



Economic and regulatory capital allocation for revolving retail exposures [☆]

Roberto Perli ^{a,*}, William I. Nayda ^b

^a Board of Governors of the Federal Reserve System, Mail Stop 75, Washington, DC 20551, USA

^b Capital One Financial Corp., McLean, VA, USA

Abstract

We present two internal capital allocation models and compare the capital ratios they generate with those prescribed by the latest revision of Basel's New Capital Accord proposal for advanced retail portfolios, which allows for explicit future margin income recognition. Given a test portfolio of credit card exposures that we assemble, we find that Basel's ratios are closer to those generated by our models for low credit risk segments. We attribute the discrepancies to the different ways Basel and our models account for future margin income, to Basel's assumptions about asset correlations and to one model taking macroeconomic conditions into account.

© 2003 Elsevier B.V. All rights reserved.

JEL classification: G0; G2

Keywords: Capital allocation; Credit risk models; Revolving retail exposures; Future margin income

[☆] The opinions expressed here are those of the authors and not those of the Board of Governors of the Federal Reserve System or those of Capital One Fin. Corp. The models discussed in the paper are illustrative in nature and should not be taken as representative of the actual models used by Capital One Fin. Corp. to allocate capital across portfolio segments. Likewise, the data used here should not be taken as representative of Capital One Fin. Corp.'s actual portfolio. We would like to thank Julie Adiletta, Paul Calem, Brad Case, Mark Carey, Michael Gordy, Bill Hood, David Jones, Loretta Mester, Marc Seidenberg, Jeff Stokes, Lance Whitaker, an anonymous referee, and participants at the Conference on Retail Credit Risk Management and Measurement organized by the Federal Reserve Bank of Philadelphia for helpful discussions and comments.

* Corresponding author. Tel.: +1-202-452-2465; fax: +1-202-452-2301.

E-mail addresses: roberto.perli@frb.gov (R. Perli), bill.nayda@capitalone.com (W.I. Nayda).

1. Introduction

Many large retail financial institutions have developed models to assess credit risk and to allocate their economic capital to different segments of their portfolios.¹ A strong incentive to develop these models was provided in part by the release by the Basel Committee on Banking Supervision (BCBS henceforth) in 1999 of a consultative paper on a New Basel Capital Accord (NBCA, see BCBS, 1999) that will eventually replace the one currently in force, and of its subsequent revisions and proposed implementation details (BCBS, 2001, 2003). It is interesting to compare the capital allocations resulting from models that could conceivably be used internally by banks to those resulting from the application of the NBCA proposal. We believe that this exercise is particularly useful for retail portfolios, which have received little attention in the literature compared to, say, commercial portfolios, and which are allowed, under the latest (as of this writing) revision of the NBCA, to subtract a proxy for future margin income from the credit loss-based capital ratios. We present two possible capital allocation models – one relatively simple, the other more complex – that try to capture some key features of retail lending, with an emphasis on the relationship between future margin income and credit losses. We assemble and segment a mini portfolio of revolving exposures consisting of test credit card accounts that span a wide range of the credit spectrum and calculate capital ratios for each segment according to the models and to the new Basel formula. Our conclusion is that the current version of the proposal produces capital allocations that are close to those generated by the models for low-risk segments, while the discrepancies can be substantial for higher-risk segments. We identify several factors that could account for the differences, including the way the NBCA approximates future margin income, its assumptions about asset correlations, and the fact that one of our models explicitly takes macroeconomic conditions into account, while the NBCA does not.

The new accord is aimed at correcting some problems with the existing accord – above all, the less than perfect consideration given to credit risk in the determination of capital ratios and the incentive for institutions to engage in so-called “regulatory arbitrage”, i.e., the selection of exposures as a function of regulatory capital ratios. To avoid these problems the new accord proposes three alternative regimes: the “standardized” (similar to the current accord but with more risk differentiation), the “foundation IRB” (for internal ratings based) and the “advanced IRB”. The IRB approaches will have banks segment the portfolio according to their own criteria (although guidelines on segmentation are provided) and then apply a given formula to determine the capital ratio for each given segment. Banks that opt for the foundation regime would have to provide only one input to the formula, namely, the probability of default for loans in each segment; banks that opt for the advanced regime will have to also provide estimates of loss given default (LGD), exposure at default, and maturity. In this paper we focus on the advanced IRB approach, as no

¹ See Treacy and Carey (2000) for a survey of the use of internal credit rating systems at the largest US banks.

foundation approach is allowed for retail institutions, and the standardized approach does not differentiate nearly enough with respect to risk to be comparable with internal models.²

Revolving retail portfolios are very different from commercial portfolios, in terms of number of accounts, exposures at default, LGD, and probabilities of default. When pricing revolving retail loans, banks take future portfolio losses into account, so accounts that do not charge off are supposed to pay for themselves and for those that default as well (and to make a profit on top of that). Given the high granularity of retail portfolios, this future income can be counted on as it does not depend on a small number of accounts. Therefore, future margin income is there to cover credit losses before a bank has to use its capital. For this reason the BCBS now allows 75% of the expected loss to be subtracted from the capital allocation resulting from the original formula because, since the loss is expected, it would have been priced for from the beginning.³

Recognizing that segments with a high probability of default might not be able to generate all the expected future margin income if too many accounts default, the BCBS introduced a provision that states that an institution will be allowed to subtract future margin income from capital only for those segments that it can prove having historically produced future margin income in excess of its expected losses plus two standard deviations of the annualized loss rate. If this provision was not there, the new proposal would say that the higher the expected loss of a segment, the higher the future margin income it will be able to generate. While it might be true that segments with high probabilities of default might be able to *assess* very high interest rates and fees, they might only be able to *collect* a small fraction of them, as accounts that default typically do so not just on their principal but also on their contractually assessed interest and fees.⁴

While there exist a few modelling frameworks for commercial loans that are relatively standardized (models such as CreditMetrics, KMV, etc.), the same cannot be said for the retail sector. Some lenders use versions of those commercial models modified to suit their specific needs. A better approach, however, either currently used by or in the works for many institutions, is to develop a framework that takes into account the peculiar features of retail lending. We present two such possible modelling frameworks in Sections 2 and 3 (a one-factor and a multi-factor model, respectively), and we compare the capital allocations they generate to those produced by the new Basel formula.

² The standardized retail approach states that all exposures should have a 6% capital ratio (or 75% risk weight), except for those more than 90 days past due, which should have a 12% ratio, and those included in BB-rated securitization tranches, which should have a 28% capital ratio.

³ The first circulated version of Basel's formula for revolving retail allowed for 100% of expected losses to be subtracted from the capital allocation, and a subsequent one allowed 90%.

