Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Journal of Banking & Finance 28 (2004) 695–720
Journal of BANKING & FINANCE
www.elsevier.com/locate/econbase

An anatomy of rating through the cycle

Gunter Löffler *

Goethe-Universität Frankfurt, Mertonstrasse 17, Frankfurt am Main 60054, Germany

Received 5 August 2002; accepted 24 December 2002

Abstract

Using a structural model of default, I derive rating characteristics if ratings are meant to look 'through the cycle' as opposed to being based on the borrowers' current condition. The through-the-cycle method, which is employed by most rating agencies, requires a separation of permanent and cyclical components of default risk. In a time series setting, this can be done through the Kalman filter. The analysis shows that several empirical irregularities of agency ratings could be the consequence of such a rating method. The stability of through-the-cycle ratings is relatively high, while their default prediction power is low. Though not predictable in the usual sense, rating changes exhibit properties that call for a reconsideration of the existing evidence.

JEL classification: G20; G33
Keywords: Credit ratings; Rating agencies; Credit cycles; Kalman filter; Efficiency; Conservatism

1. Introduction

Rating agencies such as Fitch, Moody's or Standard & Poor's play an important role for the functioning of credit markets. Their ratings are used to price risky debt, to compute economic and regulatory capital, or to calibrate internal ratings of banks and other financial institutions (for the latter see Carey and Hrycay, 2001). When using agency ratings for such purposes, two main requirements should be met. The nature of a rating should be properly understood, i.e., it should be clear what kind of information rating agencies intend to summarize. Secondly, ratings should

* Tel.: +49-69-7982-3143; fax: +49-69-7982-2143.
E-mail address: gloeffler@wiwi.uni-frankfurt.de (G. Löffler).

0378-4266/$ - see front matter
doi:10.1016/S0378-4266(03)00041-4

efficiently aggregate this information. At present, both requirements do not seem to be fulfilled.

As to informational efficiency, there is plenty of academic and anecdotal evidence which suggests that agency ratings do not fully reflect available information. Altman and Kao (1992) and Lando and Skødeberg (2002) document serial dependence in rating changes, while Delianedis and Geske (1999) conclude that borrower fundamentals predict future rating changes. Other evidence which points to inefficiencies is the high stability of agency ratings (Kealhofer et al., 1998), or the agencies' performance in the Asian crises (IMF, 1999). [1]

The peculiarities of the agencies' rating method, on the other hand, have received little attention in the literature. It is commonplace to note that agency ratings are not estimates of short-term default risk, but should rather be characterized as looking through the cycle (cf. Basel Committee on Banking Supervision, 2000). Carey and Hrycay (2001), however, are the first to study the nature and consequences of the agencies' rating architecture. They examine problems that arise if the default history of through-the-cycle ratings is used to map internal bank ratings into default probabilities. Contrary to agency ratings, bank ratings are usually based on the actual default probability over a specific horizon. In the literature, such ratings are labeled current-condition or point-in-time ratings. The category also comprises ratings based on quantitative forecasts of bankruptcy. [2]

In this context, the contribution of the present paper is twofold. I suggest a formal model of the agencies' rating process that should be helpful whenever their ratings are to be assessed, or used as an input to other models. As an application, I use Monte Carlo simulations to examine the following question: do the empirical peculiarities of agency ratings necessarily reflect informational inefficiencies, or could they be inherent to the agencies' rating system?

The model builds on the structural analysis of credit risk introduced by Merton (1974). Firms default if the value of their assets hits a default threshold which is related to their liabilities. The key variable for measuring default risk is distance to default, which is the standardized difference between the asset value and the default threshold. Differentiating between through-the-cycle and current-condition ratings is relevant when a borrower's default risk exhibits cyclical behavior. I therefore consider a case in which asset values are subject to both permanent and transitory shocks. Through-the-cycle ratings respond to permanent shocks only, current-condition ratings to both. More precisely, through-the-cycle ratings are assumed to be based on the borrower's distance to default in the event a stress scenario occurs. The stress scenario is taken to be a transitory deviation from the normal condition which occurs with a fixed probability chosen by the rating agency. In

[1] Cf. also the following quotes from the financial press: "Wall Street fixed income analysts often observe that their counterparts at rating agencies are, overly conservative" (Institutional Investor, September 1997, p. 197) and "Ratings in Asia seem to lag rather than tell you any worthwhile information in advance" (Euromoney, January 1998, p. 51).

[2] Cf. Shumway (2001) for a recent contribution to bankruptcy forecasting, and Crouhy et al. (2000) for a description of the default rate estimates produced by the financial software firm KMV.

the paper, the stress scenario is defined by assuming plausible values for the volatility of transitory shocks, and the probability of the stress scenario's occurrence. The larger the volatility of transitory shocks, the more adverse is the shock that is added to the permanent distance to default in the construction of the stress scenario. Thus, vulnerability to cycles affects the rating decision, whereas the current position in the cycle does not. Current-condition ratings are based on the current distance to default, which incorporates the permanent asset value plus cyclical deviations; they take into account that cyclical deviations tend to decline over the rating horizon.

Under both rating architectures, rating analysts need to separate permanent and cyclical components. Typically, the two components are not directly observable, but have to be estimated. In the framework of the paper, the optimal estimation procedure is to apply a Kalman filter, [3] which produces a least squares decomposition of the observed variable into permanent and transitory components. The paper does not distinguish between macroeconomic, industry-specific, or idiosyncratic shocks because the agencies' publications do not indicate that rating decisions depend on the origin of a shock. [4]

In such a setting, rating through the cycle has several implications which mirror empirical findings. Through-the-cycle ratings are relatively stable, and have a low default prediction power; rating changes are correlated with past rating changes provided contemporaneous information is controlled for. Further predictions arise if the degree of cyclicality is not known to the rating analysts. If raters overestimate the magnitude of transitory shocks, for example, rating changes can exhibit positive autocorrelation without conditioning on contemporaneous information. Finally, new information on the nature of shocks may cause drastic rating changes which seem unwarranted based on the information relevant for identifying current-condition default risk.

The findings thus suggest that the empirical evidence on ratings has to be interpreted with care. Apparent violations of informational efficiency could well result from the agencies' rating method. It is important to evaluate ratings against an appropriate benchmark, and to take their particularities into account when using them as inputs to other models. Depending on the purpose, agency ratings and current-condition ratings may not be interchangeable.

Among the related literature is Carey and Hrycay (2001), whose paper contains an empirical investigation into rating dynamics. They find that agency ratings exhibit less cyclical variation and are more stable than current-condition ratings, which is consistent with agencies following the through-the-cycle approach. Löffler (2002) examines the consequences of the agencies' tendency to reduce rating volatility by taking a rating action only when it is unlikely to be reversed shortly afterwards. Such a

[3] Cf. Hamilton (1994) for a detailed description of the Kalman filter.

[4] This does not mean that correlations across borrowers should be ignored in practice. Exploiting correlations can increase the accuracy with which permanent and transitory components of default risk are separated. In the paper, shocks are assumed to be independent across borrowers, so that separate estimation is efficient.

policy affects rating dynamics in a similar way as following a through-the-cycle approach. It appears that rating through-the-cycle can have stronger effects on rating stability, whereas avoidance of rating reversals can explain predictability of rating changes unaccounted-for by the through-the-cycle approach. Even though their effects are similar, the two rating policies are conceptually different. Avoiding rating reversals by suppressing rating changes works like a filter that leads to a loss of information. Through-the-cycle ratings, too, neglect information, but only in order to convey other information not contained in current-condition ratings.

Krahnen and Weber (2001) propose general standards for good rating practice, focusing on current-condition ratings. My analysis of cyclical components in default risk partly builds on the literature on mean reversion in asset prices (Fama and French, 1988; Poterba and Summers, 1988). The consequences of mean-reverting default risk for the prices of risky debt are studied in Collin-Dufresne and Goldstein (2001); Dangl and Zechner (2001) investigate dynamic capital structure choice and its impact on credit spreads and default risk. Recent papers on the fundamental determinants of ratings and their informational content are Blume et al. (1998) and Ederington and Goh (1998), respectively.

The paper is organized as follows. Section 2 formalizes the through-the-cycle method. Section 3 explores the ensuing rating dynamics in a Monte Carlo framework, and sets them against empirical evidence on agency ratings. Section 4 concludes.

2. Formalizing the rating process

2.1. Distance to default

I examine differences between rating methodologies in a world where ratings reflect an issuer's distance to default within a Merton (1974) type model of default. Default is triggered if the value of a firm's assets falls below a threshold which is related to the firm's liabilities. The distance to default is the standardized difference between a firm's asset value and the default threshold. The Merton model can also be adapted to sovereigns, whose assets are largely made up by the capacity to levy taxes. As will become clear in the course of the paper, the choice of this framework is not critical. In essence, the analysis requires only that ratings are based on some notion of credit quality, and that credit quality is subject to cyclical shocks.

