economic-capital and allocated to the default category. If we exclude the final default term, however, our expectation will lose a significant amount of probability mass. To solve this, we simply rescale our conditional probabilities excluding the default element. This amounts to an expected downside non-default transition. Practically, this is easily accomplished by defining,

$$
\tilde{p}_i(z^*, k) = \frac{p_i(z^*, k)}{\sum_{j=1}^{q-1} p_i(z^*, j)}, \tag{5.40}
$$

for $k = 1, \ldots, q-1$ and then writing our (revised) conditional non-default credit state expectation as,

$$
\mathbb{E}\left(S_{i+}^{(m)}\middle| z^*\right) = \sum_{k=1}^{q-1} \tilde{p}_i(z^*, k)\cdot k. \tag{5.41}
$$

[40] These credit migration notions and thresholds are extensively discussed in Chap. 2.

We have thus made a slight, but important change, from the base definition in Eq. 5.38. Equation 5.41 describes the expected (non-default) credit state for the ith obligor in the next period, conditional on an unpleasant realization of the systemic risk factor.[41] This shift in focus is captured by the use of $S_{i+}^{(m)}$ rather than the original S_{i+} notation.

Equipped with this, model consistent, estimator for the worst-case (non-default) credit deterioration, we may return to the business of approximating credit-risk economic-capital allocation. Again, we motivate our approach through the Basel IRB method. If we denote the capital allocation associated with migration risk as $\mathcal{A}_i^{(m)}(\alpha_z^*)$, we might approximate it as,

$$\begin{aligned}
\mathcal{A}_i^{(m)}(\alpha_z^*) &\approx \overbrace{\underbrace{\mathbb{E}\left(L_i^{(m)} \middle| \mathbb{E}\left(S_{i+}^{(m)} \middle| z^*\right)\right)}_{\text{Worst-case loss}} - \underbrace{\cancel{\mathbb{E}\left(L_i^{(m)} \middle| \mathbb{E}(S_{i+})\right)}}_{\mathbb{E}\left(L_i^{(m)}\right)\approx 0}}^{\text{Unexpected loss}}, \qquad (5.42)\\
&\approx \underbrace{\left(\mathbb{S}\left(\mathbb{E}\left(S_{i+}^{(m)} \middle| z^*\right)\right) - \mathbb{S}\left(S_{i0}\right)\right)}_{\Delta\mathbb{S}_i(z^*)\text{: Worst-case spread movement}} \tau_i c_i,\\
&\approx \Delta\mathbb{S}_i(z^*)\tau_i c_i.
\end{aligned}$$

Unlike the default case, we ignore the unconditional expected loss. Since it is generally quite close to zero, we exclude it from our basic structure.[42]

While Eq. 5.42 appears to be a sensible construct—built in the image of the default-risk approximation—it has an important potential structural disadvantage. There are only 20 (non-default) credit states that can be forecast with a correspondingly small number of spread outcomes. Equation 5.40 will, without some form of intervention, thus produce only a small number of discrete credit-spread outcomes for a given obligor's modified spread duration and exposure. The credit-migration economic-capital engine, however, operates under no such constraint. It can provide a (quasi-)continuum of possible migration–risk allocations. This would suggest that our approximation—even if it differs by one notch on the credit-state forecast—can exhibit rather important approximation errors. Our solution is to simply permit non-integer (i.e., fractional) credit-state forecasts and determine the appropriate spread outcome through the use of linear interpolation. This is essentially a trick—a fractional credit rating does not really exist—permitting us to transform our estimator from discrete to continuous space. This small adjustment provides our estimator with significantly more flexibility.

[41] As in the default section, z^* is the α_z^* quantile of the underlying systemic risk factor distribution.

[42] This stems from the high degree of inertia in credit-rating transition and, for most obligors, the relatively symmetric probability of upgrade and downgrade.

We have also ignored the expected-shortfall dimension. Practically, Eq. 5.42 is a quantile-based estimator. If we wish to consider the expected shortfall metric, then we will need to solve the following integral:

$$\begin{aligned}\mathcal{E}_i^{(m)}(\alpha_z^*) &= \frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^{1}\mathcal{A}_i^{(m)}(x)dx, \\ &= \tau_i c_i \cdot \underbrace{\frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^{1}\Delta\mathbb{S}_i\Big(\Phi^{-1}(1-x)\Big)dx}_{\Delta\tilde{\mathbb{S}}_i(z^*)}.\end{aligned} \tag{5.43}$$

Practically, only a small amount of incremental effort is required to numerically integrate Eq. 5.43. We thus employ the integrated worst-case spread change, $\Delta\tilde{\mathbb{S}}_i(z^*)$, as the base migration loss estimator found in Eq. 5.43.

To better understand the kernel of our approximation, we return to our simple example introduced in Table 5.1 on page 290. Naturally, we will now consider the case of credit-migration economic capital. Some of the credit obligor's details are reproduced for convenience and its modified spread duration has also been added. Unsurprisingly, the unconditional expected rating remains extremely close to the original level. Conditional on extreme outcomes of our systemic and mixing variables, however, the expected (or shocked) credit state belongs to PD class 13. If we use the integrated, expected-shortfall form, from Eq. 5.43, this increases to a 13.2 rating. Of course, such a PD categorization does not actually exist. Through appropriate interpolation along the spread curve, however, this leads to a 56 and 63 basis-point credit widening for the VaR-motivated and integrated worst-case spread movements, respectively. These values translate into forecast credit losses of about EUR 0.94 and 1.05 billion. Since the actual economic-capital value is roughly EUR 0.87 million, both of these forecasts appear to overshoot the target.[43]

Figure 5.9 provides a potentially helpful visualization of the difference between unconditional and conditional (non-default) transition probabilities. Using our sample obligor from Table 5.6, the three flavours of transition probability are plotted side-by-side across the entire (non-default) credit spectrum. The unconditional transition probabilities, as one would expect, are roughly symmetric and centred around the current credit state: PD11.[44] The α^* confidence-level (non-default) migration probabilities—both VaR-motivated and integrated from Eqs. 5.42 and 5.43—behave in a decidedly different fashion. The likelihood of remaining in the current credit state is dramatically reduced and the remaining probability

[43] Interestingly, the model-based worst-case credit-spread widening, imputed from the true credit-migration estimate, is readily calculated from Table 5.6 as $\frac{\mathcal{E}_i^{(m)}(\alpha^*)}{c_i\tau_i} = \frac{868{,}247}{65{,}000{,}000\times 2.6} \approx 51$ basis points.

[44] For this credit rating, upgrade and downgrade probabilities, on an unconditional basis, appear to be roughly equivalent. This is usually true, but it starts to break down as one moves to the boundaries of the credit spectrum.

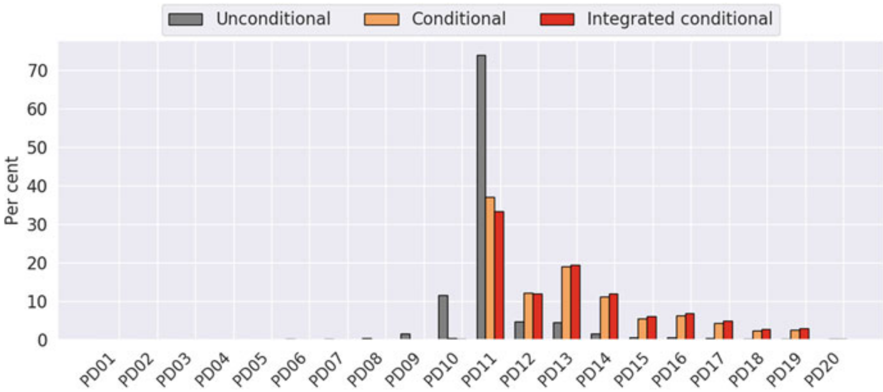

Fig. 5.9 *Credit-migration conditionality*: The preceding figure compares the unconditional transition probabilities to the $\alpha_z^* = 0.9997$ confidence-level (non-default) conditional transition probabilities for the sample obligor highlighted in Table 5.6. The unconditional, VaR-motivated, and integrated versions of these migration probabilities are displayed. The conditional quantities are, in stark contrast to their rather symmetric unconditional equivalents, skewed towards credit deterioration.