⁴ A previously circulated version of this paper made precisely this point and concluded that the version of the NBCA that was circulated at the time, which did not include this provision, grossly underestimated the capital requirements of high-credit-risk segments.

Our results indicate that Basel's ratios for low risk segments are very close to our one-factor models ratios, especially when computed assuming a constant LGD across segments. Interestingly, when we relax that assumption and recognize that different segments might have different LGDs, the relationship between probability of default and capital ratios is not monotonic for our data set, both according to Basel and according to our model. On one hand, account groups with low probability of default in our data set have an LGD that is higher than that for accounts at the other end of the spectrum. On the other hand, the pricing of products offered to groups of accounts with lower probability of default is different from the pricing of products offered to accounts with higher probability of default and typically generates less revenue as a percentage of the outstanding balances. The combination of these two facts leads to capital ratios for some higher-risk segments that can be lower than those for some lower-risk ones.

We find it also interesting to explore how capital ratios for different risk groups are affected by macroeconomic conditions. To this end we determine the capital ratios for our segments using a multi-factor model where the factors are some explicitly identified macroeconomic variables. Although our framework can be interpreted as also including risks other than credit risk, such as regulatory and legislative risk, we conclude that groups with low credit risk respond less to macroeconomic conditions than groups with high credit risk. As a consequence the capital ratios for the latter groups are much higher than those indicated by either Basel or the one-factor model.

The remainder of the paper is organized as follows. In Section 2 we present our one-factor model, which is based on the same loss distribution as Basel, but accounts for future margin income in a different and, we believe, more realistic way; we calibrate it according to our data set, compute the capital ratios it implies, and compare them to Basel's. In Section 3 we do the same for our multi-factor model. Section 4 concludes.

2. A one-factor credit risk model with future margin income

In the corporate exposure literature it has become standard to assume that a firm will default on its debt when the value of its assets falls below a certain threshold at or before time T ; it is often assumed that this threshold is the value of the firm's liabilities (KMV, 1993; RiskMetrics Group, 1997).⁵ Here we assume that there are N consumers and that the default of a consumer occurs under the same circumstances, i.e., when the value of his/her assets, denoted by $V_i(T)$, falls below a certain threshold K_i . The typical assumption is that the value of a consumer's asset at the desired horizon is standard-normally distributed, i.e., $V_i(T) \sim \Phi(0, 1)$, and that the value at the present time is zero, i.e., $V_i(0) = 0$. We also could assume that the V_i are correlated across consumers according to a certain correlation matrix Σ . The threshold below

⁵ In practice consumers might default not just because of a decline in the value of their assets but also because of cash-flow problems that might be temporary in nature.

which a consumer defaults would then be related to the probability of default p_i of that consumer: $K_i = \Phi^{-1}(p_i)$.

The model as outlined above can easily be implemented only if the number of obligors N is small, as it requires the specification of N probabilities of default and $N(N-1)/2$ correlations. For a typical retail portfolio where N can be equal to several million, this specification is clearly impractical and further assumptions are required to make even this simple version of the model tractable; here we follow the analysis in Schönbucher (2000) and Vasicek (1987). We start by assuming that the value of all consumers' assets is driven by a single common factor Y and an idiosyncratic noise component ϵ_i :

$$V_i(T) = \sqrt{\rho}Y + \sqrt{1-\rho} \epsilon_i, \quad (1)$$

where $Y, \epsilon_i \sim N(0,1)$ and i.i.d. Here ρ represents the common correlation coefficient among all consumers' assets. Note that according to (1), given a realization y of the common factor Y , the asset values and the defaults are independent.

To further simplify the model, we assume that all consumers within a risk segment have the same probability of default p and, therefore, the same default threshold K . This assumption is a reasonable approximation if applied to a sufficiently homogeneous segment of the overall portfolio. We also assume that the exposure is the same for all consumers, and we set it equal to B/N .

We are interested in determining the probability that n out of the N total consumers will default. In the case of independence across consumers, i.e., if $\rho = 0$, the probability of n defaults is given by the binomial probability function

$$f(n) = \binom{N}{n} p^n (1-p)^{N-n}. \quad (2)$$

In the general case of non-zero correlation, the probability of n defaults has to be computed by averaging over all possible realizations of Y :⁶

$$f(n) = \int_{-\infty}^{+\infty} f(n|Y=y) \phi(y) dy, \quad (3)$$

where the conditional probability of n defaults given a realization y of Y is again given by the binomial distribution

$$f(n|Y=y) = \binom{N}{n} (p(y))^n (1-p(y))^{N-n}. \quad (4)$$

⁶ See the above mentioned papers by Schönbucher (2000) and Vasicek (1987) for the details. This model is also very similar to the RiskMetrics framework (RiskMetrics Group, 1997).

Finally, the probability of default conditional on a realization y is, using Eq. (1),

$$p(y) = \Pr(V_i(T) < K_i | Y = y) = \Pr\left(\epsilon_i < \frac{K_i - \sqrt{\rho}Y}{\sqrt{1-\rho}} \mid Y = y\right) = \Phi\left(\frac{K_i - \sqrt{\rho}y}{\sqrt{1-\rho}}\right). \tag{5}$$

Substituting Eqs. (4) and (5) into (3), we obtain the probability of n defaults:

$$f(n) = \int_{-\infty}^{+\infty} \binom{N}{n} \left(\Phi\left(\frac{K - \sqrt{\rho}Y}{\sqrt{1-\rho}}\right)\right)^n \cdot \left(1 - \Phi\left(\frac{K - \sqrt{\rho}Y}{\sqrt{1-\rho}}\right)\right)^{N-n} \phi(y) dy. \tag{6}$$

Eq. (6) can be solved numerically, or the model can be simulated a large number of times to determine a probability function for losses that will converge to (6).