Since the exact modeling of the default process does not seem relevant for the purpose of this paper, I use a simple model of default. A borrower's logarithmic asset value x_t follows a random process whose increments have mean zero and constant variance. Default occurs if the asset value drops below an exogenous default threshold, whose logarithm is denoted by d. This threshold is taken to be constant.

In the Merton (1974) model, asset values follow a random walk. In consequence, the asset value process does not contain cycles, and it is not meaningful to apply a through-the-cycle perspective. To introduce cyclicality, I model asset values as the sum of a random walk x^* and an autoregressive process of order one y:

$$x_t = x_t^* + y_t; \quad x_t^* = x_{t-1}^* + \varepsilon_t, \quad y_t = \rho y_{t-1} + u_t, \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2), \quad u_t \sim N(0, \sigma_u^2), \tag{1}$$

with $0 < \rho < 1$ and $\mathrm{Cov}(\varepsilon_t, u_t) = 0$. Whereas the innovations ε_t are permanent, the u_t are not. They introduce a mean-reverting component into the asset value. Such processes have been used, for example, to describe the behavior of equity returns because equity prices have been found to be subject to both permanent and transitory shocks (see, for example, Campbell et al., 1997). The evidence on negative autocorrelation in equity returns gives one justification why asset values and thus, with a constant default threshold, the distance to default should be mean-reverting. Another possible reason is that borrowers target a certain credit quality level; survey evidence on such behavior is presented in Graham and Harvey (2001). Even if asset values follow a pure random walk, managing credit risk through changes in leverage or other means can introduce mean reversion into the default risk process.

A convenient way of summarizing the extent of cyclicality is the variance ratio, defined as the unconditional T-period variance divided by T times the one-period variance. [5] To derive the variance ratio for the process (1), we need the unconditional variance of the T-period asset value change, which is given by

$$\begin{aligned} \mathrm{VAR}(x_t - x_{t-T}) &= T\,\mathrm{VAR}(\varepsilon_t) + \mathrm{VAR}(y_t) + \mathrm{VAR}(y_{t-T}) - 2\,\mathrm{COV}(y_t, y_{t-T}) \\ &= T\sigma_\varepsilon^2 + 2\frac{\sigma_u^2}{1-\rho^2} - 2\rho^T \frac{\sigma_u^2}{1-\rho^2}. \end{aligned} \tag{2}$$

For a random walk, the variance ratio is unity. With a cyclical process like (1), short term fluctuations tend to be corrected in later periods, which makes the variance ratio smaller than one. The conditional T-period variances of changes in x and y are, by repeated substitution, given by $T\sigma_\varepsilon^2 + \sum_{t=0}^{T-1} \rho^{2t}\sigma_u^2$ and $\sum_{t=0}^{T-1} \rho^{2t}\sigma_u^2$, respectively. In the paper, the annual conditional volatility will be denoted by $\sigma(x)$ and $\sigma(y)$, respectively. Periodicity is set equal to one month.

Following industry practice, current-condition ratings are assumed to be based on the one-year probability of default. Due to mean reversion, the current asset value x_t is not sufficient to characterize the default probability. The conditional expected return on x_t varies with the transitory component y_t; over one year it is equal to $-(1-\rho^{12})y_t$. I thus suggest the following measure for the current-condition distance to default $\mathrm{CC\_DTD}_t$, which is based on the expected one year ahead asset value:

$$\mathrm{CC\_DTD}_t = \frac{E[x_{t+12}] - d}{\sigma(x)} = \frac{x_t - (1-\rho^{12})y_t - d}{\sigma(x)}. \tag{3}$$

In a world where, as in Merton (1974), default can occur only at the end of horizon, ranking borrowers according to (3) is equivalent to ranking them according to their default probabilities. (In this case, the one-year default probabilities would be given by $\Phi[-\mathrm{CC\_DTD}_t]$, with $\Phi[\cdot]$ denoting the standard normal cumulative distribution function.) To make the simulations of Section 3 more realistic, I will assume

[5] Cf. Campbell et al. (1997) for a discussion of variance ratios.

that default can occur also during the one-year horizon. If the asset value exhibits mean reversion, the relationship between distance to default and default probability is then no longer straightforward. [6] However, the measure defined in (3) will provide a good approximation to the true default risk, an assertion that will be underpinned by the results presented in Section 3. Since it will be shown that the current-condition distance to default provides a better measure of default risk than the through-the-cycle concept, any errors in approximating the true default risk through (3) lead to a bias *against* finding the results of the paper.

Note that the conclusions drawn in this paper are not based on a comparison of actual default probabilities with those arising in the model. Therefore, it is not necessary to assign an exact default probability to the distance to default measures used in the analysis. In those instances where it is illustrative to associate a default probability with a given distance to default, I state the probabilities for the case in which asset values follow a random walk and default can arise at any time during the one-year horizon. This default probability obtains as $2\Phi[-\mathrm{CC\_DTD}_t]$. [7]

I am not aware of a clear-cut description of the through-the-cycle method published by the rating agencies themselves. I therefore build on the following characterization by Carey and Hrycay (2001): agencies assign ratings based on an estimate of the borrower's default probability in a stress scenario. This can be modeled by decomposing the default probability $p(D)$ into a conditional and a marginal component: [8]

$$p(D) = p(D|S)p(S), \tag{4}$$

where $p(S)$ is the probability that the stress scenario occurs and $p(D|S)$ is the probability of default in the stress scenario. Through-the-cycle ratings are based on the latter. I complete the characterization of Carey and Hrycay (2001) by defining the stress scenario. A natural definition seems to be that it represents a deviation from the normal condition which is of a purely cyclical nature and which occurs with a certain probability over a predefined horizon. Within the framework of the asset value process introduced above, the normal condition is marked by the permanent asset value x_t^*. Cyclical variations are due to the transitory component y_t. Since changes of y are normally distributed, the deviation from the normal condition is equal to $\Phi^{-1}[p(S)]$ times the volatility of y over the horizon applied for the stress scenario. Note that a borrower's current condition may be worse than the stress scenario. In such a case, relying on the stress scenario would be inconsistent with the notion of stress because the stress scenario would correspond to an improvement rather than a deterioration of credit quality. I therefore assume that agencies assign ratings based on the minimum over the current distance to default and the distance to default in the stress scenario. Formally, if both the stress scenario and the default

[6] Available solutions for the default probability within a first-passage-time model rest on the assumption of Brownian motion (see, for example, Zhou, 2001).

[7] See Zhou (2001).

[8] For simplicity, the formula neglects cases in which borrowers jump to default without passing through the stress scenario.

probability in the stress scenario are defined over a one-year horizon, agency ratings reflect the following stressed distance to default S_DTD: [9]

$$
\begin{aligned}
S\_\mathrm{DTD} \\
&= \min\left\{\frac{x_t^* + \Phi^{-1}[p(S)]\sigma(y_t) - (1-\rho^{12})\Phi^{-1}[p(S)]\sigma(y_t) - d}{\sigma(x_t)}; \frac{x_t - (1-\rho^{12})y_t - d}{\sigma(x_t)}\right\} \\
&= \min\left\{\frac{x_t^* + \rho^{12}\Phi^{-1}[p(S)]\sigma(y_t) - d}{\sigma(x_t)}; \frac{x_t - (1-\rho^{12})y_t - d}{\sigma(x_t)}\right\}. \qquad (5)
\end{aligned}
$$

As in the definition of the current-condition distance to default, I account for the fact that the expected return on x_t varies through the cycle; in the stress scenario, it is given by $-(1-\rho^{12})\Phi^{-1}[p(S)]\sigma(y)$.

Rating agencies publish their credit assessments in the form of discrete grades. In the paper, I nevertheless examine the dynamics of the distance to default instead of those of a rating mapped thereon. Due to the discrete nature of ratings, rating changes would be serially dependent even if the underlying state variable were not (Löffler, 2002). An analysis of categorical ratings would thus obscure patterns produced by the rating methodology. I examine the distance to default rather than the associated default probability because the non-linearity of the latter would complicate the exposition and the presentation of results. Note, too, that some of the analyses are based on an ordinal ranking of borrowers, in which case it is irrelevant whether borrowers are ranked according to rating, distance to default, or default probability. In other cases, the interpretation of the results implicitly assumes that the rules for mapping distances to default into ratings are constant. [10]

At this stage, it is appropriate to check whether the definition is consistent with the agencies' descriptions of their rating methodology. Hilderman (1998) mentions sensitivity to economic cycles as a rating criterion employed by Moody's. In the framework proposed here, sensitivity is measured by $\sigma(y)$, the volatility of the cyclical component. As is evident from (5), the stressed distance to default is lower for larger $\sigma(y)$, even holding the asset volatility $\sigma(x)$ fixed. (Note that ρ is positive while $\Phi^{-1}(p(s))$ is negative.) The definition is thus consistent with cyclical borrowers receiving lower ratings. The following statement by Moody's indicates that the through-the-cycle method is particularly distinct for investment-grade issuers:

> (...) Moody's believes that giving only a modest weight to cyclical conditions best serves the interest of the bulk of investors. Investment-grade issuers presumably possess sufficient financial strength to weather a recession. Consequently, for investment grade issuers in particular,

[9] I have examined an alternative definition according to which through-the-cycle ratings are a mapping of the 'permanent' distance to default, i.e. based on $(x_t^* - d)/\sigma(x)$. Using this alternative definition instead of definition (5) does not change the conclusions drawn from the analyses in Section 3.