Table 5.6 *Extending our simple example*: The underlying table summarizes, for the same specific exposure illustrated in Tables 5.1 and 5.2, the various quantities and calculations involved in the construction of a conditional forecast of the migration loss.

Quantity	Definition	Value
Degrees of freedom	ν	70
Confidence level	α_z^*	0.9997
Exposure	c_i	65,000,000
Modified spread duration	τ_i	2.6 yrs.
Original credit state	S_{i0}	11.0
(Unconditional) expected (non-default) rating	$\mathbb{E}\left(S_{i+}^{(m)}\right)$	11.0
(Conditional) expected (non-default) rating	$\mathbb{E}\left(S_{i+}^{(m)} \middle\vert z^*\right)$	13.0
(Conditional) integrated expected (non-default) rating	$\frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^{1} \mathbb{E}\left(S_{i+}^{(m)} \middle\vert \Phi^{-1}(1-x)\right) dx$	13.2
Worst-case spread	$\Delta\mathbb{S}_i(z^*)$	56 bps.
Integrated worst-case spread	$\Delta\tilde{\mathbb{S}}_i(z^*)$	63 bps.
VaR-based forecast	$\widehat{\mathcal{A}_i^{(m)}(\alpha_z^*)}$	940,704
Expected-shortfall-based forecast	$\widehat{\mathcal{E}_i^{(m)}(\alpha_z^*)}$	1,059,842
Credit-migration economic-capital	$\mathcal{E}_i^{(m)}(\alpha^*)$	868,247

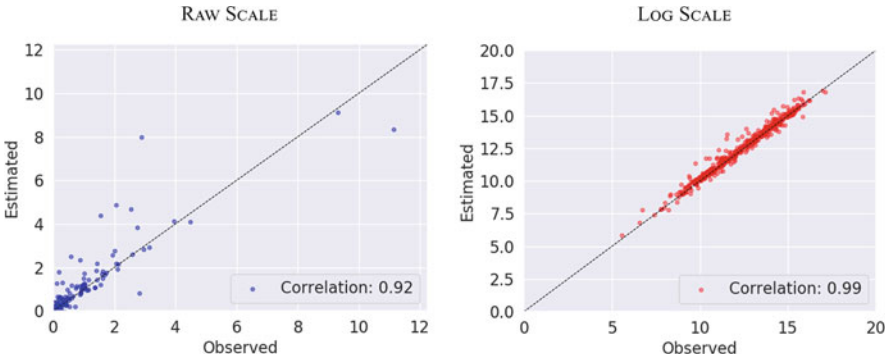

Fig. 5.10 *Degree of migration agreement*: The preceding graphics, in raw and logarithmic normalized units, illustrate the degree of agreement between our approximation in Eq. 5.43 and the observed simulation-based economic-capital contributions. The overall level of linear correlation is fairly respectable, although there is a decided tendency to overstate the level of risk.

mass is skewed strongly toward downgrade. The skew is slightly more pronounced for the integrated conditional transition likelihoods. It is precisely this information that we are attempting to harness for the approximation of credit migration. These worst-case transition probabilities thus—in a manner similar to conditional default probabilities in our earlier discussion—form the kernel of the migration-loss approximation.

Figure 5.10 generalizes the single obligor example in Table 5.6 to the entire portfolio. As before, the results are provided in raw and logarithmic normalized units. The idea is to help understand the general degree of agreement between our approximation in Eq. 5.43 and the observed simulation-based migration economic-capital contributions. The overall level of linear correlation is fairly respectable at 0.92 and 0.99 on the raw and logarithmic scales, respectively.

Our base migration-risk approximator has, similar to the result from Table 5.6, a decided tendency to somewhat misstate the precise level of risk. This is likely a structural feature of the approach culminating in Eq. 5.43. As in the default setting, conditioning solely on the systemic risk factor ignores idiosyncratic effects. Given the more symmetric nature of migration risk—relative to the skewed nature of default—restricting our approximation to downside systemic risk appears to manifest itself as overly conservative initial estimates. Figure 5.10 nonetheless suggests that, these shortcomings aside, it is a useful starting point.

Colour and Commentary 59 (The Migration Kernel)*: The decision to separately approximate migration and default, while sensible, presents a dilemma. We need to identify a reasonable first-order migration-loss approximator. A natural starting point is to lean on the default-approximation kernel*

(continued)

Colour and Commentary 59 (continued)
and see how it might be generalized. Happily, the notion of a conditional migration loss makes both logical and empirical sense. The only unknown quantity in migration loss is the obligor's future credit rating. Unconditionally, this is rather uninteresting, but the situation improves if we condition on an extreme outcome of the common systemic risk factor.[a] *Extending this idea leads to downward-skewed transition probabilities as well as worst-case rating downgrade and credit-spread estimates. These quantities are readily combined to produce an entirely reasonable, if somewhat conservative, initial migration-risk economic-capital approximator.*

[a] As in the default setting, the χ^2-distributed mixing variable is fixed at a fairly neutral value.

5.3.2 A First Migration Model

It will come as no surprise that we use the same linear regression structure—given its reasonable degree of success—as employed in the default setting. Indeed, we now have all of the necessary constituents to construct a base credit-migration approximation model. Consider the following form:

$$\begin{aligned}
\mathcal{E}_i^{(m)}(\alpha^*) &\approx \tau_i c_i \cdot \underbrace{\frac{1}{1-\alpha_z^*}\int_{\alpha_z^*}^{1} \Delta\mathbb{S}_i\Big(\Phi^{-1}(1-x)\Big)dx}_{\text{Equation 5.43}}, \qquad (5.44)\\
&\approx \Delta\tilde{\mathbb{S}}_i(z^*)\tau_i c_i,\\
\underbrace{\ln\left(\mathcal{E}_i^{(m)}(\alpha^*)\right)}_{y_i^{(m)}} &\approx \ln\underbrace{\left(\Delta\tilde{\mathbb{S}}_i(z^*)\tau_i c_i\right)}_{\text{Equation 5.43}},\\
y_i^{(m)} &= \theta_0 + \theta_1\underbrace{\ln\left(\Delta\tilde{\mathbb{S}}_i(z^*)\right)}_{X_{i,1}^{(m)}} + \theta_2\underbrace{\ln(\tau_i)}_{X_{i,2}^{(m)}} + \theta_3\underbrace{\ln(c_i)}_{X_{i,3}^{(m)}} + \epsilon_i^{(m)},
\end{aligned}$$

for $i = 1, \ldots, I$. Again applying natural logarithms, we have a similar form to the earlier default approximation in Eq. 5.18. This directly translates, of course, into the familiar OLS setting and reduces to the usual linear model,

$$y^{(m)} = X^{(m)}\Theta + \epsilon^{(m)}, \qquad (5.45)$$

where

$$y^{(m)} = \begin{bmatrix} \ln\left(\mathcal{E}_1^{(m)}(\alpha^*)\right) \\ \vdots \\ \ln\left(\mathcal{E}_I^{(m)}(\alpha^*)\right) \end{bmatrix} \tag{5.46}$$

and

$$X^{(m)} = \begin{bmatrix} 1 \ \ln\left(\Delta\tilde{\mathbb{S}}_1(z^*)\right) \ \ln(\tau_1) \ \ln(c_1) \\ \vdots \quad \vdots \quad \vdots \quad \vdots \\ 1 \ \ln\left(\Delta\tilde{\mathbb{S}}_I(z^*)\right) \ \ln(\tau_I) \ \ln(c_I) \end{bmatrix}. \tag{5.47}$$

As before, we need only employ the standard OLS estimator for $\Theta = \left[\theta_0\ \theta_1\ \theta_2\ \theta_3\right]^T$. Although the inputs are rather different, the base migration model is conceptually equivalent to the default case.

