Avoiding the rating bounce: why rating agencies are slow to react to new information — Source link ⧉

Gunter Löffler

Institutions: University of Ulm

Related papers:

- The Effect of Bond Rating Agency Announcements on Bond and Stock Prices

- An anatomy of rating through the cycle

- How rating agencies achieve rating stability

- The effect of bond rating changes on common stock prices

- Informational efficiency of credit default swap and stock markets: The impact of credit rating announcements

AVOIDING THE RATING BOUNCE: WHY RATING AGENCIES ARE SLOW TO

REACT TO NEW INFORMATION

Gunter Löffler *
Goethe-Universität Frankfurt

This version: May 2002

* Gunter Löffler, Chair of Banking and Finance, Goethe-Universität Frankfurt (Main), Mertonstr. 17, 60054 Frankfurt (Main), Germany. Tel.: ++49-69-79823143, facsimile: ++49-69-79822143. *E-mail address*: gloeffler@wiwi.uni-frankfurt.de

AVOIDING THE RATING BOUNCE: WHY RATING AGENCIES ARE SLOW TO REACT TO NEW INFORMATION

Abstract

Rating agencies state that they take a rating action only when it is unlikely to be reversed shortly afterwards. Based on a formal representation of the rating process, I show that such a policy provides a good explanation for the puzzling empirical evidence: Rating changes occur relatively seldom, exhibit serial dependence, and lag changes in the issuers' default risk. In terms of informational losses, avoiding rating reversals can be more harmful than monitoring credit quality only twice per year.

1 Introduction

Moody's Investors Service, one of the leading credit rating agencies, takes a rating action only "when it is unlikely to be reversed within a relatively short period of time" (Cantor, 2001, p. 175). As an explanation for this rating policy, Cantor cites the market's "expectation for stable ratings". Intriguingly, rating agencies are often accused of being too slow to adjust their ratings.[1] Could it be that the criticism rating agencies receive is the outcome of their desire to meet the market's preferences? This is the question I am going to address.

My analysis is built on a formal representation of the rating process. I model ratings as a mapping of a continuous variable, called credit quality, into discrete categories. Unmanaged, the discreteness produces dependencies in rating changes. The mechanics behind this feature, which is reminiscent of discreteness effects in stock price returns (see Campbell, Lo and MacKinlay, 1997), is as follows: If credit quality follows a standard probability distribution whose density declines monotonically towards the tails, a threshold triggering a rating change is more likely to be crossed by a small amount than by a large one. The closer the credit quality is to the rating boundary just crossed, however, the larger is the probability of a subsequent rating reversal relative to the probability of observing another rating change in the same direction. This bias towards rating reversals can be avoided by managing ratings as described in the above quote.

In this paper, rating management is implemented by setting tolerance regions

---

[1] Cf., for example, the Economist (1997, p. 70) on the Asian crisis: "The raters, firms such as Moody's Investors Service, Standard & Poor's, Duff & Phelps and IBCA, are supposed to be the financial markets' early warning system. Instead, the agencies have spent the past few months belatedly reacting to events."

around rating boundaries. If credit quality surpasses a boundary, but lies within the tolerance regions, the rating change is suppressed. Through simulations, I show that a policy of rating bounce avoidance could explain many of the empirical rating characteristics that have been interpreted as evidence of informational inefficiencies. With rating management, ratings are relatively stable, while rating changes are serially correlated and preceded by substantial changes in default probabilities.

Apart from rating bounce avoidance, the agencies' rating systems are characterized by another peculiarity. Most rating agencies employ a through-the-cycle approach, that is, neglect cyclical variations in credit quality when assigning ratings. Though related, the two features are distinct. They both lead to a decrease in rating volatility, but the problem of rating reversals arises even if credit quality is not cyclical.[2] Löffler (2001) shows that the through-the-cycle method, while able to explain important stylized facts like rating stability, fails to account for the predictability of rating changes.

Another possible explanation for the stylized facts is a slow processing of new information. Such an underreaction could be of a psychological nature, reflecting a common human trait (Edwards, 1968). It could also be due to infrequent revisions of ratings. The fact that rating agencies do not monitor ratings continuously is evident from their placing issuers on watchlists. Agencies would not have to devote special attention to individual issuers if all issuers were under continuous review anyway.