Eq. (6) is valid for any number of exposures N . If N is very large, as it typically is in the case of retail portfolios, a further simplification of the model is possible, as shown by Schönbucher (2000). Since, conditional on the realization of y , defaults happen independently of each other, as N tends to infinity the law of large numbers ensures that the fraction of accounts that default will be equal to the default probability: $\Pr(X = p(y) | Y = y) = 1$, where X is a random variable indicating the fraction of defaulted accounts, i.e., $X = n/N$. The expression for the probability of default is still given by Eq. (5). In this case it is easier to work out an expression for the PDF of x rather than its density. We can write

$$\begin{aligned} F(x) &= \Pr(X \leq x) = \int_{-\infty}^{+\infty} \Pr(X = p(y) \leq x | Y = y) \phi(y) dy \\ &= \int_{-\infty}^{+\infty} \mathbf{1}_{p(y) \leq x} \phi(y) dy = \int_{-y^*}^{+\infty} \phi(y) dy = \Phi(y^*), \end{aligned} \tag{7}$$

where $\mathbf{1}$ is the indicator function and y^* is defined so that $p(-y^*) = x$ and $p(y) \leq x$ for $y > -y^*$, i.e.,

$$y^* = \frac{1}{\sqrt{\rho}} \left(\sqrt{1-\rho} \Phi^{-1}(x) - K \right). \tag{8}$$

Combining all these results, we can write an expression for the PDF of the fraction of losses:

$$F(x) = \Phi\left(\frac{1}{\sqrt{\rho}} \left(\sqrt{1-\rho} \Phi^{-1}(x) - \Phi^{-1}(p) \right)\right). \tag{9}$$

If we denote the total outstanding balances by B , if we assume that each customer in a given segment of the portfolio carries the same balance, and if we express the recovery rate as $\gamma \in [0, 1]$, we have that the loss at time T implied by a fraction x of consumers defaulting is

$$L = (1 - \gamma)Bx. \tag{10}$$

Since γ and B are constants at the beginning of the period, the PDF of L has the same characteristics as (9).⁷

So far we have described the probability distribution of credit losses.⁸ To make the model comparable with the NBCA proposal for revolving retail, we need to also model FMI. Once we do this, we will be able to subtract our measure of FMI from the tail loss at a given confidence interval and obtain an *economic* measure of capital. In other words, our capital definition will be based on the tail economic loss, rather than the tail credit loss. The economic loss will be given by the difference between projected income over one year minus the tail credit loss and minus the expenses needed to generate that income.

Suppose the financial institution lends B_0 dollars to a certain segment at the beginning of the time horizon. Between time 0 and time T the company will sustain some losses L , but will also collect some revenue R and will have some expenses S . Revenue is generated by the collected finance charges (interest income) and fees (non-interest income) on performing accounts; expenses are incurred to finance the part of B_0 in excess of the company's capital (interest expenses), as well as to market the product, service the accounts and pay for overhead (non-interest expenses). Interest income is obviously related to the initial balances B_0 , but it is also related to the loss L , since finance charges assessed on non-performing accounts will not in general be collected. If we denote the rate applied to outstanding balances by r_f and collected interest income by R_f , we have

$$R_f = r_f B_0 - r_f L = r_f (B_0 - (1 - \gamma) B_0 x) = r_f (1 - x(1 - \gamma)) B_0. \quad (11)$$

In a similar way, non-interest income is related to both initial balances and losses. Here we assume that non-interest income is a constant fraction of outstanding balances, even if fees are assessed in dollar terms rather than as a percentage of balances. Given the assumption that all accounts within the same segment carry the same balance, this is equivalent to assuming that a constant fraction of consumers in each segment pays the annual fee, the late fee, the over-limit fee, and all the other applicable fees. If we denote this constant fraction by λ and non-interest income by R_λ , we have⁹

$$R_\lambda = \lambda B_0 - \lambda L = \lambda (B_0 - (1 - \gamma) B_0 x) = \lambda (1 - x(1 - \gamma)) B_0. \quad (12)$$

⁷ Here we assume that, at default, the loss will be the balance B minus the recovery rate γ . The fact that, in the real world, people might default for more than the average balance can be dealt with in our framework by modifying γ (we can denote it by γ^*). For example, if the average balance for a certain segment is \$1000, the credit limit is \$2000, the recovery rate γ is 20%, but people who default really default for the whole \$2000, then we would set $\gamma^* = -0.6$. In general γ^* can be computed as $\gamma^* = 1 - k(1 - \gamma)$, where k is the multiple of the average balance that is lost at default.

⁸ Gordy (2002) provides a general framework for risk-factor models and shows that, if the portfolio is infinitely granular, a capital allocation based on the tail of the credit loss distribution is portfolio invariant if there is only one factor.

⁹ Note that in both Eqs. (11) and (12) we assume that, if γ is the recovery rate, not only is the principal recovered but so are the corresponding interest and non-interest income.

Total revenue is therefore given by $R = R_f + R_\lambda$:

$$R = (r_f + \lambda)(1 - x(1 - \gamma))B_0. \quad (13)$$

Interest expense is also related to the initial outstanding balances, since that (minus the capital the firm holds) is the amount that needs to be financed. If we call C the capital held, the amount to be financed is $B_0 - C$. Since financing has to occur at the beginning of the period, any loss incurred after that still needs to be financed. If r_b is the average cost of funds applicable to that particular segment,¹⁰ interest expense will therefore be¹¹

$$S_r = r_b B_0 - r_b C = r_b (B_0 - C). \quad (14)$$

As with non-interest income, we assume that non-interest expenses are incurred on a per-account basis and, therefore, are a constant percentage, denoted as ψ , of outstanding balances, since we assume a constant per-account balance. Again, recoveries and losses are not assumed to affect ψ .¹² Non-interest expenses are therefore $S_\psi = \psi B_0$, and total expenses are

$$S = r_b (B_0 - C) + \psi B_0. \quad (15)$$

Putting together Eqs. (10), (13) and (15) we can write an expression for the balances at the end of the period, B_T :

$$\begin{aligned} B_T &= B_0 - L + R - S \\ &= B_0 - x(1 - \gamma)B_0 + (r_f + \lambda)(1 - x(1 - \gamma))B_0 - r_b(B_0 - C) - \psi B_0 \\ &= B_0((1 + r_f + \lambda)(1 - x(1 - \gamma)) - r_b - \psi) + r_b C. \end{aligned} \quad (16)$$

Since x is stochastic with a PDF given by (9), B_T is stochastic as well, and its distribution is completely determined by that of x . If $B_T < B_0$, the company or segment will have suffered an economic loss; it will have shown a profit if $B_T > B_0$. In percentage terms the profit or loss is

$$\pi_B = \left(\frac{B_T}{B_0} - 1 \right) \sim G(F(x), c), \quad (17)$$

where $G(F(x), c)$ is the probability distribution of π_B , which depends on Eq. (9), and $c = C/B_0$. The capital charge will be given by the left tail of $G(F(x), c)$ at an appropriate percentile.

¹⁰ The cost of fund needs not be constant for all segments: internal treasury units can charge lower funds transfer prices to better quality segments to reflect the easier access to external financing (for example, securitization) for those accounts and the fact that they resemble high quality instruments.

¹¹ This assumes, again for simplicity, that there is no balance paydown within the time horizon at issue.