Table 5.7 provides the usual summary for our base, structurally motivated, credit-migration model shown in Eq. 5.45. All *four* regression coefficients are strongly statistically significant. The percentage of variance explained (i.e., R^2) attains a decent level. The other goodness-of-fit measures, however, show room for improvement. The proportional root-mean-squared value of almost 40% is somewhat better than in the base default model, but it suggests the presence of important outliers. Having already been in this situation, let's skip directly to the extension of Eq. 5.45 to include a wider range of idiosyncratic explanatory variables.

Table 5.7 *A first migration-approximation model*: This table summarizes the results of the base migration economic-capital approximation model. It has an intercept and three logically motivated explanatory variables. The results, while not overwhelming, are reasonably encouraging.

Quantity	Description	Estimate	Standard error	t-statistic	p-value
θ_0	Intercept	−2.04	0.17	−11.90	0.00
θ_1	Worst-case spread	0.69	0.02	29.37	0.00
θ_2	Spread duration	1.01	0.02	64.62	0.00
θ_3	Exposure	1.01	0.01	143.70	0.00
Root mean *squared* error		40%			
Mean absolute error		23%			
Median absolute error		17%			
R^2		0.886			

5.3.3 The Full Migration Model

As in the default case, an extension of the base model is necessary. The full migration-risk regression model—which is once again the result of a significant amount of analysis—is described as:

$$\overbrace{\ln\left(\mathcal{E}_i^{(m)}(\alpha^*)\right)}^{y_i^{(m)}} = \overbrace{\underbrace{\theta_0 + \theta_1 \ln\left(\Delta\tilde{\mathbb{S}}_i(z^*)\right) + \theta_2 \ln(\tau_i) + \theta_3 \ln(c_i)}_{\text{Base migration model}} + \underbrace{\sum_{k=4}^{\kappa} \theta_k X_k^{(m)}}_{\substack{\text{Additional}\\ \text{explanatory}\\ \text{variables}}}}^{\sum_{k=0}^{\kappa} \theta_k X_i^{(m)}} + \epsilon_i^{(m)}, \tag{5.48}$$

for $i = 1, \ldots, I$ risk-owner risk contributions. The interesting part of Eq. 5.48 relates to the identity of the additional explanatory variables.

Table 5.8 displays the results of the final credit-migration economic-capital approximation model. All model parameters are statistically significant. Coincidentally, *five* additional explanatory variables were also selected. As in the default setting, the integer-valued credit rating and the concentration index turn out to be surprisingly useful in explaining credit-migration risk outcomes. Two additional indicator (i.e., dummy) variables also prove helpful. The first identifies public-sector

Table 5.8 *Full migration-approximation model*: This table provides the key summary statistics for the full credit-migration economic-capital approximation model. It does not fit quite as well as the default approach, but it appears to do reasonably well.

Quantity	Description	Estimate	Standard error	t-statistic	p-value
θ_0	Intercept	−0.62	0.17	−3.70	0.00
θ_1	Worst-case spread	0.93	0.02	51.40	0.00
θ_2	Spread duration	0.98	0.01	123.43	0.00
θ_3	Exposure	1.01	0.00	239.04	0.00
θ_4	Obligor credit rating	−0.08	0.00	−17.66	0.00
θ_5	Concentration index	1.42	0.07	19.74	0.00
θ_6	Public-sector indicator variable	−0.33	0.03	−12.32	0.00
θ_7	High-risk indicator variable	−0.27	0.05	−4.98	0.00
θ_8	Systemic weight	−1.73	0.13	−13.03	0.00
Root mean *squared* error		15%			
Mean absolute error		11%			
Median absolute error		9%			
R^2		0.978			

entities in the portfolio exactly as in Eq. 5.32 in the default setting.[45] The second high-risk indicator is defined as

$$\mathbb{I}_{H_i} = \begin{cases} 1 : S_i \geq \text{PD14} \\ 0 : S_i < \text{PD14} \end{cases} . \tag{5.49}$$

Obligors at the very low end of the credit spectrum are, as one might expect, dominated by default risk. As a consequence, identifying them for different treatment improves the overall approximation. As before, these indicator variables are based on a certain logic, but their success likely depends on rather complicated non-linear aspects of the production t-threshold model.

The systematic weight parameter, the final addition to our approximating regression, also provides value in describing migration risk. Its role in determining the relative importance of systemic and idiosyncratic factors makes it important in both the migration and default settings. It is excluded in the full default model, because it is ultimately captured by other explanatory variables.

Figure 5.11 provides a heat-map of the cross correlation associated with our migration explanatory variables. The only potential danger point—in terms of multicollinearity—arises between the credit rating and public-sector variables. With this borderline exception, the remaining descriptive variables all appear to capture different systemic and idiosyncratic elements of migration-risk economic capital. This is underscored by the goodness-of-fit measures in Table 5.8. In addition to a fairly healthy increase in the R^2 measure, the proportional root-mean-square errors are reduced by a factor of three; the mean and median absolute errors, perhaps the most important dimensions, are cut in half. The results are broadly comparable—although marginally worse—than the full-default model outcomes presented in Table 5.5.

Figure 5.12 concludes our examination of the migration-approximation model through a detailed examination of the error profiles. In particular, it highlights the error plots (estimated less observed) for the base and full migration-risk approximation models. Each perspective considers the individual obligor and (normalized) size dimensions. A small number of model outliers are identified in the base model and inherited in the extended model graphics; in all cases, the additional explanatory variables improve the situation. Overall, a threefold reduction in the scale of error is provided with the extended model. Although it struggles in certain situations, and there is stubborn tendency to overestimate risk, the general performance of the approximation can be considered to be satisfactory.

[45] There are, it should be stressed, the only four of eight (non-intercept) variables arising in both default and migration models. This observation underscores the benefits of separate treatment of migration and default risk.

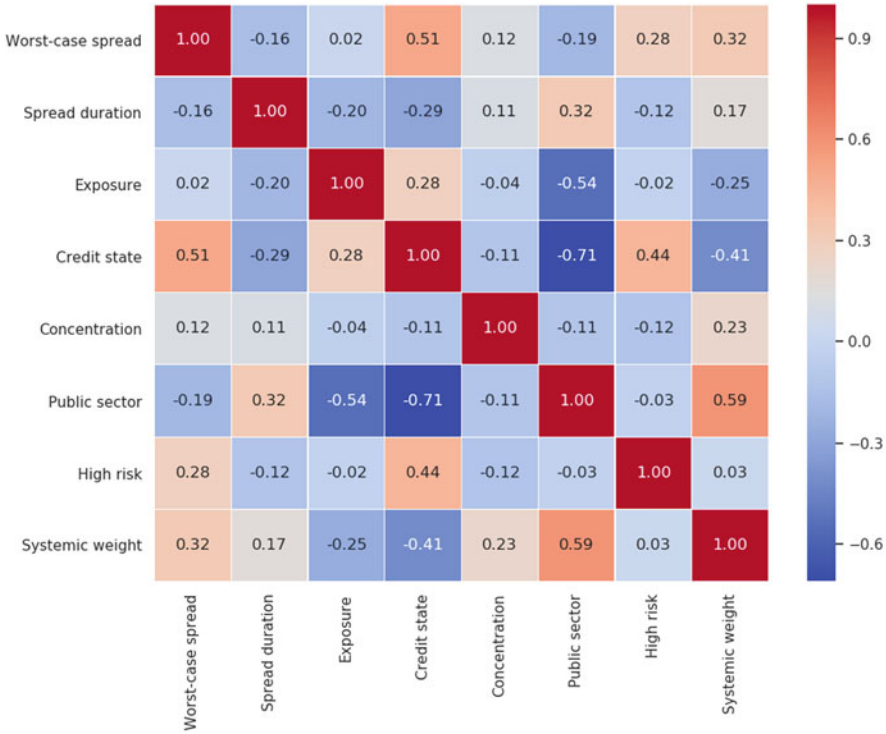

Fig. 5.11 *Migration instrument correlation*: This heat-map illustrates the cross-correlation between the key migration risk explanatory variables included in the extended regression model summarized in Table 5.8.