One could suspect that putting an issuer on credit watch indicates a situation in which the credit quality is no longer in line with the current rating but where the rating change is suppressed in order to avoid its likely reversal. If used in this way, credit watch could mitigate informational losses from rating bounce avoidance because it

---

[2] Consistent with this view, Cantor (2001) states that the avoidance of rating reversals supports the

would signal the true credit quality to outside observers. However, this is not what agencies claim to do:

> "These Watchlists list the names of credits whose Moody's ratings have a likelihood of changing. These names are actively under review because of developing trends or events which, in Moody's opinion, warrant a more extensive examination."[3]

> CreditWatch highlights the potential direction of a short- or long-term rating. It focuses on identifiable events and short-term trends that cause ratings to be placed under special surveillance by Standard & Poor's analytical staff. These may include mergers, recapitalizations, voter referendums, regulatory action, or anticipated operating developments.[4]

Putting a borrower on watch indicates a situation in which the probability of a change in credit quality is - due to imminent events - relatively high, not one where credit quality has already changed.

This paper cannot answer the question whether it is the raters' policy or informational inefficiencies that underlie the stylized facts. What it does show, for example, is that the effects of rating bounce avoidance can lead to substantial informational losses, similar to those brought about by infrequent monitoring. Any critique of the rating agencies runs the risk of being partial as long as it does not take the official rating policy into account. In consequence, market participants should ask rating agencies to reveal their rating policy in sufficient detail. Otherwise the market will not know what it gets, nor will it be able to evaluate the quality of rating agencies.

---

through-the-cycle approach.
[3] http://www.moodys.com/moodys/cust/watchlist/watchlist.asp, 23/01/2001.
[4] http://www.standardandpoors.com/ResourceCenter/RatingsDefinitions.html#creditwatch, 23/01/2001.

The related literature includes papers on empirical characteristics of agency ratings. Carey and Hrycay (2001) and Kealhofer, Kwok and Weng (1998) find that agency ratings are relatively stable compared to alternative rating systems. Altman and Kao (1998) and Lando and Skødeberg (2002) document the existence of serial dependence in rating changes. Delianedis and Geske (1999) show that rating changes lag changes in default probabilities. Their evidence is in line with previous findings that stock prices and stock analyst forecasts predict rating changes (e.g. Holthausen and Leftwich, 1986, and Ederington and Goh, 1998). Extant normative or descriptive papers on rating systems do not address the problem of rating bounces. Krahnen and Weber (2001) propose general standards for good rating practice. A report by the Basel Committee on Banking Supervision (2000) provides a comprehensive overview of rating practices. Crouhy, Galai and Mark (2001) describe the rating system of Moody's and Standard & Poor's and propose a prototype rating system for bank internal ratings.

The remainder of the paper is organized as follows. Section 2 presents a formalization of rating processes. Section 3 uses simulations to quantify the effects of rating policies on rating dynamics. Section 4 concludes.

2 Formalizing rating policies

Credit ratings can be viewed as a mapping of credit quality into discrete categories $i=1,..., N$. A borrower is assigned rating grade i if the credit quality z lies within the rating boundaries for that grade, b_{lower}^{i} and $b_{upper.}^{i}$ Many, but not all rating systems take this credit quality to be the probability of default.[5] Credit quality and default probability

---

[5] Cf. Basel Committee on Banking Supervision (2000) for a description of rating systems.

can be linked by introducing a default barrier d. A borrower defaults if the credit quality z falls below d. The probability of crossing the barrier is thus decisively affected by the value of z.

I consider a rating system with 17 rating categories (N=17) excluding default, which is the number of (modified) rating grades for which Moody's and Standard & Poor's publish default rate statistics. Without loss of generality, the default barrier d is set to zero. The random process describing the credit quality as well as the mapping rule are calibrated to meet the following requirements:

i)   the one-year probabilities of default for the median borrowers of each category are close to the historical default rates of corresponding agency ratings.

ii)   for median borrowers the probability that the rating remains stable is uniform across rating categories; absent rating management, it is equal to 35% on a one-year horizon. (A median borrower of grade i has a credit quality that is right in the middle of the credit qualities marking the boundaries to neighboring rating grades $i - 1$ and $i + 1$.)