[10] While there is evidence for shifts in rating standards (Blume et al., 1998), it is not apparent that they are recurring, or follow predictable patterns. One possible reason for systematic shifts in rating standards could be changes in average credit quality. In the simulations conducted in this paper, average credit quality is constant because innovations are independent across firms.

> Moody's ratings do not automatically change with business cycles. (Cantor and Fons, 1999, pp. 6–7)

As noted by Carey and Hrycay (2001), many non-investment grade issuers are already in a condition of stress, which blurs differences between the agencies' method and the current-condition approach. In definition (5), this is accounted for by the minimum condition which transforms through-the-cycle ratings to current-condition ones. Since the current position in the cycle does not enter the definition other than through the minimum condition, ratings "do not automatically change with business cycles". Reasons for rating changes are mentioned in the following quote from Standard & Poor's (2000):

> Although exogenous factors (such as weakening terms of trade, higher global interest rates, financial distress among external creditors, threat of war) can limit policymakers' degrees of freedom, they usually do not undermine a sovereign's creditstanding and were not motivating factors behind rating actions in 1999. When exogenous events do alter (normally lower) a sovereign's creditstanding, it is because either the policymakers' capacity to respond was originally overestimated or the amplitude of the commodity or interest rate cycle was underestimated (Standard & Poor's, 2000, p. 8).

Overestimating the government's capacity to respond corresponds to underestimating the autocorrelation coefficient ρ, which determines the speed of adjustment to the normal condition. Consistent with this description, the stressed distance to default decreases if ρ is increased. (The numerator in the first term of (5) is decreasing in ρ, while the denominator is increasing in ρ). The amplitude of cycles corresponds to $\sigma(y)$, the volatility of the transitory component. As mentioned above, the stressed distance to default decreases if $\sigma(y)$ goes up.

2.2. Separating permanent and temporary components

Both in the framework of this paper and in reality, [11] raters need to separate transitory and permanent components. Current-condition ratings require an estimate of the transitory component y_t to assess the conditional expected return while through-the-cycle ratings require an estimate of the permanent asset value x_t^*. If the current asset value x_t is the only observable variable, the standard solution to these problems is to apply a Kalman filter which produces a prediction of x_t^* and y_t based on the observed x_t. For the process used here, the intuition behind the statistical procedure is as follows. If the observed value x_t deviates from its unconditional mean, deviations are attributed to both permanent and transitory shocks. The precise split is determined by the relative variance of accumulated permanent and transitory shocks. The longer deviations persist, and the lower the variance of transitory shocks is,

[11] "The greatest challenges facing analysts in rating a cyclical company are to capture its equilibrium and distinguish fundamental changes in credit quality from cyclical trends" (Fitch IBCA, 1999, p. 1).

the less likely are deviations to be transitory. Formally, the dynamics of the observed variable x_t can be represented through the following system of equations (the exposition closely follows Hamilton, 1994):

$$\text{State equations:} \quad \begin{bmatrix} x_t^* \\ y_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix} \begin{bmatrix} x_{t-1}^* \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ u_t \end{bmatrix}, \tag{6}$$

$$\text{Observation equation:} \quad x_t = [1 \quad 1] \begin{bmatrix} x_t^* \\ y_t \end{bmatrix}. \tag{7}$$

Now define the following matrices:

$$\xi_t = \begin{bmatrix} x_t^* \\ y_t \end{bmatrix}; \quad F = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}; \quad Q = \begin{bmatrix} \sigma_\varepsilon^2 & 0 \\ 0 & \sigma_u^2 \end{bmatrix}; \quad H = \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \tag{8}$$

The Kalman filter begins with $\hat{\xi}_{1|0}$, a forecast of ξ_1. If there are no observations on x prior to $t = 1$, $\hat{\xi}_{1|0}$ will be the unconditional mean of ξ_1. Having observed x_t, the estimate is updated with the following formula:

$$\hat{\xi}_{t+1|t} = F\hat{\xi}_{t|t-1} + FP_{t|t-1}H(H'P_{t|t-1}H)^{-1}(x_t - H'\hat{\xi}_{t|t-1}), \tag{9}$$

where $P_{t|t-1}$ is the mean-squared error of $\hat{\xi}_{t|t-1}$. If, for example, the initial state of (6) is known to the analyst, the first period mean-squared error will be equal to Q:

$$P_{1|0} = \begin{bmatrix} \sigma_\varepsilon^2 & 0 \\ 0 & \sigma_u^2 \end{bmatrix}. \tag{10}$$

Before repeating the update in $t + 1$, one has to compute the mean-squared error of the updated prediction:

$$P_{t+1|t} = F(P_{t|t-1} - P_{t|t-1}H(H'P_{t|t-1}H)^{-1}H'P_{t|t-1})F' + Q. \tag{11}$$

The matrix $P_{t|t-1}$ is not a function of the realizations before t, but determined solely by the parameters of the process. Finally, note from (9) that unexpected changes in the predictions $\hat{\xi}$ have a variance of

$$\mathrm{Var}(\hat{\xi}_{t+1|t} - F\hat{\xi}_{t|t-1}) = \mathrm{Var}(FP_{t+1|t}H(H'P_{t+1|t}H)^{-1}(x_t - H'\hat{\xi}_{t|t-1})). \tag{12}$$

Since $H'\hat{\xi}_{t|t-1}$ is the forecast of x_t made in $t - 1$, unexpected changes in the Kalman estimates are solely driven by contemporaneous innovations in the observed variable.

2.3. Parameterizing the model

The parameters of the state equation are determined by choosing values for the autocorrelation parameter ρ, the conditional volatility of the observed asset value x_t, and the variance ratio of T-month asset returns relative to one-month asset returns. This gives two equations (for the volatility and for the T-month variance ratio) with two unknowns (σ_ε^2 and σ_u^2), which are solved numerically.

In the remainder of the paper, I will illustrate the effects of rating through the cycle with the help of examples. In these examples, the annual conditional volatility of x_t is set to 0.15. In fact, the precise value of this volatility does not influence the results because the analyses focus on changes in the distance to default rather than on absolute default probabilities. The autocorrelation parameter ρ is set to 0.98 in most cases, but variations will consider values of 0.99 and 0.96. The resulting half-life of transitory shocks is approximately 2.9, 5.7 and 1.4 years for ρ equal to 0.98, 0.99 and 0.96, respectively. In their analysis of the dynamics of book leverage, Fama and French (2002) estimate the annual adjustment to a target leverage ratio to be 10% per annum for dividend paying firms and 18% for non-payers. With the framework of this paper, a given annual rate of adjustment can be produced by setting the monthly autocorrelation parameter ρ equal to $(1 - \text{annual rate of adjustment})^{1/12}$. The values obtained by Fama and French thus correspond to ρ equal to 0.991 and 0.984, respectively. As some characteristics of through-the-cycle ratings are more pronounced if the pace of mean reversion is slow, the value of 0.98 is chosen to be at the lower end of the empirical evidence to avoid overestimating the effects of rating through the cycle.

For the variance ratio I choose four different values based on the five-year variance. The associated ratio is set to 0.6, 0.7, 0.8 or 0.9. I am not aware of empirical estimates of the extent of mean reversion in the distance to default. For a sample of US stocks, Poterba and Summers (1988) obtained an average five-year variance ratio of 0.89 for individual excess equity returns. Since the standard error of this estimate is fairly large (0.20) and, more importantly, the distance to default can be managed by the firm, the chosen values for the five-year variance ratio seem to be representative. In the analyses of this paper, I will mostly assume that rating agencies know the structure and parameters of the system (Eqs. (6)–(8)). [12] As to $p(S)$, the probability that the stress scenario occurs, there is little information about which value provides a representative description of the agencies' rating method. [13] I will set it equal to 0.05 in most cases, considering alternative values of 0.2 and 0.01. Both the stress scenario and the distance to default in the stress scenario are defined over a one-year horizon. Note that, from Eq. (5), the effects of applying a different horizon for the stress scenario can be reproduced by applying a different $p(S)$.