Colour and Commentary 60 (A WORKABLE MIGRATION APPROXIMATION MODEL): *Migration-risk economic capital is not easier to predict than its default equivalent.*[a] *Analogous to the default case, one can nonetheless construct a fairly defensible regression model involving meaningful systemic and idiosyncratic explanatory variables. The results are satisfactory; in particular, the* R^2 *approaches 0.98 and the median proportional approximation error comes out—not quite on par with the full default model—at around 9%. Migration risk approximation nevertheless remains a constant struggle. The model is re-estimated daily, parameters are stored in our database, and the results are carefully monitored via our internal diagnostic dashboards. A more formal model-selection procedure is performed on an annual basis.*

[a] To repeat, if it was, then we could perhaps dispense with our complicated simulation algorithms.

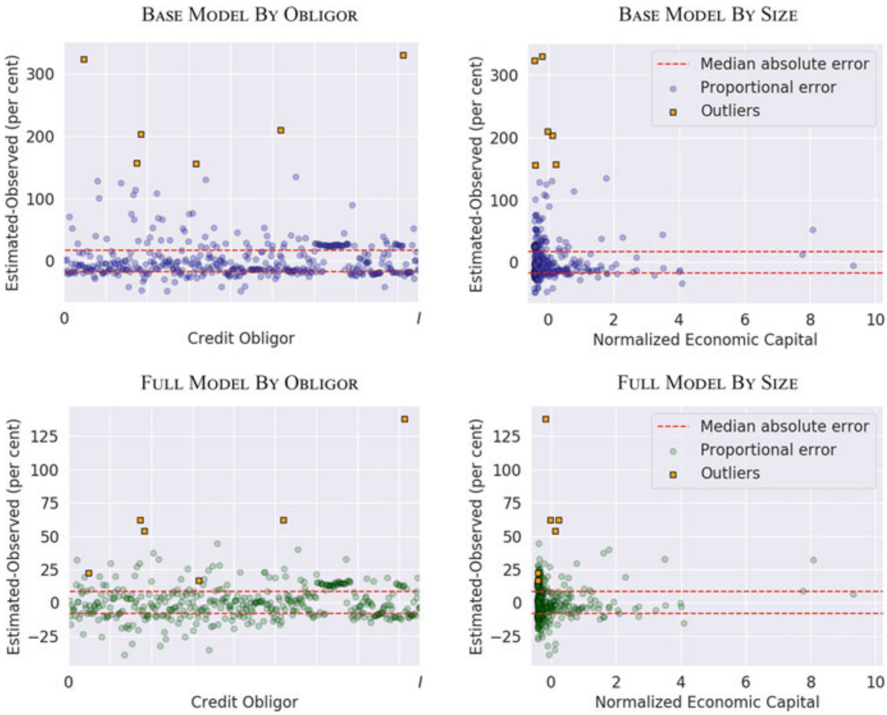

Fig. 5.12 *Migration-error analysis*: The preceding graphics highlight the error plots (estimated less observed) for the base and full migration-risk approximation models. Each perspective consider the individual obligor and (normalized) size dimensions. A threefold reduction in the scale of error is provided with the extended model.

5.4 Approximation Model Due Diligence

Lawyers talk about the idea of a standard of care and due diligence in a very technical and specific way. These ideas can, if we're careful with them, be meaningfully applied to the modelling world. We have, as model developers, a certain responsibility to exercise a reasonable amount of *prudence and caution* in model selection and implementation. Despite having invested a significant amount of time motivating and deriving our approximation models, we cannot yet be considered to have entirely performed our due diligence. The reasons are fairly simple: to this point, we have only examined their performance on an in-sample basis for a single point in time. To appropriately discharge our responsibility for due diligence, this section will address the *two* following questions:

1. How do our default and model approximations perform over a lengthier time interval?

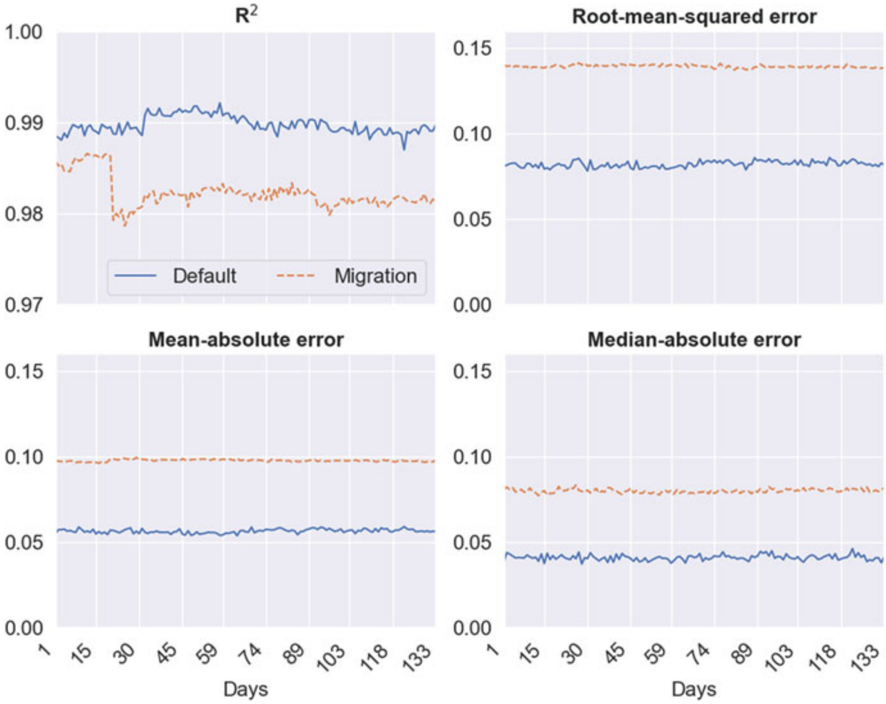

Fig. 5.13 *A historical perspective*: Using the portfolio data over each date in a six-month period straddling 2020 and 2021, the full default and migration models are estimated and the evolution of our four main goodness-of-fit measures are displayed. The results indicate a relatively high degree of stability over time.

2. What is the relative performance of our approximations—again over a reasonable length of time—along *both* the in- and out-of-sample dimensions?

Let's begin with the first, more easily addressed, issue. The default and migration approximation models were fit—using the ideas from the previous sections—to a sequence of portfolios. Figure 5.13 illustrates the results by focusing on the daily evolution of the four resulting goodness-of-fit criteria over a roughly 130 business-day (i.e., six-month) period. The results are fairly boring; the root-mean, mean-absolute, and median-absolute error estimates remain resolutely stable over the entire period. The default R^2 never deviates, in any meaningful way, from 0.99. The only slightly interesting result is a downward jump in the migration R^2 stemming from a hard-to-approximate addition to the portfolio earlier during our horizon of analysis.

Boring, in the context of Fig. 5.13, is a good thing. The results strongly suggest that our two approximations are satisfactorily robust. Caution is, of course, required. A few new observations could undermine these positive results and lead us to question our conclusion. This is important to keep in mind and argues for constant oversight. This point notwithstanding, we may cautiously infer, from the analysis in Fig. 5.13, that our approximations have value beyond the single datapoint presented in the previous sections.

The second question is harder to answer and takes us further afield. The natural starting point is to provide a bit more precision on the difference between in- and out-of-sample estimates. As the name suggests, in-sample fit describes errors stemming from observations used to fit (or train) one's approximation algorithm. Out-of-sample fit, by extension, applies to errors for observations falling outside, or excluded from, the estimation dataset. Any information embedded in out-of-sample observations is not, literally by construction, captured in one's model parameters. It is precisely this characteristic that makes out-of-sample prediction both difficult and interesting.

Out-of-sample prediction is a powerful technique to guard against the over-fitting of an estimation model. This is a concern shared by both statisticians and data scientists. Over-fitting basically describes an excessive specialization of the model parameters to one's specific data sample; the consequence is the identification of trends that do not generalize more broadly to other samples or the overall data population.[46] It is easy to see how out-of-sample error analysis is precisely the right tool to assess the degree of model over-fitting in one's approach.