¹² In practice non-interest expense includes costs such as marketing and setup costs, etc., that are sustained initially and therefore apply to both performing and non-performing accounts. Other costs, such as phone calls, mailing of statements, etc., apply only to accounts that are not charged off, and others, such as the cost of recoveries, apply only to charged off accounts. Here we assume that the latter two types of costs are the same and, therefore, non-interest expense is constant with the loss. Barakova and Carey (2003), however, find that banks that survived after sustaining high credit losses experienced strong increases in non-interest expense coincident with and following the bad tail event. This “cost shock” was enough to wipe out the banks’ net income even after accounting for provisions.

Note that in Eq. (17) π_B depends on c , the capital ratio, but c depends in its turn on π_B . The determination of the capital charge, and of the distribution of π_B , is therefore an iterative process: one can start with an arbitrary capital charge, compute the percentage return distribution and from it a new capital ratio; then use this new capital ratio to repeat the process until the capital ratio converges. Formally, we can write

$$c_k = f_\alpha(G(F(x), c_{k-1})), \tag{18}$$

where f_α denotes the α -percentile of $G(F(x), c_{k-1})$. Eq. (18) shows that the capital ratio c is the fixed point of a difference equation. Since $r_b < 1$, the solution to (18) will converge to a certain c , which will depend on all the other parameters. We can easily solve for such c by noting that $f_\alpha(G(F(x), c_{k-1})) = B_T(x_\alpha, c_{k-1})/B_0 - 1$, where x_α is that fraction of defaults such that the probability of x_α or fewer defaults happening will be exactly α , or $F(x_\alpha) = \alpha$, with $F(x_\alpha)$ given by (9). Using Eqs. (16) and (18), we can therefore write

$$c_k = (1 + r_f + \lambda)(1 - x_\alpha(1 - \gamma)) - r_b - \psi + r_b c_{k-1} - 1 \tag{19}$$

from which it follows that the capital ratio c for each segment is

$$c = \frac{(r_f + \lambda - r_b - \psi) - (1 + r_f + \lambda)(1 - \gamma)x_\alpha}{1 - r_b}. \tag{20}$$

Note that c is defined as the α -percentile of the earnings distribution.¹³ A bank will have to hold capital to face potential negative earnings and, therefore, c is in general a negative number, which becomes larger in absolute value (i.e., more negative) as x_α increases and γ decreases. It is possible, at least in principle, that a certain portfolio generates a very high net income relative to its tail loss and that, based on Eq. (20), c will be zero or positive.¹⁴ Since a positive c would imply that a bank will have to hold *negative* capital, we redefine c as

$$c = \min \left(\frac{(r_f + \lambda - r_b - \psi) - (1 + r_f + \lambda)(1 - \gamma)x_\alpha}{1 - r_b}, 0 \right). \tag{21}$$

The new Basel proposal uses the following formula to obtain the capital ratio for each given probability of default:

$$c = \text{LGD} \cdot \Phi \left(\frac{\sqrt{\rho} \Phi^{-1}(\alpha) + \Phi^{-1}(p)}{\sqrt{1 - \rho}} \right) - 0.75 \cdot p \cdot \text{LGD}, \tag{22}$$

¹³ Here earnings are intended in an economic sense rather than GAAP.

¹⁴ This does not typically happen in the real world, where the competition for accounts with low tail losses forces the net income generated by those portfolios to be relatively low.

where $p \cdot \text{LGD}$ is the expected loss and α is the desired percentile, i.e., $\alpha = 0.999$ for a 99.9% confidence level.¹⁵

The Basel formula is obviously related to our model, as it can be obtained from the same one-factor loss model.¹⁶ Note that the first term of Eq. (22) can be derived from Eq. (9). The latter gives the probability that the fraction of defaults will be less than any given number x . We are interested in particular in x_x , which can be found by inverting Eq. (9):

$$F^{-1}(x_x) = \frac{\sqrt{1-\rho}\Phi^{-1}(x_x) - \Phi^{-1}(p)}{\sqrt{\rho}} = \Phi^{-1}(\alpha) \quad (23)$$

and therefore,

$$x_x = \Phi\left(\frac{\sqrt{\rho}\Phi^{-1}(\alpha) + \Phi^{-1}(p)}{\sqrt{1-\rho}}\right). \quad (24)$$

Eq. (22) is obtained by multiplying the loss by the LGD (or $1 - \gamma$ in our notation) and by approximating future margin income by 75% of the expected loss. The only difference with our model is in the way we account for FMI.

2.1. Determination of the models parameters

We calibrate our model based on a mini-portfolio that we assembled, consisting of six different risk groups from the overall Capital One database of “test” solicitations.¹⁷ We selected accounts from two different products, which we label Product 1 and Product 2; as described in Table 1, Product 1 is a low-margin, low-credit-risk product, whereas Product 2 has a high margin and is riskier, in terms of probability of default.¹⁸ Also note that Risk Group 3 is the only one in our data set that does not satisfy the FMI qualification criterion that future revenue should be in excess of twice the loss standard deviation. We then segmented each product in three homogeneous risk groups, based on probability of default; given the nature of the two products, it turns out that the lowest p for Product 2 is higher than the highest p for Product 1. Note that the products and groups within each product were purposefully selected to span the whole credit spectrum, rather than to match the actual

¹⁵ One odd feature of Eq. (22) is that for probabilities of default higher than about 67%, the capital that a bank is required to hold actually starts decreasing. This obviously might leave some risk managers feeling a little uneasy, but, also 67% is obviously a huge probability of default, not likely to be seen in most portfolios. Moreover, the FMI qualification criterion likely would disqualify any hypothetical segments with that kind of probability of default from revolving retail treatment.

¹⁶ Although they do not focus on retail portfolios, Gordy (2002) and Wilde (2001), also show the relationship between a one-factor model and previous versions of the Basel formula.

¹⁷ Capital One performs numerous product tests that, if successful, will eventually develop into full-scale lending. At the same time the company maintains a rich database monitoring the performance of those accounts over time.

¹⁸ We are reluctant to label the products with any of the common designations of superprime, prime, near prime, or subprime because no precise industry or regulatory definition of those terms has ever been agreed upon.

Table 1
Product characteristics

Risk group	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Product	1	1	1	2	2	2
Probability of default	p_1	$p_2 > p_1$	$p_3 > p_2$	$p_4 > p_3$	$p_5 > p_4$	$p_6 > p_5$
Loss given default	High	High	High	Low	Low	Low
Basel FMI criterium	Yes	Yes	No	Yes	Yes	Yes

portfolio of any real-world credit card company; this is ideal for our analysis, since we can study capital allocations for a large range of probabilities of default not normally seen in traditional, actual portfolios.

Because of confidentiality issues, we cannot disclose the information relative to the exact probabilities of default of each risk group, or to their income, expenses, and specific losses. However, below we use data pertinent to each segment as inputs to the model to obtain theoretical capital ratios for each of them. We use historical averages over the time period spanned by our data set, which goes from January 2000 to February 2003 and therefore includes the 2001 recession.