To illustrate the effects of rating through the cycle, I resort to Monte Carlo simulations because the minimum condition in (5) makes it difficult to derive analytic results. To mimic the real world, in which borrowers can be at different positions in the cycle, and rating agencies have typically only imperfect knowledge on these positions, I choose the following set-up for the simulations: The initial value of the transitory component y is drawn from its unconditional distribution, which is normal with mean zero and variance $\sigma_u^2/(1-\rho^2)$. The initial value of y is assumed to be known to the raters; uncertainty about the position in the cycle is introduced

[12] Unreported analyses show that conclusions do not change if the true variance ratio is set to 0.75, but raters estimate it to be 0.6, 0.7, 0.8 or 0.9, each with probability 0.25.

[13] I am only aware of the following statement by Fitch IBCA (1999, p. 2): "Fitch IBCA's ratings reflect the fundamentals of the company in the midpoint of its downturn and recovery."

Table 1
Parameter specifications

Model parameters	Chosen parameter values (alternative values separated by commas)
Variance ratio of five-year asset returns	0.6, 0.7, 0.8, 0.9, 1.0
Autocorrelation parameter ρ	0.96, 0.98, 0.99
Annual conditional asset volatility $\sigma(x)$	0.15
Variance of permanent shocks σ_ε^2	Follows from variance ratio, ρ and $\sigma(x)$
Variance of transitory shocks σ_u^2	Follows from variance ratio, ρ and $\sigma(x)$
Probability of stress scenario $p(S)$	0.01, 0.05[a], 0.2
Initial value of the cyclical component y	Drawn from unconditional distribution
Default threshold d	0
Initial Kalman prediction error	0

[a] Parameter chosen in the base case.

by starting the analysis of rating dynamics in the fourth year of the simulated data set. [14] The set-up thus corresponds to a situation in which raters knew the state of the cycle three years ago, but then had to rely on the Kalman filter predictions.

The magnitude of the prediction errors varies with the parameters chosen for the state equation. With a five-year variance ratio of 0.8 and an autocorrelation parameter of 0.98, for example, the mean-squared error of the Kalman prediction made for y at the beginning of year four is 0.012, which is almost 50% of the unconditional variance of y (0.025).

The description of the framework for the simulation experiments is now complete. In the presentation of results, I will refer to the *base case* as comprising the following assumptions: (i) the initial value of y_t is drawn from its unconditional distribution, and known to the rating agency; (ii) the analysis starts in the fourth year of each simulated asset value path; (iii) the probability for a stress scenario $p(S)$ is 0.05; (iv) the parameters of the asset value process are known to the rating agency; (v) the default threshold d is zero. Table 1 provides a summary of the parameter specifications.

For the fourth year of the simulated data, Table 2 lists the conditional volatilities of the observed asset value, the permanent asset value, and the estimated asset value in the base case. The volatility of the Kalman prediction for the permanent asset value can be obtained by summing up the one-period variances given in (12). Intriguingly, the volatility of the estimated permanent asset value is lower than the volatility of the permanent asset value itself. To see why this can be so, consider a period in which the transitory component does not change, while the permanent does. The Kalman filter will classify a part of this permanent shock as transitory, and the change in the estimated permanent asset value will be smaller than the actual change. Another way of explaining the picture is that permanent and transitory components of the asset value are independent, whereas their estimates are, from Eq. (9), positively correlated. In addition, the sum of the estimates is equal to the sum of the true values. $\mathrm{Var}[\hat{x}_t^* + \hat{y}_t]$ plus a positive covariance term is therefore equal to

[14] Conclusions do not change if the analysis starts at the beginning of year two.

Table 2
Conditional asset volatilities

Variance ratio of five-year asset returns	Autocorrelation parameter ρ	Volatility		
		Observed asset value x_t	Permanent asset value x_t^*	Estimated permanent asset value $\hat{x}_t^*$
1	–	0.150	0.150	0.150
0.9	0.98	0.150	0.134	0.130
0.8	0.98	0.150	0.114	0.103
0.7	0.98	0.150	0.085	0.066
0.6	0.98	0.150	0.031	0.010
0.8	0.99	0.150	0.067	0.038
0.8	0.96	0.150	0.131	0.129

$\mathrm{Var}[x_t^* + y_t] = \mathrm{Var}[x_t^*] + \mathrm{Var}[y_t]$. The relation can only hold if the variances of the estimates are smaller than the variances of the underlying variables.

3. Rating through the cycle as an explanation of stylized facts

3.1. The default prediction power of ratings is low

Agency ratings are often used to infer individual obligor default probabilities. Examples are portfolio credit risk models such as CreditMetrics (Gupton et al., 1997) or the recent proposal of the Basel Committee on Banking Supervision (2001). Proponents of alternative rating methods point out that the predictive quality of agency ratings can be substantially improved (cf. Kealhofer et al., 1998). The fact that Moody's has developed a statistical model for predicting short term default risk, which complements the traditional through-the-cycle rating, indicates that even the rating agencies share this view (cf. Sobehart et al., 2000).

By definition, through-the-cycle ratings are not based on the current default probability. This alone does not question their use for assessing default risk because the correlation between through-the-cycle ratings and current-condition default probabilities could still be very high. One can indeed construct simple examples in which the correlation is perfect. Consider M borrowers, which in period $t = 0$ are identical in every respect. Now simulate individual asset value changes, which may be cross-sectionally dependent or not. If, in the next period, borrowers are ordered according to their default risk, current-condition and stressed distance to default produce identical orderings. The reason is that the estimates $\hat{x}^*$ and $\hat{y}^*$, which are alone responsible for differences in default risk, are monotonically related to the observed asset value x (see Eq. (9)).

In reality, the parameters of the asset value process will not be identical across borrowers, and borrowers will be situated at different points in the cycle. In consequence, through-the-cycle ratings will lose discriminatory power. While the default probability depends on the current position in the cycle, through-the-cycle ratings are only affected by sensitivity to cycles. In addition, once the extent of cyclicality differs between borrowers, the Kalman predictions of the permanent asset value will differ even if the borrowers are initially at the same position in the cycle, and then experience an identical change in their observed asset values. Consider two borrowers whose variance ratios are 0.7 and 0.9, respectively. If the asset values of both borrowers increase by 10%, the increase in the estimated permanent asset value will be smaller for the borrower with the smaller variance ratio. This effect is also relevant for the current-condition distance to default because it includes an estimate of the cyclical component y. The magnitude of the effect, however, is smaller. From Eq. (3), estimates of the transitory component y_t are scaled down by $(1-\rho^{12})$ when entering the definition of the distance to default. For $\rho = 0.98$, this scaling factor is 0.22.

To gauge the magnitude of these effects, and to compare it to existing evidence on the predictive quality of agency ratings, I perform the following experiment. I look at a sample of 50,000 borrowers. In $t = 0$, the asset value of each borrower is drawn from a uniform distribution over the interval [0.15, 0.55]. For a variance ratio of one, the continuous time default probabilities associated with these asset values would be 31.73% ($= 2\Phi[-0.15/0.15]$) and 0.025%, respectively. The range thus roughly covers the default rates of issuers rated AAA to CCC by Standard & Poor's. As in the base case, the initial state of the cycle y_0 is drawn from its unconditional distribution. To account for the fact that the extent of cyclicality will differ across borrowers in practice, I randomly assign each borrower a variance ratio of 0.6, 0.7, 0.8 or 0.9. Based on these initial conditions, asset values are simulated until the end of year four. [15] Measures of the distance to default are determined using the Kalman filter estimates of x^* and y. I store the values of the current-condition and the stressed distance to default from the end of year three, as well as a variable indicating whether a default has occurred in year four; a default is taken to occur when at least one of the monthly asset values simulated for year four is below zero. [16]

The simulation is repeated 50,000 times. For the analysis, I take only those runs where the asset value at the end of year three lies in the interval [0.15, 0.55]. Again, this restriction is meant to bring the simulated sample in line with the data sets used in empirical studies of rating performance. Following Sobehart et al. (2000), the predictive quality of ratings is measured through power curves, which are constructed as follows. For a given rating system, borrowers are sorted according to their distance to default. If the latter predicts defaults, a large fraction of defaults will occur among the borrowers with a small distance to default. This relation can be checked by examining the percentage of defaulters whose distance to default is lower than the α quantile of the population.

[15] The impact of the variance ratio can be illustrated through the rank correlation coefficients between current-condition and stressed distances to default at the end of year three. They amount to 0.73 (variance ratio = 0.6), 0.85 (0.7), 0.94 (0.8), and 0.98 (0.9).

[16] The discreteness of the simulation will lead to default rates which are lower than the ones that would be observed in a continuous time setting. On the other hand, a continuous time analysis is unlikely to provide a perfect description of a world in which, for example, payment obligations are due only in discrete intervals.