Over-fitting is a particularly depressing and dangerous model shortcoming. One's in-sample goodness-of-fit measures look great, but the model actually fails dreadfully out of sample. Given the high levels of R^2 and low proportional errors in the preceding approximation models, it is entirely natural to worry about this problem. Have we, for example, over-specialized our model structure to the provided sample data? Addressing this entirely rationale fear is an important part of our due-diligence process. Moreover, given that the loan-pricing problem is, by definition, an out-of-sample prediction problem, we cannot ignore this dimension.

How should we proceed? Stone [26] and Geisser [11], almost 50 years ago, proposed a general technique to systemically manage in- and out-of-sample datasets and thereby control the over-fitting problem. The idea is quite simple: one first partitions the given dataset into k (roughly) equally sized subsets. Each of these subsets is referred to as a *fold*. One then excludes one of these folds—let's start with $k = 1$—and uses the remaining data to estimate the model. In statistical or machine learning, these are generally referred to as the training and testing datasets, respectively. One proceeds to compute two sets of goodness-of-fit measures: one for the estimation dataset and another for the set of excluded data. These are the in- and out-of-sample measures, respectively. One then puts the excluded data back into the sample, and excludes another fold (say, $k = 2$) and repeats the exercise. This is performed k times until each of the partitions has been excluded once. The result is

[46] A less charitable way to describe over-fitting is the characterization of trends that do *not* actually exist.

Table 5.9 *Goodness-of-fit statistics*: This table summarizes a range of goodness-of-fit summary statistics for both in- and out-of-sample estimates. The values are based on a five-fold cross-validation exercise performed for every business day over a six-month period.

		Default				Migration			
Measure	Perspective	Mean	σ	Min	Max	Mean	σ	Min	Max
R^2	In-sample	0.990	0.002	0.982	0.996	0.982	0.004	0.969	0.994
	Out-of-sample	**0.990**	**0.007**	**0.965**	**0.999**	**0.979**	**0.013**	**0.927**	**0.999**
Root-mean-squared error	In-sample	0.081	0.004	0.066	0.091	0.137	0.008	0.114	0.150
	Out-of-sample	**0.086**	**0.015**	**0.050**	**0.158**	**0.149**	**0.039**	**0.092**	**0.309**
Mean absolute error	In-sample	0.056	0.002	0.047	0.062	0.097	0.003	0.085	0.106
	Out-of-sample	**0.059**	**0.006**	**0.041**	**0.090**	**0.102**	**0.011**	**0.073**	**0.145**
Median absolute error	In-sample	0.041	0.002	0.035	0.047	0.080	0.003	0.069	0.092
	Out-of-sample	**0.042**	**0.005**	**0.028**	**0.059**	**0.081**	**0.008**	**0.056**	**0.105**

a broad range of both in- and out-of-sample estimates.[47] This is referred to—again following from Stone [26] and Geisser [11]—as k-fold cross validation.[48]

Table 5.9 illustrates the results of a k-fold cross validation exercise performed for every day during the same six-month time interval employed for the construction of Fig. 5.13. On each distinct date, the original dataset was randomly reshuffled to ensure different fold compositions over time. With roughly 130 days and $k = 5$ folds per dataset, the consequence is about $5 \times 130 = 650$ separate sets of in- and out-of-sample goodness-of-fit statistics. In Table 5.9 we find the mean, volatility (i.e., σ), minimum and maximum in- and out-of-sample good-of-fit measures for both migration and default approximation models.

Although there are admittedly many numbers to peruse, a few conclusions are possible. First of all, there is a universal degradation of fit when moving from the in- to the out-of-sample perspective. Second, and more interestingly, the magnitude of the deterioration of fit is (for the most part) relatively modest. Finally, the default model appears to be more robust, along this dimension, than their migration equivalents.

Figure 5.14 seeks to help us better digest all of the figures in Table 5.9 by illustrating the competing distributions of in- and out-of-sample good-of-fit measures for the default approximation model. Visually, the in-sample fit measures are clearly tighter than the out-of-sample values. The centre of the two distributions do not appear to be dramatically different, although there is a decided right skew.[49] In the default setting, the overall worsening of out-of-sample prediction accuracy appears to be reasonably limited.

[47] It is also possible to repeat this again by randomly reshuffling one's original dataset and repartitioning.

[48] See Hastie et al. [16, Section 7.2] for a helpful introduction to cross validation.

[49] Or, in the case of the R^2, a left skew.

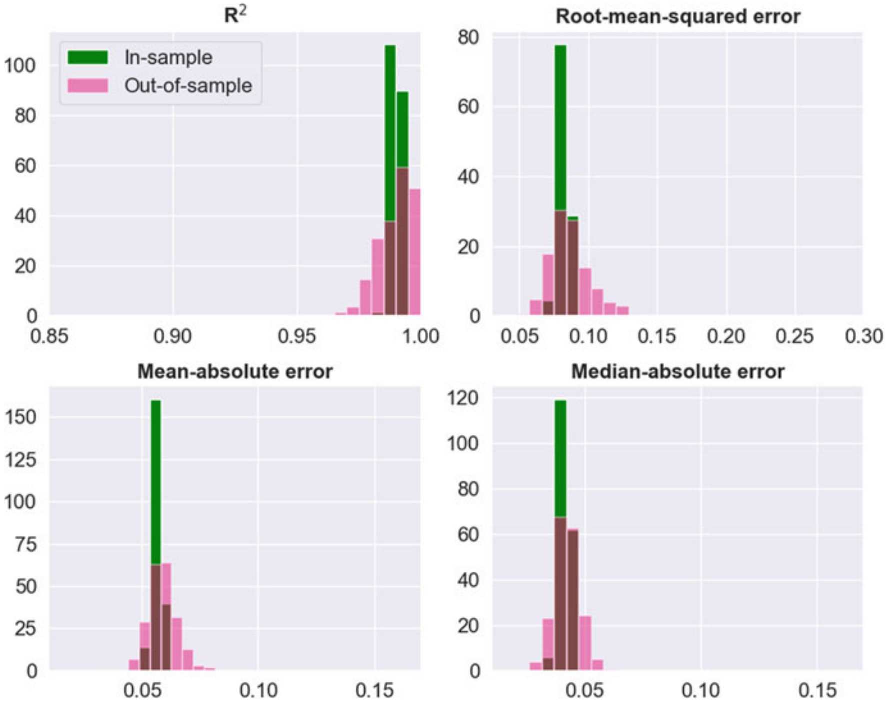

Fig. 5.14 *Default out-of-sample fit*: Computed using a five-fold cross validation of the full default approximation model for each day over a six-month period, the graphics display the distribution of in- and out-of-sample goodness-of-fit measures. Each day, the order of the observations using the cross validation were randomly shuffled.

Figure 5.15 repeats the analysis found in Fig. 5.14, but instead examines the *migration* approximation model. Once again, the in-sample fit measures are fairly closely clustered around a central point. The dispersion of the out-of-sample metrics is qualitatively similar, but substantially broader than in the default case. The root-mean-squared error, in particular, exhibits a bimodal distribution. This stems from the inclusion or exclusion of certain outlier observations from the training and testing samples. While the out-of-sample performance of the migration approximation is significantly worse than the default model, the analysis does not reveal any terrible surprises. We may, therefore, cautiously provide fairly positive responses to the two due-diligence questions posed at the beginning of this section.