The only other input needed by the model is the asset correlation. Estimating the asset correlation coefficient is problematic. The BCBS has provided specific mappings between probability of default p and asset correlation ρ for different portfolios for every iteration of the consultative paper (see BCBS, 1999, 2001, 2003). The latest mapping (BCBS, 2003) for revolving retail portfolios is

$$\rho = 0.02 \cdot \frac{1 - e^{-50p}}{1 - e^{-50}} + 0.11 \cdot \left(1 - \frac{1 - e^{-50p}}{1 - e^{-50}} \right). \quad (25)$$

According to this formula, the asset correlation varies between 11% for very low probability of default segments to 2% for segments with p in excess of about 10%.¹⁹ As a first approximation we will use it in the calculations here. We would like to spend a few words, however, cautioning that the model is sensitive to this parameter. In particular, if ρ is even a few percentage points higher than what Basel assumes, capital for high- p segments can be sharply higher. Similarly, if ρ is lower than what Basel says for low- p segments, their capital can be drastically reduced.

We also note that the price to be paid in case asset correlation is, say, constant within a portfolio, or, even worse, in case it actually slopes up with p , is potentially very high. This is because a downward sloping relationship between ρ and p mitigates the size of extreme losses, and a lower ρ would make the loss distribution less tail-heavy. If in reality the two effects do not offset each other, or worse if they reinforce each other, this could lead to a gross underestimation of the size of extreme losses. We do not attempt here to estimate asset correlations from our data, since they span a time horizon of only a few years, with only one minor recession (at least from a consumer-retail point of view). We will, however, derive a set of “implied” asset correlations from our multi-factor model in Section 3.1 below.

¹⁹ Previous versions of the proposal had a ρ varying between 15% and 2%.

2.2. Capital ratios

In this subsection we apply the model above to our data set and determine the capital ratio for each segment at a confidence interval of 99.9% over a one-year period and we compare them with those generated by Basel. Note that since the credit loss framework and the asset correlation assumptions are the same for the two approaches, the differences are mainly due to future margin income accounting. We compute our results under two assumptions concerning the LGD.

2.2.1. Constant LGD

First, we use a constant LGD for all segments, representative of assumptions typically made for credit card portfolios;²⁰ the results are reported in Table 2.²¹ According to the model, capital ratios for Product 1 range between 1.95% and 8.33%, while for Product 2 they range between 4.91% and 21.67%.

One interesting thing to note is that for Product 1 the ratios are relatively close to those indicated by Basel, even when taking the FMI provision into account. The ratios are higher than the models for the first group because of the way Basel gives credit for FMI, as a fraction of the expected loss: for groups with very low expected loss, such as Group 1, the FMI credit is almost insignificant. On the other hand, according to our approach, even in the event of a tail loss, the product could generate income that, although not as large as that generated by Product 2, is still greater than what Basel assumes.

Another interesting thing to note is that for the best Product 2 risk group (group 4), the capital ratio according to our model is actually lower than that for the worst segment of Product 1, even if the probability of default is higher. This is due to the fact that the revenue that Product 2 generates is significantly higher than that of Product 1; in particular the higher revenue more than offsets the higher losses due to the higher probability of default. Basel, again because of the way it accounts for FMI, produces a monotonic capital curve.

2.2.2. Product-specific LGD

Table 3 reports the capital ratios using each segment's estimated LGD instead of a constant LGD. The LGD was estimated for each risk group as the ratio of the average net loss per *defaulted* account and the average daily balance per *booked* account over the life of the products from origination to the last month in our sample. Product 1 has in general much higher credit lines than Product 2 and a much lower aver-

²⁰ Risk Management Association (2003) reports capital ratios for different probabilities of default averaged over a number of large US banks, computed under the constant LGD assumption.

²¹ Column (2) of the table reports the Basel capital ratios according to the "other retail" capital formula, which would apply to those segments that do not meet the FMI provision (BCBS, 2003). It is the same as the revolving retail formula, except that there is no subtraction of 75% of the expected loss and asset correlation is higher for low probabilities of default.

Table 2
Capital ratios, one-factor model vs. Basel, constant LGD

Product	Risk group	OF model (1)	Basel-No FMI (2)	Basel-FMI (3)	Difference (3) or (2)–(1)
1	Group 1	1.95%	5.39%	3.15%	+1.20%
1	Group 2	4.19%	7.29%	4.24%	+0.05%
1	Group 3	8.33%	10.66%	6.05%	–2.33%
2	Group 4	4.91%	18.30%	10.68%	+5.77%
2	Group 5	9.81%	24.80%	14.52%	+4.71%
2	Group 6	21.67%	32.84%	18.34%	–3.33%

OF model: one-factor model; Basel-No FMI: Basel formula without FMI provision; Basel-FMI: Basel formula with FMI provision.

Table 3
Capital ratios, one-factor model vs. Basel, product-specific LGD

Product	Risk group	OF model (1)	Basel-No FMI (2)	Basel-FMI (3)	Difference (3) or (2)–(1)
1	Group 1	7.59%	10.97%	6.41%	–1.18%
1	Group 2	11.79%	14.94%	8.69%	–3.10%
1	Group 3	17.76%	20.52%	11.64%	+2.76%
2	Group 4	3.53%	17.31%	10.11%	+6.58%
2	Group 5	3.21%	20.61%	12.06%	+8.85%
3	Group 6	4.66%	22.32%	12.46%	+7.80%

OF model: one-factor model; Basel-No FMI: Basel formula without FMI provision; Basel-FMI: Basel formula with FMI provision.

age utilization rate; accounts in default, however, tend to have a higher than average balance and, therefore, even after accounting for recoveries, the LGD is high. Product 2, on the other hand, is characterized by low credit lines and high utilization rates; accounts in default cannot and do not have balances much higher than the typical average balance, and after accounting for recoveries the LGD is low.