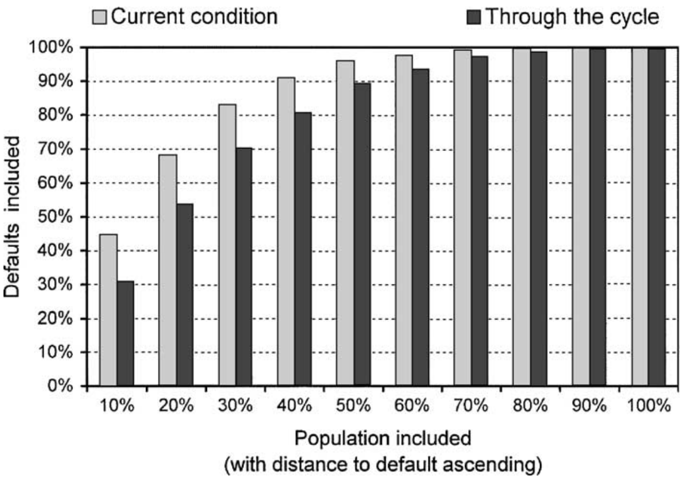

Fig. 1. Simulated power curves for current-condition and through-the-cycle distance to default.

Fig. 1 contains the simulated power curves for the current-condition and the stressed distance to default. The current-condition approach is clearly superior in predicting defaults. 68% of all defaults occur among the 20% of all borrowers with the lowest current-condition distance to default; borrowers with a stressed distance to default lower than the 20% quantile, by contrast, contain only 54% of all defaulted borrowers. Kealhofer (2000) presents the results of a similar, empirical study that compares the performance of S&P ratings with KMV default predictions. The KMV default rate predictions, which are based on the Merton (1974) model, are essentially current-condition. The figures in Kealhofer (2000) closely resemble the ones obtained here. The 20% of all borrowers with the lowest KMV (S&P) rating contain 72% (61%) of all defaults.

To assess the impact of assumptions on the relative predictive power of internal and external ratings, it is useful to condense the information contained in power curves into a single figure. Sobehart et al. (2000) propose the following accuracy ratio: for a given rating system, determine the area above the diagonal, and relate it to the maximum area that can be enclosed above the diagonal. This maximum would be achieved by a rating system in which no defaulting firm gets a better rating than a non-defaulting firm. The accuracy ratio thus lies between zero and one, and larger values indicate better predictive power. Sobehart et al. (2000) compute empirical accuracy ratios for various quantitative default prediction models and obtain figures between 0.43 and 0.73.

Table 3 contains the accuracy for the assumptions described above as well as for various modifications. They confirm the conclusion that current-condition ratings are superior in predicting defaults. As should be expected, the advantage is smaller if there is less cyclicality (the variance ratio is set to 0.8 or 0.9 instead of being 0.6, 0.7, 0.8 or 0.9), or cyclical shocks have a shorter half-life (the autocorrelation parameter is set to 0.96 instead of 0.98). The choice of the stress scenario probability $p(S)$

Table 3
Simulated accuracy ratios[a]

Assumptions					Simulated accuracy ratios	
Variance ratios	Autocorrelation ρ	Asset value	Stressed default probability	Probability of stress scenario	Current condition	Through-the-cycle
Assumptions as in Fig. 1						
0.6–0.9	0.98	0.15–0.55	Unrestricted	0.05	0.72	0.54
Variations (bold face)						
0.8–0.9	0.98	0.15–0.55	Unrestricted	0.05	0.71	0.60
0.6–0.9	**0.96**	0.15–0.55	Unrestricted	0.05	0.71	0.66
0.6–0.9	0.98	0.15–0.55	Unrestricted	**0.01**	0.72	0.51
0.6–0.9	0.98	0.15–0.55	**∼Investment grade**	0.05	0.63	0.36
0.6–0.9	0.98	0.15–0.55	**∼Speculative grade**	0.05	0.64	0.40

[a] Having arranged default and rating data in a power curve (cf. Fig. 1) accuracy ratios relate the area above the diagonal to the maximum area that can be enclosed above the diagonal. The larger the accuracy ratio, the better the predictive power of a rating system.

has little influence on the accuracy ratio of through-the-cycle ratings. Finally, I assess the default prediction power of through-the-cycle ratings separately for borrowers with a large or low stressed distance to default at the end of year three. [17] The cut-off value is taken to be the stressed distance to default that obtains for a borrower with a permanent asset value of 0.45 and a variance ratio of 0.75, i.e., the average across the four ratios chosen for the simulation. With an asset value of 0.45, the continuous-time default probability is 0.27%. The cutoff value should thus provide a reasonable split into borrowers that would receive agency ratings in the investment grade and speculative grade domain, respectively. [18] The analysis does not reveal any striking differences between the two domains, which does not mean that one should not observe such differences in practice. Contrary to the assumptions made here, borrowers with a relatively low permanent asset value could tend to be more vulnerable to cycles, e.g. because market imperfections reduce the ability or willingness to respond to adverse shocks if financial distress is imminent. In this case, the predictive advantage of current-condition ratings would be higher for low-quality borrowers.

I also introduce cross-sectional correlation in asset value changes (not reported in Table 3). For groups of 2500 borrowers, the correlation of both permanent and transitory shocks is set to 20%; the inter-group correlation is zero. The simulation thus produces data sets similar to a 20-year rating history; it is repeated 100 times. Even in

[17] As before, the asset value at the end of year three has to be in the interval [0.15, 0.55] for inclusion in the analysis.

[18] The average historical default rate of issuers rated BBB is 0.27% (Standard & Poor's, 2002).

the simulation where the through-the-cycle approach performs best relative to the current-condition one, the former is clearly inferior. [19]

Since the simulated figures closely mirror empirical ones, the latter cannot be interpreted as evidence of informational inefficiency. Even if rating agencies efficiently use available information and incorporate it timely into ratings, ratings will not optimally predict defaults. If one aims at obtaining estimates of current-condition default probabilities, agency ratings could, through their architecture, be inferior to other rating systems.

3.2. *Ratings are relatively stable*

Kealhofer et al. (1998) and Carey and Hrycay (2001) find that agency ratings exhibit a much larger stability than current-condition ratings. Kealhofer, Kwok and Weng derive current-condition ratings from an application of the Merton (1974) model and compare their stability with the one of S&P ratings. Carey and Hrycay use a logit model to categorize issuers rated by Moody's. Typically, 40–50% of current-condition ratings remain stable over a one-year horizon, compared to 80–90% in the case of agency ratings. Carey and Hrycay attribute this discrepancy to the agencies' rating methodology, but they do not examine whether the potential effects of rating through the cycle are large enough to account for the evidence. An alternative explanation could be that agencies consistently underreact to new information.

The following analysis quantifies the effects of rating through the cycle. As above, the initial asset value of each borrower is drawn from a uniform distribution over the interval [0.15, 0.55]; the initial state of the cycle y_0 is drawn from its unconditional distribution. Since the results depend heavily on the nature of cyclicality, I conduct the simulation experiment separately for different assumptions about variance ratios and autocorrelation coefficients. Based on 50,000 four year asset value paths for parameterization i, stability is assessed as follows. For borrowers not defaulted until year three, I determine α_i such that 50% of all borrowers whose current-condition distance to default lies within the α_i, and $(1-\alpha_i)$ quantiles at the end of year three are situated in the same range at the end of year four. [20] In the next step, I examine which fraction of borrowers fulfil this condition if the same α_i is used to group borrowers according to the stressed distance to default. A 50% stability of current-condition ratings is thus used as a benchmark, and rating grades under the two architectures are made comparable through the fraction of borrowers they comprise.

Table 4 summarizes the simulated one-year transition probabilities. Generally, through-the-cycle ratings are more stable than current-condition ratings. With a variance ratio of 0.7 and an autocorrelation parameter $\rho = 0.98$, for example, rating stability increases from 50% to 70.3% when moving from a current-condition rating

[19] In this case, the 20% of borrowers with the lowest current-condition (stressed) distance to default contain 66% (53%) of all defaulters.

[20] The condition is not met if a borrower defaults within year four.

Table 4
Simulated one-year rating stability of through-the-cycle ratings if the stability of current-condition ratings is 50%[a]

Assumptions for simulations		Simulated rating stability
Variance ratio	Autocorrelation parameter ρ	
1	–	0.500
0.9	0.98	0.526
0.8	0.98	0.592
0.7	0.98	0.703
0.6	0.98	0.877
0.8	0.99	0.782
0.8	0.96	0.514

[a] Simulations are initialized by drawing the cyclical component from its unconditional distribution. Rating stability is the probability of no rating changing in year four; the width of the rating grade is the same as the one that makes the stability of current-condition ratings equal 50%. For the stressed distance to default, the probability that the stress scenario occurs is set to 5%.

to a through-the-cycle approach. The differences are more pronounced if the variance ratio is smaller, or the autocorrelation coefficient ρ is larger. With a variance ratio of 0.6 or an autocorrelation parameter of 0.99, the magnitude of the differences is close to the ones documented in empirical studies. The findings are robust with respect to $p(S)$, the probability that the stress scenario occurs. If $p(S)$ is taken to be 0.01 or 0.2 rather than 0.05, the stability of through-the-cycle ratings remains largely unchanged. With a variance ratio of 0.7 and $\rho = 0.98$, it amounts to 71.1% ($p(S) = 0.01$) and 66.3% ($p(S) = 0.2$).