Colour and Commentary 61 (DUE-DILIGENCE RESULTS)*: The notion of due-diligence basically sets a standard for the requisite degree of caution and prudence for a reasonable person finding themselves in a position of*

(continued)

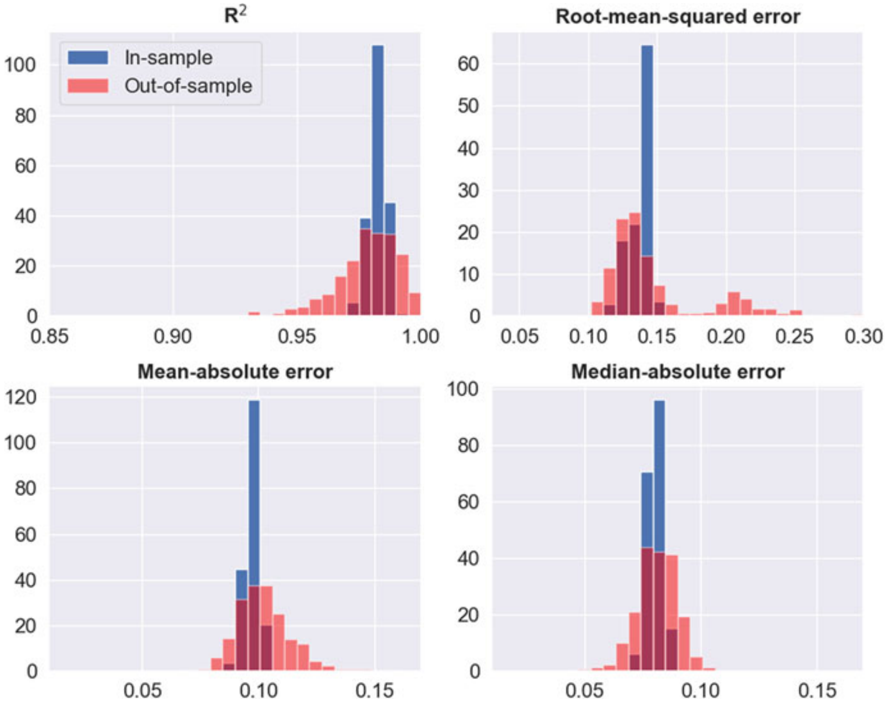

Fig. 5.15 *Migration out-of-sample fit*: Analogous to Fig. 5.14, the graphics apply to the migration case. Computed using a five-fold cross validation of the full migration approximation model for each day over a six-month period, the graphics display the empirical distribution of in- and out-of-sample goodness-of-fit measures. Each day, the order of the observations using the cross validation were randomly shuffled.

Colour and Commentary 61 (continued)
responsibility. Given its central importance for a range of important credit-risk economic capital applications—including, most centrally, loan pricing and stress testing—we identify a strong need for model due-diligence. This is operationalized in two *questions: model-prediction observations over a broader time period* and *out-of-sample performance. Examining daily portfolios over a six-month period reveals an admirable degree of robustness in our in-sample goodness-of-fit measures. Analyzing the same six-month data window within the context of a (daily) five-fold cross validation exercise further demonstrates only moderate degradation associated with out-of-sample goodness-of-fit in our approximation models.*[a] *This extensive due-diligence analysis provides a degree of comfort in our approximation models.*

[a] The default model nonetheless significantly outperforms the migration approximator.

5.5 The Full Picture

The previous discussion hopefully underscores that, despite rather strenuous analytical efforts and a small amount of creativity, it is *not* particularly easy to construct a fast, accurate, and defensible semi-analytic approximation of the default and migration economic-capital estimates. If it was, of course, then we might ultimately be able to avoid lengthy and computationally expensive simulation algorithms. The challenges certainly stem from the many complex non-linear relationships between the individual obligors and their characteristics. Despite these complications, we are able to offer *two* fairly defensible approximation approaches.

It should be fairly obvious, by this point, how total credit-risk economic capital is approximated. Nevertheless, for completeness, the full approximator is written as

$$\begin{aligned} \mathcal{E}_i(\alpha^*) \approx \widehat{\mathcal{E}_i(\alpha_z^*)} &= \underbrace{\sum_{k=0}^{\kappa} \hat{\xi}_k X_{i,k}^{(d)}}_{\substack{\text{Default}\\ \text{risk}}} + \underbrace{\sum_{k=0}^{\kappa} \hat{\theta}_k X_{i,k}^{(m)}}_{\substack{\text{Migration}\\ \text{risk}}}, \\ &= \widehat{\mathcal{E}_i^{(d)}(\alpha_z^*)} + \widehat{\mathcal{E}_i^{(m)}(\alpha_z^*)}, \end{aligned} \tag{5.50}$$

for $i = 1, \ldots, I$. This requires two main inputs: the regression coefficients and the explanatory variables. For various downstream applications the parameters—estimated and stored daily—are readily available. Separate functions, involving a modicum of numerical integration, are available to collect the descriptive variable matrices (i.e., $X^{(d)}$ and $X^{(m)}$) as required.

Both approximations rely upon structurally motivated linear-regression estimators. This offers an important advantage: it allows us to exploit our knowledge of the underlying problem. Other choices are nonetheless possible. The presence of complex non-linear relationships strongly suggests consideration of other, perhaps more empirically motivated, techniques. The area of machine-learning, for example, offers a number of potential methods to address these features of our model. Internal assessment of popular machine-learning approaches has, in the current implementation, nevertheless been somewhat disheartening. While generally reasonable, they fail to outperform the previously presented techniques. One reason appears to relate to relatively small datasets; this is, to a certain extent, a constraint associated with the need to condition on the current portfolio structure. Another explanation is that our proposed approximators are actually rather sensible and correspondingly difficult to beat.

Colour and Commentary 62 (ECONOMIC-CAPITAL APPROXIMATION CHOICES)*: Building a defensible and workable mathematical approximation of a complicated object is part art and part science. Numerous choices are required. The current economic-capital approximation model is no exception. At least* three *explicit strategies have been followed. First of all, separate approximation models are constructed for both default and migration effects. This enhances analytical flexibility, but roughly doubles the total amount of effort. Structurally motivated linear models, as a second point, have been employed. These high-bias (hopefully low-variance) estimators permit exploitation of model knowledge, but may ignore important non-linear elements.*[a] *Finally, and perhaps most importantly, model selection was verified over an extended time period with the consideration of both in- and out-of-sample perspectives. The combination of these three elements does not guarantee optimal approximation models—far from it—but it does represent a comprehensive framework for their defence and ultimate improvement.*

[a] There is, for example, no reason one might not be able to exploit model knowledge in the context of more flexible estimators.

5.5.1 A Word on Implementation

The actual implementation of the migration and default approximation models is, after all the detailed discussion of their construction and defensibility, rather anticlimactic. Figure 5.16 summarizes the various algorithmic steps. Conceptually, the code follows the same *three* steps for each model: compute and collect the instrument (or explanatory) variables, calculate the dependent variables, and then estimate the regression coefficients. All these activities are assigned to separate functions since each requires a reasonable amount of effort. `getXDefault` and `getXMigration`, which manage the computation of response variables, perform a moderate number of numerical integrations. Moreover, the logarithmic transformations require a bit of caution.

The primary output of the approximation algorithm takes the form of two Python dictionaries: `vd` and `vm`. These two objects are populated with the default and migration economic-capital estimates, respectively. Although they usefully contain model forecasts, residuals, standard errors, test statistics and p values, the most important data members are the parameters.[50] These are used extensively in

[50] In addition, all of these diagnostics are computed in both logarithmic and currency space. Although our principal focus is the currency estimates, there is useful information in both perspectives.

```
begin approximateEc
    getXDefault: this function collects the default instrument variables            DEFAULT
        and places the result into Xd∈ ℝ^{n×κ}                                      LOGIC
    getY: assigns the dependent variable to Yd∈ ℝ^{n×2}; we need two
        dimensions to represent logarithmic and transformed currency space
    estimateModel: solves multivariate regression Yd = β_d Xd + ε_d and places
        parameters and goodness-of-fit results into vd data object

    getXMigration: this function collects the migration response variables          MIGRATION
        and places the result into Xm∈ ℝ^{n×κ}                                      LOGIC
    getY: assigns the dependent variable to Ym∈ ℝ^{n×2}; the same
        function is employed with a different input flag
    estimateModel: solves multivariate regression Ym = β_m Xm + ε_m and places
        parameters and goodness-of-fit results into vm data object; the same
        estimation logic (and function) used for both default and migration approximation

    Clean-up: Write the regression coefficients to the appropriate database
end approximateEc
```

Fig. 5.16 *The EC approximation*: Using instrument variables derived from the structural elements of the default and migration computations, economic-capital is approximated using a standard OLS framework. Since a logarithmic transformation is employed and it is useful to separately predict their outcomes, the default and migration approximations are individually performed.

subsequent applications and, upon the completion of `approximateEc`, are saved into their own database table.