A rating stability of 35% is below the figures reported in Standard & Poor's (2001), where the median stability across grades AAA to CCC is equal to 74.1%. However, the empirical stability of agency ratings is likely to be affected by the rating policy. Kealhofer, Kwok, and Weng (1998) report a transition matrix for ratings which are based on statistical estimates of default probabilities, that is, a transition matrix not affected by active rating management. There, the median stability is 44.4%. Since I use 17 rating grades instead of seven as in Kealhofer, Kwok and Weng, I regard a value lower than 44.4% to be appropriate. Empirically, rating stability tends to decline with credit quality. I choose not to model this feature because the differences are relatively

small and irregular, at least with respect to the neighboring rating grades where most transitions end. In Standard & Poor's (2001), for example, the rating stability for grades BBB-, BB+, BB, BB- and B+ is 76.7%, 69.3%, 72.8%, 71.8% and 78.2%, respectively.[6]

Many credit risk models[7] contain analogues to the state variable z used here. These variables are often interpreted as the borrowers' asset values, and mostly taken to be normally distributed. As shown by Gordy and Heitfeld (2001), however, empirical rating transition matrices are difficult to replicate with structural models that are based on a normally distributed state variable. Gordy and Heitfeld propose fat-tailed distributions of the alpha-stable family as an alternative. In this paper, I model the credit quality z as following a Student t distribution, whose application for credit portfolio risk analysis has been investigated by Frey and McNeil (2001).

To determine default probabilities for a given $z,$ and d equal to zero, I follow Merton (1974) and assume that default can only occur at the end of the horizon, instead of allowing for defaults within the horizon. With this assumption, it is straightforward to translate a given credit quality into a default probability.[8] Clearly, the assumption is critical when an asset value model is used to estimate the default risk of an issuer (see Zhou, 2001). For the purpose of this paper, however, it does not seem to be relevant whether defaults occur continuously or only at the end of horizon. The model is a stylized representation of the rating process; rating boundaries will be calibrated such that default probabilities are comparable to actual default rates. Some of the conclusions drawn from the analysis are based on changes in default probability that are observed within the model. These conclusions would be questionable if the relation between default probabilities prevailing under the two concepts were highly non-linear, which

---

[6] The figures in Standard & Poor's are adjusted by spreading the frequency of rating withdrawals proportionally across rating grades.
[7] See, for example, Gupton, Finger and Bhatia (1997).

they are not. In the case where credit quality follows a Brownian motion, for example, the one-year default probability is 2 $\Phi(-z_t)$ if one allows for continuous defaults, Φ denoting the standard normal distribution function (see Zhou, 2001). This is exactly twice the default probability that obtains when default is restricted to occur at the end of horizon. Note, too, that the incompatibility of normally distributed asset values with empirical rating data is not due to the assumption that default occurs at the end of horizon.[9]

The degrees of freedom of the asset value distribution are determined by minimizing the sum of squared differences between empirical default rates and the model default probabilities that are associated with rating grade midpoints, subject to the constraint that condition ii) is met.[10] With the t distribution, the probability of being in default at year-end is given by $T(-z_t, n)$, where $T(. , n)$ denotes the cumulative distribution function of a t distribution with n degrees of freedom; to obtain a rating stability of 35% for median borrowers, the width of a rating boundary has to be set equal to $2 \, T^{-1}(1 - 0.35 / 2, n)$. The lower boundary of rating grade 1, which corresponds to CCC, is zero. Thus, choosing the degrees of freedom n uniquely defines the rating boundaries as well as the default probability of median borrowers. The calibration is conducted separately for default rates from Moody's Investors Service (2001) and Standard & Poor's (2001). I let the degrees of freedom vary from 1 to 1000 and choose the one which minimizes the least squares criterion defined above. In either case, it results in a t distribution with 3 degrees of freedom. The goodness of the fit can be judged from Table 1, which also contains numerical values for rating boundaries that follow from choosing 3 degrees of freedom.

---

[8] Available solutions for first-passage-times are restricted to Brownian motion (see Zhou, 2001).
[9] This is shown in Löffler (2002).
[10] I minimize squared logarithmic differences summed over all ratings with a positive empirical default

Based upon the result of the minimization procedure, I assume that for a given horizon m, where m is measured in years, scaled changes in credit quality, $(z_{t+m} - z_t)/\sqrt{m}$, follow a t distribution with 3 degrees of freedom. Since the t distribution is a continuous mixture of normals, the framework can be formally represented as follows:

Credit quality: $\qquad\qquad\qquad z_t = \sqrt{\dfrac{3}{w}}\, x_t, \quad w \sim \chi^2(3)$

Default barrier $\qquad\qquad\qquad d = 0$

Lower boundary for grade i: $\qquad b^i_{lower} = 2\,(i-1)\,T^{-1}(1-0.35/2,\ 3), \quad i = 1,..., N$

Upper boundary for grade i $\qquad b^i_{upper} = 2\,i\,T^{-1}(1-0.35/2, 3), \quad i = 1,..., N-1$