With the LGD differentiated by risk group, the capital ratios are significantly different from those in Table 2. In particular, those for Product 1 are now much higher, whereas those for Product 2 are significantly lower. Aside from the magnitude of the capital ratios, note the qualitative fact that now it is not just our model that says group 4 should have lower capital ratios than group 3, but also Basel's formula. This is the result of the combined effect of high margin income and low LGD: the latter, with respect to the constant LGD case, basically returns to the segments of Product 2 not only a share of the principal but also a corresponding share of the revenues. Note that Basel's formula approximates FMI with a percentage of the expected loss and, therefore, capital decreases one-for-one with the LGD. In our one-factor model, as the LGD decreases, capital ratios decrease faster, because FMI increases as γ decreases. Independently of our modelling assumptions, the non-monotonicity of the Basel capital curve questions the commonly held belief that products with higher probabilities of default are necessarily riskier. Importantly, it also seems to be at

odds with what is known as “subprime regulatory guidance”, whereby segments with FICO scores below 660 should have capital ratios two to three times higher than those for “superprime” segments.²²

Given the importance of the LGD in determining the capital ratios, an interesting question to ask is whether a segment’s LGD can be controlled by the lending institution, at least to some extent, at the outset. Of course, once lending is extended and default is suspected or under way, aggressive account management practices and recovery strategies might help reduce the LGD. However, a careful selection of the credit limit might help reduce the LGD *ex ante*. Lower-risk segments typically have high credit limits. Performing accounts use only a fraction of that credit limit, so the outstanding balances are smaller, and possibly much smaller, than the available credit, or “open to buy”. Accounts that default, however, typically do so for amounts higher than the average outstanding balance and often close to the limit, hence producing LGDs in excess of 100% of average outstandings. Higher-risk segments, on the other hand, typically have lower credit limits, and the average outstanding balance is much closer to the available credit, producing LGDs typically lower than 100% after accounting for recoveries. The benefits of keeping credit limits under control are clear. For low risk segments, they help keep the LGD, and therefore the capital ratio, low.²³ For high-risk segments low limits are even more important, as the combination of high probability of defaults and high LGDs could be fatal for a lending institution, unless extremely high levels of capital are held as a cushion.

3. A multi-factor model of the income statement

So far we have assumed that defaults depend on only one unnamed factor. That setup has the advantage of being relatively simple and very tractable. One shortcoming, however, is that the probability of default is the same at all times: recessions, for example, or the value of any other macroeconomic or financial variable, had no effect on defaults. In the same way, the collection of revenues and the expenses are completely independent of macroeconomic conditions. Furthermore, defaults could occur only at time T and, in general, the assumptions we made concerning the structure of the model were very strong and assumed to be stable over time.

In this section we address these issues by presenting a multi-factor model of the income statement,²⁴ where the factors are some specified variables. We model the whole income statement of a given risk segment, since we believe, as in Section 2, that economic capital should be allocated not based on credit losses but on economic

²² See Supervisory Letter SR01-04, Board of Governors (2001).

²³ Of course there are other reasons that determine the credit limit extended to certain segments. For example, offering too low a limit to good credit risks might be impractical or impossible for an institution if the competitors are all offering higher limits.

²⁴ Again, from an economic perspective rather than GAAP.

losses. According to this new approach, every broad category of the income statement is modelled separately over a certain time horizon (one year in our case) and for several sub-periods of the time horizon (12 months): there will therefore be probability distributions for losses, interest income, fee revenue, etc. When aggregated, these will generate a probability distribution of earnings; the economic capital allocated to a particular segment will be the left tail of this distribution, at a certain appropriately chosen confidence level. The fact that the income statement is modelled month by month is important because in this way capital will be held to face economic losses that could occur over periods shorter than one year and that could be missed by models that look only at events at the horizon T .

We choose to estimate a relationship between each of the income statement variables and the economy, as represented by some key variables. The choice of these variables is largely dependent on the individual institution and even on the particular product. For example, losses on mortgages might be influenced by different macro-economic variables than losses on credit cards or auto loans. From an analytic point of view, we still assume that defaults are triggered when the value of the assets of an individual falls below a certain threshold, as in the previous section. Now, however, we assume that ²⁵

$$V_t = \sum_{j=1}^h \delta_j m_{j,t-l} + \zeta_t, \quad (26)$$

where V_t is the value of assets for consumers in each segment (assumed to be homogeneous), ²⁶ m_j , $j = 1, \dots, h$, is the set of driving variables, and ζ_t is the idiosyncratic error for each consumer. We allow for a lag $l > 0$ in the model, as charge-offs can contractually occur only a number of months (usually four or six) after the default decision was actually made. ²⁷ We assume that the driving variables have a correlation matrix Σ and that $\zeta_t \sim N(0, 1)$.

The problem that we have now is to estimate the parameters δ_j . This can obviously not be done directly since the assets value V is unobservable, but it can be done using a probit approach. ²⁸ We know that each segment has a probability of default p and that default will occur when assets fall below a threshold K related to p in the usual way. Then the probability of default, conditional on the realization m of the driving variables, is ²⁹

²⁵ Our multi-factor credit loss framework again follows Schönbucher (2000).

²⁶ For notational convenience we drop the subscript i to indicate each segment. Equation (26) and all other equations below will be the same for all segments but will have different parameters and possibly different sets of driving variables.

²⁷ The assumption that $l > 0$ also has the advantage of considerably simplifying the analysis, as the $m_{j,t-l}$ are known at time t and, therefore, we do not have to worry about their distribution.

²⁸ This is a latent variable problem and can be addressed by either probit or logit regression.

²⁹ Note that this specification is very similar to CreditPortfolioView, a model proposed by McKinsey; see Wilson (1997), and Crouhy et al. (2000). That model uses a logit, rather than probit approach.

$$p_t(m_t) = \Pr \left(\zeta_t < K - \sum_{j=1}^h \delta_j m_{j,t-l} \right) = \Phi \left(K - \sum_{j=1}^h \delta_j m_{j,t-l} \right). \tag{27}$$

Let us denote the default rate in any given month t as $x_t(m_t)$. Then $x_t(m_t)$ will be our estimate of $p_t(m_t)$ and the parameters δ_j can be estimated for each risk group. Once the coefficients are known, a probability distribution of default fractions x can then be generated.

If we assume again that all consumers within a segment will carry the same balance, then the dollar loss will be $L_t = x_t B_t (1 - \gamma)$, where γ is the recovery rate. Let us denote the income statement (IS henceforth) variables other than losses for each segment at time t as $s_{k,t}$, with $k = 1, \dots, n$, where n is the number of variables that are explicitly modelled. Each of them will depend on losses (to capture the fact that the higher the losses, the lower collected revenue will be), as well as on the driving variables (to capture possible dependency of consumers' spending patterns and card usage, non-interest and interest expenses, etc. on the same driving variables). Since the loss itself depends on the driving variables, we can write the following reduced form for the IS variables:

$$s_{k,t} = f(m_{t-l}, \epsilon_{k,t}, \rho_k). \tag{28}$$

Here ϵ denotes the error term for each IS variable for each segment and ρ_k denotes the autocorrelation coefficient of the error term. We allow again for lags, as we did for losses.