Having computed the approximation models, we can immediately put them to work. More complex approximation-related applications are considered in future chapters, but a small question needs to be addressed. The first application of our approximation methodologies, touches on the relationship between credit obligors and individual loans (or treasury) positions.

5.5.2 An Immediate Application

Credit-risk economic-capital computations operate, as previously discussed, at the risk owner level. A risk owner, to be clear, is an entity whose credit deterioration, or failure to pay, would lead to financial loss. It may be a single corporation or public institution or, in some cases, it can cover a parent company and its individual subsidiaries. While these details are important, for the purposes of this discussion, one simply needs to understand the relationship between risk-owner and individual exposures. Finding the correct terminology for this latter quantity is a challenge; we might refer to it as a trade, a transaction, an instrument, or a position. None of these is entirely satisfactory. However you wish to refer to it, it is the most atomistic element of one's portfolio. In the lending book, these are essentially loans. On the treasury side, however, this can involve a broad range of positions in various

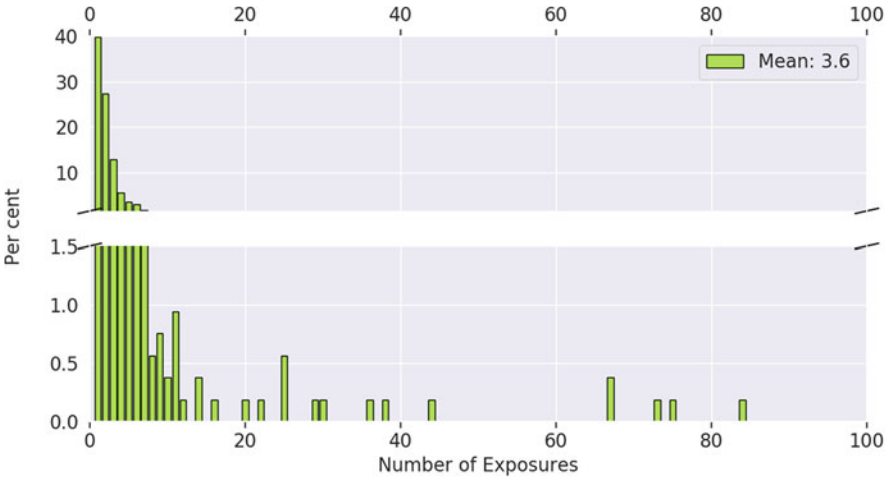

Fig. 5.17 *Exposure counts*: Each individual risk owner has, on average, four separate individual exposures. The preceding graphic illustrates the significant dispersion around this central value. Two thirds of the portfolio's risk owners have only one or two exposures, whereas others have dramatically more.

instruments such as swaps, bonds, forward contracts, repos, or deposits. At the NIB, for lack of a better term, we use *individual exposure* (or just exposure) to refer to these elements of our portfolio.[51]

To actually compute economic capital, these individual exposures are aggregated to the risk-owner level. In recent years, our portfolio has had in the neighbourhood of 600 individual risk owners, but more than 2000 exposures implying roughly four exposures per risk owner.[52] In its base implementation, however, economic-capital is only assigned at the risk-owner level; it is, after all, at this level where default (or credit transition) occurs. This creates a small challenge. For analytic purposes, reporting, and loan-pricing activities, it is necessary to allocate the risk-owner economic capital to each of its underlying exposures.

How are these exposures distributed by risk owner? Figure 5.17 answers this question by displaying, for the usual arbitrary date in 2020, the range of exposure numbers across individual risk owners. These results are presented as a percentage of the total number of exposures. Although the mean outcome is roughly four, it varies dramatically. About two thirds of the portfolio's risk owner have only one or two exposures, whereas some have dozens. The largest number of exposures per risk owner approaches 100. These are certainly associated with treasury counterparties

[51] The reader should, of course, feel free to use (or invent) a better descriptor.

[52] Other financial institutions may have different distributions, but face the same conceptual issue.

and include a range of instruments such as swaps, deposits, and repo exposures.[53] Given the presence of netting and collateral arrangements, the sheer number of exposures, in isolation, does not tell a very useful story.[54]

Since the model does not provide any explicit guidance as to how we should allocate the overall risk-owner economic-capital to individual exposures, we need to address this ourselves. After some deliberation, a few general allocation principles were identified to guide us. More specifically, there are *three* key elements, in our view, that are required:

1. **Conservation of Mass**: Economic-capital can be neither created, nor destroyed, by the allocation procedure.
2. **Risk Proportionality**: Riskier exposures, all else equal, should receive a larger proportion of the risk allocation than their safer equivalents.
3. **Weak Positivity**: A lower bound on the exposure-level risk allocation is zero; we do *not* allocate negative economic capital to an exposure. This might appear somewhat controversial. Negative economic capital, after all, is frequently assigned in a market-risk setting. Such positions, or risk-factor exposures, can broadly be interpreted as a form of hedge. Outside of default insurance, which we have not traditionally employed, such dynamics do not appear to be sensible in the credit-risk setting.

We will employ these basic axioms to direct our choice of allocation methodology.

Some notation is now required. Let us represent risk-owner economic capital, once again, as $\mathcal{E}_i$ for $i = 1, \ldots, I$. Here we refer to the total risk-owner economic-capital allocation including both default and credit-migration effects. Our task is to attribute $\mathcal{E}_i$ to each of its $\mathrm{N}_i \geq 1$ exposures. Clearly, if $\mathrm{N}_i = 1$, this is a trivial task. If $\mathrm{N}_i = 10$, however, it will be rather less obvious how to apportion the total economic capital.

Whatever approach one selects, it is basically a weighting problem. Practically, therefore, any allocation algorithm can be described as

$$\mathcal{E}_{ij} = \omega_{ij}\mathcal{E}_i, \tag{5.51}$$

where $\mathcal{E}_{ij}$ is the jth exposure's allocation of the ith risk-owner's economic capital. ω_{ij} is the weight of the jth exposure to the ith risk owner. These weights need to be identified. In other words, the entire exercise is about finding sensible ω's.

The good news is that we may now employ our previously described principles to identify reasonable ω choices. The conversation-of-mass axiom, for example, can be translated into a requirement that the full risk-owner level economic capital (i.e.,

[53] In some large institutions, there may be thousands of exposures with a single risk-owner.

[54] This question is addressed in detail in Chap. 10.

$\mathcal{E}_i$) is assigned to the exposures. Mathematically, this amounts to

$$\begin{aligned}\mathcal{E}_i &= \sum_{j=1}^{\mathrm{N}_i} \mathcal{E}_{ij}, \\ &= \sum_{j=1}^{\mathrm{N}_i} \underbrace{\omega_{ij}\mathcal{E}_i}_{\substack{\text{Equation}\\ 5.51}}, \\ &= \mathcal{E}_i \underbrace{\sum_{j=1}^{\mathrm{N}_i} \omega_{ij}}_{\substack{\text{Our}\\ \text{choice}}},\end{aligned} \tag{5.52}$$

which implies that we merely require that

$$\sum_{j=1}^{\mathrm{N}_i} \omega_{ij} = 1, \tag{5.53}$$

for $i = 1, \ldots, I$. Quite simply, the condition in Eq. 5.53 ensures that the total economic-capital amount is assigned.

If we define $\mathcal{R}_{ij}$ as the riskiness of the jth exposure associated with the ith risk-owner, then we also require that:

$$\mathcal{R}_{ij} \geq \mathcal{R}_{ik} \Rightarrow \omega_{ij} \geq \omega_{ik}. \tag{5.54}$$

for all $j, k = 1, .., \mathrm{N}_i$. A riskier exposure receives a greater economic-capital allocation; this addresses the risk-proportionality principle. Our final axiom, which deals with weak positivity, implies that the individual weights must respect:

$$\omega_{ij} \in [0, 1], \tag{5.55}$$

for all $j = 1, \ldots, \mathrm{N}_i$ and $i = 1, \ldots, I$.