There are two explanations for the stability of through-the-cycle ratings. In the current-condition approach, the volatility of the state variable is larger because it is subject to both permanent and transitory shocks. Through-the-cycle ratings, by contrast, are not affected by the current position in the cycle. The effect is corroborated by the fact that the permanent asset value is not known to the rating agency, but has to be estimated. As shown above (see Table 2) the variance of Kalman filter estimates can be significantly smaller than the variance of the predicted variable. Since rating changes are triggered by a change in the estimated permanent asset value, this leads to a further increase in stability when rating through the cycle. To assess the magnitude of this effect, I calculated the rating stability under the assumption that the permanent asset value is observable. With a variance ratio of 0.7 and $\rho = 0.98$, stability is 63.3% instead of 70.3%.

Contrary to what is observed empirically, the simulations do not produce a rating stability above 90%. A possible explanation for this discrepancy is that the through-the-cycle approach is not the only mechanism that reduces the volatility of agency ratings. Rating agencies suppress rating changes when they are likely "to be reversed within a relatively short period of time" (Cantor, 2001, p. 175). Though related, the two features are distinct. The latter is primarily due to the discrete nature of ratings (see Löffler, 2002), which is not modeled in this paper.

3.3. Rating changes are predictable

Empirical studies of rating changes have documented a significant positive autocorrelation (Altman and Kao, 1992; Lando and Skødeberg, 2002). A partial explanation for this phenomenon could be that rating agencies "dole out the bad news in small doses rather than savaging the bond issuer – who is, after all, their customer – all in one go" (Economist, December 13, 1997, p. 70). Serial correlation might also be due to horizon effects. Consider a firm which gradually expands into a new, risky business segment, repeatedly issuing new debt to finance necessary investments. Over time, the default probability will rise. Even if the rating analyst perfectly predicts this development, she will not completely incorporate it in the current rating if the rating horizon is shorter than the time span in which the firm's restructuring is completed. Rating changes will then exhibit positive autocorrelation.

Both effects are at work regardless of the rating methodology chosen, even though their magnitude is likely to be different. Gradualism in negative rating changes is likely to be less visible to uninformed outsiders when rating through the cycle. As the stressed distance to default can only be estimated with noise, not reacting to publicly available negative information can be justified by classifying the shock as transitory.

Since the magnitude of such effects is difficult to assess, I turn to the question of whether the through-the-cycle approach itself can produce serially correlated rating changes. If (i) the Kalman filter is used to infer the stressed distance to default, (ii) the parameters are specified correctly, and (iii) innovations are serially independent and identically distributed, the answer is no within the framework of this paper. The Kalman filter yields the least-squares estimator of the permanent asset value. It is thus not possible to explain future changes in the estimate using past data. If it were, the Kalman estimate would fail to be efficient. This non-predictability result continues to hold if the information set used for testing predictability is expanded to include all observable information; such tests have been performed by Delianedis and Geske (1999), who estimate default probabilities based on borrower fundamentals, and show that rating changes lag changes in the estimated default risk. In the model, the only observable variable is the asset value itself; due to the efficiency of the Kalman filter, this variable cannot be used to predict future changes in the Kalman estimates.

This statement holds for the Kalman estimate as such, and under the assumptions made in the base case. It requires several modifications.

3.3.1. The current-condition component in through-the-cycle ratings

Due to mean reversion, changes in the observed asset value are negatively autocorrelated, and so are changes in the current-condition distance to default, which is driven by the observed asset value. [21] Since the stressed distance to default can,

[21] The current-condition distance to default accounts only for the one-year effects of mean reversion, not the long-term ones.

through the minimum condition modeled in (5), turn into a current-condition one, one might therefore detect negative serial correlation in a sample of through-the-cycle ratings. With the chosen parameterizations this effect is visible, but small. Using a sample simulated within the base case setting (10,000 independent trials), I perform the following regression:

$$S\_\mathrm{DTD}_{48} - S\_\mathrm{DTD}_{36} = \alpha + \beta(S\_\mathrm{DTD}_{36} - S\_\mathrm{DTD}_{24}) + \omega. \tag{13}$$

One-year changes in the stressed distance to default are regressed on the lagged one-year change in the stressed distance to default. With $\rho = 0.98$, a variance ratio of 0.8 leads to a β coefficient of −0.007, statistically insignificant (t-value = −0.70) despite the large sample size (10,000); choosing a variance ratio of 0.7 produces a coefficient of −0.03 (t-value = −2.94).

3.3.2. Dependencies beyond autocorrelation

Another modification extends the traditional notion of predictability. If changes in the distance to default are regressed on contemporaneous observables and the lagged rating changes, the lagged rating changes might be significant even if there is no simple autocorrelation. The intuition for this pattern is as follows: if the asset value experiences a permanent shock in one period, and stays constant at the new level in subsequent periods, the Kalman filter will classify the shock only gradually as permanent. Fig. 2 shows such an adjustment path. With a variance-ratio of 0.7 and $\rho = 0.98$ it takes more than four years until 50% of the shock are classified as permanent. Empirically, the gradual processing of a permanent shock can be revealed by controlling for new information as is done in a regression.

Such dependencies are also evident from an inspection of the Kalman equations. First differencing Eq. (9) and using the shortcut $K_t = FP_{t|t-1}H(H'P_{t|t-1}H)^{-1}$ yields

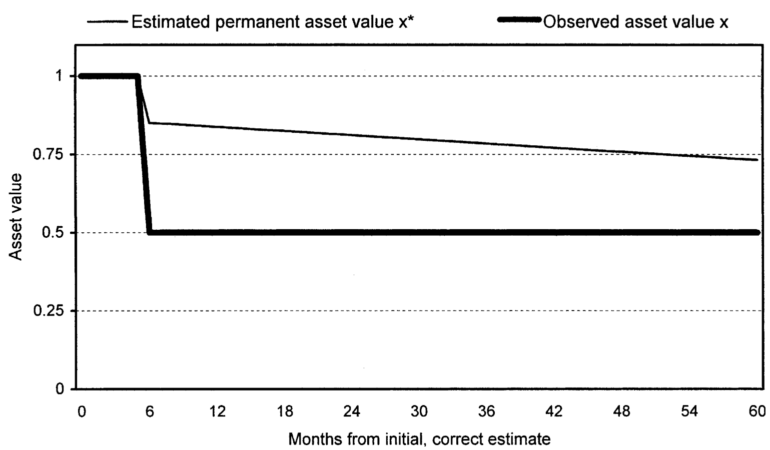

Fig. 2. Kalman estimate of the permanent asset value following a one-time shock to the observed asset value (variance ratio = 0.7, autocorrelation parameter $\rho = 0.98$).

$$\begin{aligned}\hat{\xi}_{t+1|t} - \hat{\xi}_{t|t-1} &= F\hat{\xi}_{t|t-1} + K_t(x_t - H'\hat{\xi}_{t|t-1}) - F\hat{\xi}_{t-1|t-2} - K_{t-1}(x_{t-1} - H'\hat{\xi}_{t-1|t-2}) \\ &= K_t x_t - K_{t-1}x_{t-1} + (F - K_t H')\hat{\xi}_{t|t-1} - (F - K_{t-1}H')\hat{\xi}_{t-1|t-2}. \qquad (14)\end{aligned}$$

As can be seen from Eq. (14), current changes in the Kalman prediction depend not only on changes in the observed asset value, but also on the lagged Kalman predictions.