Taking the possibility of autocorrelation in the error term into account is important when simulating the model. Since there is no particular theory behind a relationship such as (28), one does not have any good reason to believe that the error is pure white noise, and indeed it will most likely not be, since there are conceivably many variables that determine the behavior of consumers in relation to chargeoffs and other IS variables than those included in Eq. (28). If the error term is autocorrelated, and such autocorrelation is not taken into account, simulations based on an equation such as (28) will tend to systematically over or underestimate the magnitude of the variable being simulated, depending on whether the error at the starting point of the simulation is positive or negative. Estimating Eq. (28) with the explicit consideration of ρ_k guarantees that there will not be such systematic bias.

From an operational point of view, a linear functional form for (28) is assumed:

$$s_{k,t} = \beta_{k,0} + \sum_{j=1}^h \beta_{k,j} m_{j,t-l} + \rho_k \epsilon_{k,t-1} + v_{k,t}, \tag{29}$$

where it is assumed that there is just first-order autocorrelation in the error term, i.e., $\epsilon_{k,t} = \rho_k \epsilon_{k,t-1} + v_{k,t}$. Eq. (29) can be estimated in a number of ways, including maximum likelihood, two-step iterative Cochrane–Orcutt, etc. (see Hamilton (1994) or many other possible references).

Once the choice of which income statement variables to model has been made, and once the driving variables to use have been identified, the parameters in Eq. (29) can easily be estimated given a data set of past values of the driving and the

IS variables. Simulated paths for each of the IS variables can then be generated as a function of simulated paths for the key driving variables.³⁰ One therefore has to specify some process for each of those variables which should depend on what that variable is. For example, a stochastic process with a unit root could be a good approximation for GDP, unemployment, etc. If interest rates are among the h driving variables, however, some information about the future values of rates for any given maturity is contained in the current yield curve; a better framework than a random walk could therefore be the one proposed by Hull and White (1990) or any of the many other term structure models available in the literature (see Hull, 1997, for a survey).

Once simulated paths for all the h driving variables have been generated and the relationships in (29) between the income statement variables and the macro variables have been estimated, a simulated path for earnings can easily be generated by adding together all the income statement variables:

$$\theta_t = \sum_{k=1}^n s_{k,t}(\beta, \rho, \Sigma) + L_t(\delta, \Sigma), \quad (30)$$

where θ denotes earnings and the income statement variables are taken with the appropriate sign (i.e., positive for revenues and negative for expenses and losses). Eq. (30) makes explicit the fact that the resulting paths for earnings will depend on the simulated conditions of the economy through the regression parameters β , ρ , and δ and through the correlation matrix Σ among the driving variables. If sufficiently many paths for earnings are generated, a probability distribution function for θ , $F(\theta; \beta, \delta, \rho, \Sigma)$, can be estimated. Then the capital that a firm will hold is

$$c = \theta^{(\alpha)} = \min(F^{-1}(\alpha, \beta, \delta, \rho, \Sigma), 0), \quad (31)$$

where α is an appropriately chosen percentile, such as 99.9%. As in the one-factor case, Eq. (31) says that the capital ratio will be the left α -tail of the distribution of earnings; if simulated earnings for some reason turn out to be always positive, capital will be zero.

3.1. Capital ratios

We applied the model described in this section to our sample data set consisting of the same six risk segments used for the one-factor model. Confidentiality agreements prevent us from discussing the exact variables used in the estimation, the estimated parameters and other details. The capital ratios for the nine segments are reported in Table 4, column (1), together with the capital ratios obtained in the previous section for the one-factor model with product-specific LGD (column (2)) and those generated by the Basel formula (column (3)).

³⁰ The paths for the driving variables are simulated taking their correlation into account via a Cholesky decomposition of the correlation matrix Σ .

Table 4

Capital ratios, multi-factor model vs. one-factor model and Basel, product-specific LGD

Product	Risk group	MF model (1)	OF model (2)	Basel (3)	Difference (3)–(1)
1	Group 1	1.00%	7.59%	6.41%	+5.41%
1	Group 2	6.60%	11.79%	8.69%	+2.09%
1	Group 3	15.85%	17.76%	20.52%	+4.67%
2	Group 4	11.12%	3.53%	10.11%	–1.01%
2	Group 5	20.97%	3.21%	12.06%	–8.91%
2	Group 6	22.71%	4.66%	12.46%	–10.25%

MF model: multi-factor model; OF model: one-factor model; Basel: Basel formula with or without FMI provision.

The multi-factor model ratios are lower than both the Basel's and the one-factor model ratios for high credit quality segments (Product 1), and drastically so for group 1; they are higher for the other segments, especially for groups 5 and 6. Note that the non-monotonicity of the capital curve persists also in the multi-factor model, although the large ratios for groups 5 and 6 make the results more compatible with subprime guidance, not only qualitatively (in that high probability of default segments hold more capital than low probability of default segments) but also quantitatively, in that the capital ratios assigned to groups 4–6 are within (or even in excess of) the 1.5–3 times range of the capital ratios assigned to groups 1–3.

The way we model future margin income no doubt accounts for part of the differences in capital ratios between our multi-factor model and Basel, as was the case for the one-factor model. Here, however, the assumptions underlying the model are not the same as Basel and, therefore, several other factors can account for the discrepancies. First, consider that the estimation period contains a recession; everything else being equal, the multi-factor model says that economic conditions affect high-risk segments more than low-risk segments. Recall also that while the one-factor model contained an explicit asset correlation parameter, which was calibrated according to Basel's assumptions, the multi-factor model does not. The results might therefore suggest that asset correlations, at least over our estimation period, do not follow Basel's assumed pattern of declining as probabilities of default increase; they might actually increase with risk, leading to the problem discussed earlier of Basel's possibly underestimating capital for riskier segments. To further expand on this point, we show on Table 5 what we call the "implied" asset correlations. These are obtained by asking what asset correlation parameter would force Basel's formula to assign to each risk group the same capital ratio assigned by our multi-factor model.³¹ As one can see, not only the implied asset correlation for the groups in Product 1 (low risk) are lower than those for corresponding groups in Product 2 (higher-risk), but they are also increasing with p within each product; both facts are not in agreement with Basel's assumed calibration.

³¹ No positive asset correlation could make Basel's capital ratio for group 1 equal to 1%.

Table 5
Implied vs. Basel asset correlations

Product	Risk group	Implied asset correlations (1)	Basel asset correlations (2)	Difference (2) – (1)
1	Group 1	–	9.39%	–
1	Group 2	6.17%	8.46%	+2.29%
1	Group 3	9.05%	6.19%	–2.86%
2	Group 4	2.56%	2.13%	–0.43%
2	Group 5	6.08%	2.01%	–4.07%
2	Group 6	7.41%	2.00%	–5.41%

Second, the multi-factor models assumptions about the structure or the probability distribution of losses are less stringent than in the one-factor case, as the driving variables play a large role in it; it might be the case that the actual probability distribution of losses is different from the one implied by the one-factor model or Basel. This is related to the asset correlation problem discussed above, as the asset correlation parameter directly affects the shape of the probability distribution of losses, as Eq. (9) shows. However, the probability distribution could be different even if the real asset correlation coefficients scaled with p as postulated by Basel.