The three conditions in Eqs. 5.53 to 5.55 provide us with all that we require to write a concrete expression for our individual weights. The weights must be positive, sum to unity, and respect our risk ordering. The following expression is consistent with all of these constraints:

$$\omega_{ij} = \frac{\max\left(\mathcal{R}_{ij}, 0\right)}{\sum_{k=1}^{N_i} \max\left(\mathcal{R}_{ik}, 0\right)} \tag{5.56}$$

The only element that is missing is a choice of risk assessment. There exist a number of possibilities—such as the total exposure or market value—but there is already an obvious candidate. In the previous sections, a detailed description of the default and migration economic-capital approximations was provided. These would appear to be a highly sensible choice for each individual $\mathcal{R}_{ij}$ assessment. More specifically, using Eq. 5.50, we require that

$$\begin{aligned}\mathcal{R}_{ij} &= \overbrace{\underbrace{\sum_{k=0}^{\kappa} \hat{\xi}_k X_{ij,k}^{(d)}}_{\substack{\text{Default}\\ \text{risk}}} + \underbrace{\sum_{k=0}^{\kappa} \hat{\theta}_k X_{ij,k}^{(m)}}_{\substack{\text{Migration}\\ \text{risk}}}}^{\widehat{\mathcal{E}_{ij}(\alpha_z^*)}}, \\ &= \widehat{\mathcal{E}_{ij}^{(d)}(\alpha_z^*)} + \widehat{\mathcal{E}_{ij}^{(m)}(\alpha_z^*)}.\end{aligned} \tag{5.57}$$

In short, we first compute the necessary instrument variables for each ith risk owner and $j = 1, \ldots, N_i$ exposures. We then use our regression coefficients to produce associated migration and default economic-capital estimates; this serves admirably as a consistent measure of position risk for our allocation algorithm.[55]

Colour and Commentary 63 (Exposure-Level Allocation)*: Since default and migration occur at the risk-owner level, they must be computed and allocated from a risk-owner perspective. Economic-capital allocations—for a variety of practical applications—are nonetheless required at the*

(continued)

[55] Our estimation algorithms are indifferent as to whether they are performed at the risk-obligor or instrument level. One need only be able to compute the necessary explanatory variables. We can also, should we so desire, separately allocate default and migration economic capital.

Colour and Commentary 63 (continued)
individual exposure level.[a] *A downgrade or a default will impact all of a risk-owner's underlying exposures, albeit in potentially different ways.*[b] *In other words, not all exposures possess the same risk profile and, as a consequence, assignment of risk-owner economic-capital is not an uninteresting and routine task. Identifying three commonsensical principles governing this operation and, leaning heavily on our approximation models, an internally consistent methodology is presented. We can see that our extensive effort to construct a reasonable economic-capital estimator is already bearing fruit.*

[a] Individual exposure is a potentially confusing term. As a trade, transaction, position, or instrument, it refers to the most atomistic elements of our portfolio.

[b] Different loans, for example, may exhibit differences in guarantees or collateral thus influencing loss-given-default values.

5.6 Wrapping Up

After having gone to such lengths to construct a functioning simulation-based engine for our credit-risk economic-capital model in Chap. 4, it is somewhat peculiar to immediately allocate an entire chapter towards finding semi-closed-form analytic approximations of the same object. The driving reason is speed. Even with clever use of parallel-processing and variance-reduction techniques, we cannot produce credit-risk economic capital estimates in seconds.[56] Certain applications, however, demand rapidity. On one hand, it is necessary to perform timely and flexible loan-pricing computations. Loan origination, as a fundamental activity, simply cannot operate without the ability to quickly consider a multitude of lending options. On the other hand, some endeavours (such as stress-testing) involve a huge scale of calculation that cannot be handled at simulation-model speeds. Whatever the reason, there is a clear demand for a reasonably accurate, instrument-level, fast, semi-analytical approximation of both default and migration economic capital. Their construction is tedious, time-consuming, and we are ultimately never entirely satisfied with their performance. The effort is nonetheless worthwhile. Our investment in this chapter—which has already found one small application—will pay generous dividends as we turn to various economic-capital applications in the forthcoming discussion and analysis.

[56] Even doing so in a small number of minutes is, for a reasonable price, almost impossible.

References

1. M. Bellalah, M. Zouari, A. Sahli, and H. Miniaoui. Concentration measures in risk management. *International Journal of Business*, 20(2):110–127, 2015.
2. BIS. The internal ratings-based approach. Technical report, Bank for International Settlements, January 2001.
3. BIS. International convergence of capital measurement and capital standards: A revised framework. Technical report, Bank for International Settlements, January 2004.
4. BIS. An explanatory note on the Basel II IRB risk weight functions. Technical report, Bank for International Settlements, July 2005.
5. D. J. Bolder. *Credit-Risk Modelling: Theoretical Foundations, Diagnostic Tools, Practical Examples, and Numerical Recipes in Python*. Springer-Nature, Heidelberg, Germany, 2018.
6. D. J. Bolder. Statistics and Machine-Learning: Variations on a Theme. In P. Nymand-Andersen, editor, *Data Science in Economics and Finance for Decision Makers*. Risk Books, 2021.
7. D. J. Bolder and T. Rubin. Optimization in a Simulation Setting: Use of Function Approximation in Debt-Strategy Analysis. Bank of Canada: Working Paper 2007-13, February 2007.
8. L. Breiman. Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3):199–231, 2001.
9. S. Emmer and D. Tasche. Calculating credit risk capital charges with the one-factor model. Technical report, Deutsche Bundesbank, January 2005.
10. S. Figini and P. Uberti. Concentration measures in risk management. *The Journal of the Operational Research Society*, 64(5):718–723, May 2013.
11. S. Geisser. The predictive sample reuse method with applications. *Journal of the American Statistical Association*, 70(350):320–328, 1975.
12. G. H. Golub and C. F. V. Loan. *Matrix Computations*. The John Hopkins University Press, Baltimore, Maryland, 2012.
13. M. B. Gordy. A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12:199–232, 2003.
14. M. B. Gordy and E. Lütkebohmert. Granularity Adjustment for Basel II. Deutsche Bundesbank, Banking and Financial Studies, No 01-2007, January 2007.
15. M. B. Gordy and E. Lütkebohmert. Granularity adjustment for regulatory capital assessment. University of Freiburg, January 2013.
16. T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer-Verlag, New York, first edition, 2001.
17. G. James, D. Witten, T. Hastie, and R. Tibshirani. *An Introduction to Statistical Learning*. Springer-Verlag, New York, first edition, 2013.
18. I. T. Jollliffe. *Principal Component Analysis*. Springer, New York, New York, 2002.
19. G. G. Judge, W. Griffiths, R. C. Hill, H. Lütkepohl, and T.-C. Lee. *The Theory and Practice of Econometrics*. John Wiley and Sons, New York, New York, second edition, 1985.
20. E. Lütkebohmert. *Concentration Risk in Credit Portfolios*. Springer Verlag, Berlin, 2008.
21. R. J. Martin and T. Wilde. Unsystematic credit risk. *Risk*, pages 123–128, November 2002.
22. B. Milanovic. The Gini-type functions: An alternative derivation. *Bulletin of Economic Research*, 46(1):81–89, 1994.
23. M. Pykhtin. Multi-factor adjustment. *Risk*, (3):85–90, March 2004.
24. A. Ralston and P. Rabinowitz. *A First Course in Numerical Analysis*. Dover Publications, Mineola, New York, second edition, 1978.
25. T. Ribarits, A. Clement, H. Seppälä, H. Bai, and S.-H. Poon. Economic Capital Modeling: Closed form approximation for real-time applications. European Investment Bank and Manchester University, June 2014.
26. M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):111–147, 1974.

27. B. Torell. Name concentration risk and pillar 2 compliance: The granularity adjustment. Technical report, Royal Institute of Technology, Stockholm, Sweden, January 2013.
28. O. A. Vasicek. Limiting loan loss probability distribution. KMV Corporation, August 1991.
29. S. Yitzhaki and I. Olkin. Concentration indices and concentration cures. *Stochastic Orders and Decision under Risk*, IMS Lecture Notes – Monograph Series 1991.