I conduct Monte Carlo simulations to assess the explanatory power of lagged changes in the stressed distance to default. I simulate 10,000 independent asset value paths, setting the variance ratio to 0.6, 0.7, 0.8, or 0.9, respectively. With the data from these random samples, I perform the following regressions:

$$\begin{aligned}S\_\mathrm{DTD}_{48} - S\_\mathrm{DTD}_{36} &= \alpha + \beta_1(\mathrm{CC\_DTD}_{48} - \mathrm{CC\_DTD}_{36}) \\ &\quad + \beta_2(S\_\mathrm{DTD}_{36} - S\_\mathrm{DTD}_{24}) + \omega. \qquad (15)\end{aligned}$$

One-year changes in the stressed distance to default are regressed on the contemporaneous change in the current-condition distance to default and the lagged one-year change in the stressed distance to default. Abstracting from the fact that agencies report only ratings, not the underlying distance to default, this is a regression an outside observer of rating changes might perform in order to test whether the rating agency, beside incorporating new information, also makes up for information previously neglected. Panel A of Table 5 reports the estimated coefficients. The as-

Table 5
Regression analysis of simulated changes in the stressed distance to default[a]

Assumptions for simulations		Regression results		
Variance ratio of five-year asset returns	Autocorrelation parameter ρ	$\hat{\beta}_1$	$\hat{\beta}_2$	R^2
Panel A: all observations				
0.9	0.98	0.883	0.022	0.995
0.8	0.98	0.732	0.047	0.984
0.7	0.98	0.526	0.058	0.921
0.6	0.98	0.211	−0.121	0.395
0.8	0.99	0.419	−0.028	0.712
0.8	0.96	0.894	0.043	0.997
Panel B: only observations with stressed distance to default < current-condition distance to default				
0.9	0.98	0.881	0.026	0.996
0.8	0.98	0.718	0.059	0.991
0.7	0.98	0.477	0.106	0.979
0.6	0.98	0.082	0.169	0.957
0.8	0.99	0.268	0.067	0.981
0.8	0.96	0.894	0.044	0.997

[a] Using a simulated sample, one-year changes in the stressed distance to default are regressed on the contemporaneous change in the current-condition distance to default (β_1), and the lagged one-year change in stressed distance to default (β_2). For the stressed distance to default, the probability that the stress scenario occurs is set to 5%.

sumed autocorrelation parameter ρ is 0.98. Simulation errors are negligible as the t-statistics are mostly above 20.

The estimated coefficients $\hat{\beta}_2$ show that changes in the stressed distance to default are explained by lagged values of the same variable provided new information is controlled for. With a variance ratio of 0.7, for example, a 10% change in the stressed distance to default produces a 0.58% change in the following year, ceteris paribus. For some of the parameter combinations, the effect is negative. This is due to situations in which the stressed distance to default turns into a current-condition one. I therefore run the regressions (15) separately for those observations where the minimum condition in (5) does not bite at any of the dates $t = 24$, 36 and 48. The results are also reported in Table 5, Panel B. The influence of the lagged stressed distance is generally larger, and always positive. With a variance ratio of 0.6, a 10% change in the stressed distance to default leads to a 1.69% change in the following year.

It is difficult to gauge the magnitude of the effects, especially because there are no comparable empirical studies. Even if the effects were too small to become visible in a statistical analysis of rating dynamics, one should still be careful when interpreting anecdotal evidence. An observation like "The agencies are continuing to downgrade Japanese financial institutions for reasons that have been widely commented upon for months, if not years" (Economist, December 13, 1997, p. 71), is not necessarily evidence of informational inefficiency. If a downgrade is followed by another downgrade, it can well be that the two changes go back to a single shock. The second downgrade could occur because the rating agency learns that it underestimated the persistence of the shock. In the presence of uncertainty about the nature of shocks, such learning is not in itself a sign of informational efficiency. Also, in order for such learning to occur, it is not necessary that the initial shock is corroborated by new developments pointing in the same direction. Even if overall conditions do not change after the shock, it will increasingly be classified as permanent. One single event can underlie a series of rating changes, and rating agencies explaining their rating decisions might well cite the same reasons again and again.

3.3.3. Estimation errors

So far, I have assumed that rating agencies know the parameters of the asset value process. This assumption is likely to be unrealistic. Estimates of long-term dynamics are imprecise even if the relevant variables are observed over a period of 50 years or more (see Campbell et al., 1997). The estimation problem is aggravated because parameters will typically not be stationary. A change in a firm's financing strategy or its business activities can change both the speed of adjustment and the magnitude of cyclical shocks.

The likely magnitude of estimation errors is difficult to assess. Thus, I only present an example which illustrates their potential effects. I examine a case in which the rating agency misestimates the average degree of cyclicality. Due to cross-sectional correlation of shocks as well as changes in corporate financial strategy, estimation errors are unlikely to cancel out across borrowers. Assume that the rating agency uses an aggregate time series, say a stock market index, to estimate the five-year-variance ratio of asset values. If the data covers 50 years of monthly observations, and the rating agency employs the estimator in Campbell et al. (1997, p. 52), the asymptotic standard error of the estimate is 0.36.

Now assume that the true five-year variance ratio is 0.9, whereas the agency has estimated it to be 0.7. The estimate is less than one standard error away from the true value, but the error can have significant consequences. For a sample of 10,000 issuers, I simulate independent four-year asset value paths with a variance-ratio of 0.9. The stressed distance to default is estimated through the Kalman filter, assuming the five-year variance ratio to be 0.7. With the simulated data, I run the following regression:

$$S\_\mathrm{DTD}_{48} - S\_\mathrm{DTD}_{36} = \alpha + \beta(S\_\mathrm{DTD}_{36} - S\_\mathrm{DTD}_{24}) + \omega. \tag{16}$$

The estimated coefficient β is 0.06 (t-value = 6.05), revealing positive autocorrelation. In the example, the rating agency underestimates the magnitude of permanent shocks. On average, deviations will persist longer than expected, and the agency will gradually adjust its estimates in the same direction as previous changes. [22] Regressing the stressed distance to default on the lagged current-condition distance to default produces similar results.

The results are potentially relevant for the design and interpretation of statistical tests. Consider the study by Delianedis and Geske (1999). Using 10 years of data from 1987 to 1996, the authors show that rating changes lag changes in estimated current-condition default probabilities. Their significance tests assume independent observations. The validity of such an assumption should be checked. Rating agencies who efficiently use available information could still misestimate the average degree of cyclicality over a period of several years, leading to predictability in rating changes. Note that similar problems arise if the rationality of macroeconomic forecasts is tested for a small number of years, and errors in the forecasters' models cannot be ruled out (see, for example, Batchelor and Dua, 1991).

3.4. Some rating changes are unrelated to new information

While rating agencies are typically criticized for being predictable and conservative in their actions, sometimes rating changes are also regarded as unwarranted or excessive. In 1997 for example, the sharp downgrade of several east Asian countries puzzled some observers (cf. IMF, 1999, S.207). Within the framework of this paper, adjustments that are seemingly unrelated to news can be explained as follows: there may be information events that have only a small impact on current-condition default risk, but that are highly relevant for separating permanent and transitory components. If the new information necessitates a change in the estimate of permanent components, agency ratings may react strongly even though the current-condition default risk, on which outside observers often focus, remains largely unaffected.

Such information can work in two ways. First, it can cause raters to alter their estimates of the parameters of the asset value process. A firm's announcement that

[22] If the agency overestimated the magnitude of permanent shocks, serial correlation would be negative.

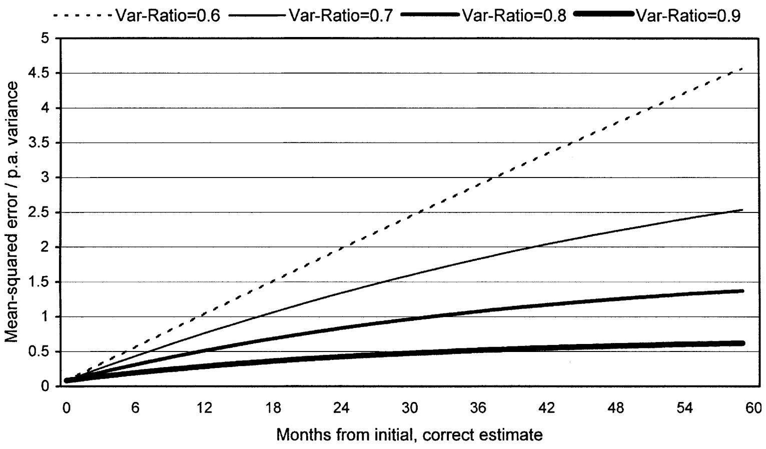

Fig. 3. Mean-squared error of the estimated permanent asset value relative to the one-year variance of the permanent asset value (autocorrelation parameter $\rho = 0.98$).

it will take longer than expected to bring leverage down to the target level would be one example; another is given in the quote from Standard & Poor's discussed in Section 2. Second, there may be concrete information on the current position in the cycle. Consider a firm which experiences a decrease in sales during a recession. Rating analysts then have to assess whether the decrease is cyclical, or due to permanent structural changes which coincided with the recession. Often, information necessary for such an assessment will only be revealed through time, [23] but certain firm announcements or actions (e.g. closing down production facilities) could reveal the nature of the shock.

The potential impact of the latter type of information can be gauged by examining the mean-squared error of the Kalman filter prediction. If there is information that allows to identify permanent and transitory components, previous estimates will be adjusted, and the average magnitude of adjustments will be related to the mean-squared error of the estimates. Whether outside observers consider adjustments to be large depends on the normal degree of variability. I thus examine the ratio of the mean-squared error of the predicted permanent asset value to the one-year variance of the permanent asset value, $\mathrm{MSE}(\hat{x}_t^*)/\mathrm{VAR}(x_t^* - x_{t-12}^*)$. This is done for four variance ratios, and assuming the forecast error to be zero in $t = 0$. The autocorrelation parameter ρ is set to 0.98.