Third, the multi-factor model allows revenue to evolve randomly over time. The volatility of revenue turns out to be much higher for high- p segments than for low- p segments,³² in part because their response to changes in macroeconomic conditions seems to be stronger. This high volatility generates some simulated paths for the various components of revenue, and for non-interest income in particular, that are well below the historical averages and hence leads to higher capital ratios. This is interesting per se, but we find it to be particularly appropriate and desirable, as it effectively forces riskier segments to hold capital in case something unforeseen, such as changes in the legal or regulatory environment, cuts revenue drastically and permanently.³³

The magnitude of the capital ratios generated by the model relative to Basel's poses some incentive problems. Without entering into the details about securitization, which is outside the scope of this paper, and to the extent that the results of our multi-factor model are representative of the true risks facing credit card issuers in general, we argue that there could be incentives for retail banks to take on more risky loans, or to take off balance sheet less risky ones, in a way that is not dissimilar from the situation generated by the 1988 accord currently in force. Moreover, since Basel capital charges for loans that are usually considered relatively safe, such as

³² It is indeed extremely low for groups 1–3, even lower than what is implied by Eq. (13) in the one-factor model.

³³ For example, some countries have recently begun to question the size of the interchange fee that credit card companies and/or the international circuits such as Visa or Mastercard, collect from merchants for the privilege of allowing them to accept their card. There is no way to account for this risk in the one-factor model or in Basel, except to classify it as an operations risk. We prefer to think about it as an economic risk, as it affects the business directly rather than being incidental to it.

“platinum” credit cards, can be quite high, certain banks might be induced to opt for the standardized approach rather than the advanced IRB approach.³⁴

4. Conclusions

We presented two models that could in principle be used to allocate economic capital across risk segments and compared the capital ratios with those generated by the most recent version of the Basel formula for revolving retail exposures, which allows for the recognition of future margin income in the capital calculations. We found that Basel’s capital ratios are close to those generated by the one-factor model for low-risk segments. Moreover, both sets of capital ratios, which are based on similar assumptions, sometimes generate counterintuitive capital ratios. Specifically, the capital ratios of risk segments with high probability of default can be lower than those for segments with low probability of default, if the loss-given-default for the former is significantly lower and the revenue they generate significantly higher. The multi-factor model ties the capital ratios to economic conditions and relaxes many assumptions and generates capital ratios that are more in agreement with the common belief that low-credit-risk segments should hold less capital than high-credit-risk segments. In addition, the capital ratios obtained from the multi-factor model could indicate that Basel’s assumptions about how asset correlations change with the probability of default might be inaccurate, especially at the low and high end of the credit spectrum.

References

- Altman, E.I., Saunders, A., 2001. An analysis and critique of the BIS proposal on capital adequacy and ratings. *Journal of Banking and Finance* 25, 25–46.
- Barakova, I., Carey, M., 2003. How quickly do troubled banks recapitalize? With implications for portfolio VaR credit loss horizons. Working Paper, Federal Reserve Board.
- Basel Committee on Banking Supervision, 1999. A new capital adequacy framework. Available online at <http://www.bis.org>.
- Basel Committee on Banking Supervision, 2001. The new Basel Capital Accord, second consultative paper. Available online at <http://www.bis.org>.
- Basel Committee on Banking Supervision, 2003. The new Basel Capital Accord, third consultative paper. Available at <http://www.bis.org>.

³⁴ Other authors have found incentive problems with the new accord. Notably, Altman and Saunders (2001) found that reliance on agency ratings rather than internal ratings could produce cyclically lagging capital requirements. Calem and LaCour-Little (2001) argue that appropriate risk-based capital ratios for mortgage loans are generally below current regulatory ratios and may help explain the high degree of securitization of those loans. Kupiec (2001a,b) observes that, given the current proposed calibration, IRB banks will tend to concentrate on high quality lending, whereas standardized approach banks will tend to concentrate on risky lending, and that IRB banks will tend to prefer loans that are more likely to default in recessions.

- Board of Governors of the Federal Reserve System, 2001. Subprime lending. Supervisory Letter SR01-04. Available at <http://www.federalreserve.gov/boarddocs/srletters>.
- Calem, P.S., LaCour-Little, M., 2001. Risk-based capital requirements for mortgage loans. Financial and Economics Discussion Series 2001-60, Board of Governors of the Federal Reserve System.
- Crouhy, M., Galai, D., Mark, R., 2000. A comparative analysis of current credit risk models. *Journal of Banking and Finance* 24, 59–117.
- Gordy, M.B., 2002. A risk-factor model foundation for ratings-based bank capital rules. Finance and Economics Discussion Series 2002-55, Board of Governors of the Federal Reserve System. *Journal of Financial Intermediation*, Forthcoming.
- Hamilton, J.D., 1994. *Time Series Analysis*. Princeton University Press, Princeton.
- Hull, J.C., 1997. *Options, Futures, and Other Derivatives*. Prentice Hall, New York.
- Hull, J.C., White, A., 1990. Pricing interest rate derivative securities. *Review of Financial Studies* 3, 573–592.
- KMV Corporation, 1993. Portfolio management of default risk. Working paper, Available online at <http://www.kmv.com>.
- Kupiec, P.H., 2001a. The new Basel Capital Accord: The devil is in the (calibration) details. IMF Working Paper, Available at www.imf.org.
- Kupiec, P.H., 2001b. Is the new Basel Accord incentive compatible? IMF Working Paper, Available at www.imf.org.
- Risk Management Association, 2003. Retail credit economic capital estimation—best practices. Available online at <http://www.rmahq.org>.
- RiskMetrics Group, 1997. CreditMetrics, technical document. Available online at <http://www.riskmetrics.com>.
- Schönbucher, P.J., 2000. Factor models for portfolio credit risk. Manuscript, University of Bonn, available online at <http://www.defaultrisk.com>.
- Treacy, W.F., Carey, M., 2000. Credit risk systems at large US banks. *Journal of Banking and Finance* 24, 167–201.
- Vasicek, O.A., 1987. Probability of loss on loan portfolio. KMV Corporation, Available at <http://www.kmv.com>.
- Wilde, T., 2001. IRB approach explained. *Risk* 14 (5).
- Wilson, T., 1997. Portfolio credit risk. *Risk* 10 (9 and 10).