Fig. 3 depicts the evolution of the ratio over time. If transitory effects are large (variance ratio = 0.6 or 0.7), it takes less than 18 months until the mean-squared error is larger than the annual variance of the underlying variable. That is, if raters were to learn the true permanent asset value after 18 months during which they based their ratings on estimates, the average magnitude of the ensuing rating

[23] Cf. Fitch IBCA (1999, p. 3): "(...) structural changes occur over long periods and are sometimes difficult to pinpoint until they are well advanced."

adjustment would be similar to the one brought about by one year's innovations in the state variable. For higher variance ratios, the mean-squared error can still amount to more than 50% of the annual variance of permanent effects. Thus, the potential for extreme rating changes that seem unwarranted by changes in current-condition default risk is large.

4. Summary and conclusion

The aim of this paper was to contribute to a better understanding of the rating methodology employed by rating agencies such as Moody's or Standard & Poor's. Using a structural model of default, I derived predictions about rating characteristics if ratings are meant to look 'through the cycle' as opposed to being based on the borrowers' current condition. I assumed that default risk is subject to both permanent and transitory shocks, and that rating agencies focus on the default probability in a stress scenario whose reference point is the permanent credit quality of a borrower. This necessitates a separation of permanent and transitory components. Within the model, the optimal way for doing this is the Kalman filter.

The analysis showed that empirical irregularities of agency ratings could be a consequence of the through-the-cycle method. Rating stability is significantly higher than with a current-condition approach. Ratings are not perfectly correlated with actual default risk, and they are correlated with past rating changes provided contemporaneous information is controlled for. Predictability in the usual sense can stem from errors in assessing the degree of cyclicality. The empirical evidence on ratings should therefore be interpreted with care. Apparent shortcomings of agency ratings might well be inherent to the rating method. Rating through-the-cycle does not per se lead to predictability of the type tested by Altman and Kao (1992), Lando and Skødeberg (2002), and Delianedis and Geske (1999). However, statistical tests of predictability should not presuppose independent observations across borrowers and time. Rating agencies use one method for rating many borrowers, and it may take considerable time until errors in the application of this method become evident to the raters.

Additional research could examine the pros and cons of rating through the cycle from the perspective of lenders, borrowers and regulators. In contrast to the through-the-cycle approach, assessing default risk based on the borrowers' current condition usually does not involve an analysis of long-term default risk dynamics. [24] Since prices of risky debt are affected by mean reversion (cf. Collin-Dufresne and Goldstein, 2001), agency ratings can contain valuable information beyond the one contained in current-condition ratings. This is not so much an intrinsic advantage of the through-the-cycle approach as a shortcoming of existing current-condition approaches. Under a current-condition architecture, the information on long-term

[24] The default rate model used by KMV, for example, does not account for differences in mean reversion across firms.

dynamics could be conveyed through a term-structure of ratings. If lenders prefer to base their decisions on a single rating, however, adopting a through-the-cycle approach may help to efficiently summarize the relevant information.

The Basel Committee on Banking Supervision (2001) has proposed to tie bank capital requirements closer to default risk. This poses the question whether banks should measure individual credit risk with a through-the-cycle or a current-condition approach. Even though through-the-cycle ratings are incomplete measures of short-term default term risk, they need not be inferior for the purpose of bank regulation. Among others, Estrella (2001) and Catarineu-Rabell et al. (2002) argue that regulators should avoid procyclicality in capital requirements. Relying on through-the-cycle ratings would be one way of achieving this objective.

Acknowledgements

I wish to thank an anonymous referee, Joost Driessen, Ralf Elsas, Jan-Pieter Krahnen, Markus Nöth, Peter Raupach, Christian Schlag, Mark Wahrenburg, participants of the 2001 European Finance Association meeting in Barcelona, the 2001 banking workshop in Münster, and seminars at Frankfurt and Mannheim Universities for helpful comments.

References

Altman, E., Kao, D.L., 1992. Rating drift in high-yield bonds. Journal of Fixed Income 2, 15–20.

Basel Committee on Banking Supervision, 2000. Credit ratings and complementary sources of credit quality information. Basel.

Basel Committee on Banking Supervision, 2001. A new capital adequacy framework. Basel.

Batchelor, R., Dua, P., 1991. Blue chip rationality tests. Journal of Money, Credit, and Banking 23, 692–705.

Blume, M., Kim, F., MacKinlay, C., 1998. The declining credit quality of US corporate debt: Myth or reality? Journal of Finance 53, 1389–2013.

Campbell, J., Lo, A., MacKinlay, C., 1997. The econometrics of financial markets. Princeton University Press, Princeton, NJ.

Cantor, R., 2001. Moody's investors service response to the consultative paper issued by the Basel Committee on Banking Supervision and its implications for the rating agency industry. Journal of Banking and Finance 25, 171–186.

Cantor, R., Fons, J.S., 1999. The evolving meaning of Moody's bond ratings. Moody's Investors Service.

Carey, M., Hrycay, M., 2001. Parameterizing credit risk models with rating data. Journal of Banking and Finance 25, 197–270.

Catarineu-Rabell, E., Jackson, P., Tsomocos, D., 2002. Procyclicality and the new Basel Accord – banks' choice of loan rating system. Working paper, Bank of England.

Collin-Dufresne, P., Goldstein, R., 2001. Do credit spreads reflect stationary leverage ratios? Journal of Finance 56, 1929–1958.

Crouhy, M., Galai, D., Mark, R., 2000. A comparative analysis of current credit risk models. Journal of Banking and Finance 24, 59–117.

Dangl, T., Zechner, J., 2001. Credit risk and dynamic capital structure choice. Working paper, University of Vienna.

Delianedis, G., Geske, R., 1999. Credit risk and risk neutral default probabilities: Information about rating migrations and defaults. Working paper, UCLA.

Ederington, L., Goh, J., 1998. Bond rating agencies and stock analysts: Who knows what when? Journal of Financial and Quantitative Analysis 33, 569–586.

Estrella, A., 2001. The cyclical behavior of optimal bank capital. Working paper, Federal Reserve Bank of New York.

Fama, E., French, K., 1988. Permanent and temporary components of stock prices. Journal of Political Economy 96, 246–273.

Fama, E., French, K., 2002. Testing tradeoff and pecking order predictions about dividends and debt. Review of Financial Studies 15, 1–33.

Fitch IBCA, 1999. Challenges in rating cyclical companies.

Graham, J., Harvey, C., 2001. The theory and practice of corporate finance: Evidence from the field. Journal of Financial Economics 60, 187–243.

Gupton, G.M., Finger, C.C., Bhatia, M., 1997. CreditMetrics – Technical document, New York.

Hamilton, J., 1994. Time Series Analysis. Princeton University Press, Princeton, NJ.

Hilderman, M., 1998. Industrial company rating methodology. Moody's Investors Service.

International Monetary Fund, 1999. International capital markets – developments, prospects, and key policy issues. IMF World Economic and Financial Surveys.

Kealhofer, S., 2000. Comment on "A new capital adequacy framework". KMV Corporation.

Kealhofer, S., Kwok, S., Weng, W., 1998. Uses and abuses of default rates. CreditMetrics Monitor (1st Quarter), 37–55.

Krahnen, J.P., Weber, M., 2001. Generally accepted rating principles: A primer. Journal of Banking and Finance 25, 3–24.

Lando, D., Skødeberg, T., 2002. Analyzing rating transitions and rating drift with continuous observations. Journal of Banking and Finance 26 (2/3), 423–444.

Löffler, G., 2002. Avoiding the rating bounce: Why rating agencies are slow to react to new information. Working paper, University of Frankfurt.

Merton, R.C., 1974. On the pricing of corporate debt: The risk structure of interest rates. Journal of Finance 29, 449–470.

Poterba, J., Summers, L., 1988. Mean reversion in stock prices. Evidence and implications. Journal of Financial Economics 22, 27–59.

Shumway, T., 2001. Forecasting bankruptcy more accurately: A simple hazard model. Journal of Business 74, 101–124.

Sobehart, J.R., Keenan, S.C., Stein, R.M, 2000. Benchmarking quantitative default risk models: A validation methodology. Moody's Investors Service.

Sobehart, J.R., Stein, R.M., Mikityanskaya, V., Li, L., 2000. Moody's public firm risk model: A hybrid approach to modeling short term default risk. Moody's Investors Service.

Standard & Poor's, 2000. Sovereign 1999 recap: The ups, the downs, and the moral of it all. CreditWeek (January 12), 6–8.

Standard & Poor's, 2002. Ratings Performance 2001.

Zhou, C., 2001. An analysis of default correlations and multiple defaults. Review of Financial Studies 14, 555–576.