Mapping rule (absent rating management):

$$b^i_{lower} < z_t \leq b^i_{upper} \Rightarrow Rating_t = i$$

where x is standard Brownian motion. Credit quality thus follows a random walk. For the purpose of the paper, this assumption is useful because it implies that there is no role for a through-the-cycle approach. It introduces an inconsistency because credit quality diverges to extreme levels as time passes. However, unreported analyses show that choosing a mean-reverting process does not change conclusions.[11]

Even though the state variable z follows a random walk, the rating derived from this variable will not be. Consider a borrower whose credit quality crosses a rating boundary. Since credit quality changes follow a bell-shaped distribution, the boundary is more likely to be exceeded by a small than by a large amount. Conditional on a rating

rate.

[11] Specifically, I choose an Ornstein-Uhlenbeck process with an annual variance of one and a half-life of random innovations equal to three years.

change, the probability that the rating change is reversed is thus larger than the probability that the rating change is followed by another change in the same direction. In the limit, when the credit quality just hits the boundary, the probability of a reversal is 50%, while the probability of observing another change in the same direction is, on a one-year horizon, equal to $T(-2\ T^{-1}(1 - 0.35/2, 3), 3) = 19.45\%$.

As noted in the introduction, Moody's claims to take a rating action only when it is unlikely to be reversed within a relatively short period of time. The description does neither specify the time horizon nor what is exactly meant by unlikely. Assume that a rating agency tolerates a reversal probability of p^* within a period of m years. Within the rating model described above, such a policy can be formulated as follows: The probability p that a rating change is reversed in the next m years depends on the difference between the credit quality and the rating boundary just crossed:

$$\text{Prob(reversal)} = p = \begin{cases} \text{Prob}\left(z_{t+m} \leq b^{crossed\ in\ t}\right) = T\left(\left(b^{crossed\ in\ t} - z_t\right)/\sqrt{m}, 3\right) \ \text{for upgrades} \\ \\ \text{Prob}\left(z_{t+m} > b^{crossed\ in\ t}\right) = T\left(\left(z_t - b^{crossed\ in\ t}\right)/\sqrt{m}, 3\right) \ \text{for downgrades} \end{cases}$$

The rating policy prescribes that the actual probability of reversal p is smaller than the target probability p^*. It can be implemented by requiring credit quality to exceed a rating boundary by at least $T^{-1}(p^*, 3)\sqrt{m}$. What happens if, in one period, credit quality crosses two boundaries, but fails to exceed the second boundary by the critical amount? In this case, the rating will be adjusted by one grade rather than two. Figure 1 shows the various possibilities for a single-period change in credit quality.

The following example illustrates the conservatism that can be introduced by such a rating policy. Assume $p^*=0.2$ and $m=1$, that is, the rating agency wants to avoid situations where rating changes are reversed with a probability of 20% within one year. A rating boundary then has to be exceeded by $T^{-1}(0.2, 3)=0.978$. This tolerance region

is almost as wide as the interval pertaining to one rating grade (1.007). In effect, such a rating policy would blur differences between neighboring rating grades.

3 Rating bounce avoidance as an explanation of stylized facts

In this section I use the rating model described above to assess whether the desire to avoid frequent rating reversals could underlie the peculiarities of agency ratings that have been documented in the literature. Notably,

- agency ratings appear relatively stable compared to other rating systems;

- ratings exhibit drift. Subsequent rating changes in the same direction are more frequent than subsequent rating changes in opposite directions;

- ratings lag changes in issuers' default probabilities.

These stylized facts will be addressed one after another within the framework laid out in the previous section. I assume that rating agencies pursue a policy of avoiding rating reversals. The tolerated probability for rating reversals p^* is set at 0.1, 0.2 or 0.3; the time period m is chosen to be 0.25, 0.5 and 1, corresponding to time intervals of three, six and twelve months respectively.

The effects of rating bounce avoidance are assessed through Monte Carlo simulations. Periodicity is one month. In one run of the simulations, I generate a random path for the credit quality z. According to the mapping rules from section 2, the credit quality z is translated into ratings. Since the assumed credit quality dynamics are independent of the current and past credit qualities, and the width of rating boundaries is uniform across grades, the starting value for the credit quality is not decisive for the results. I choose the initial credit quality to be 5.53, which is the credit quality of a median borrower within rating category BB. On a one-year horizon, the associated

default probability is 0.58%. One run of the Monte Carlo simulations extends over a period of three years. I perform 10,000 replications for each parameterization.

3.1 Ratings are relatively stable

Kealhofer, Kwok and Weng (1998) and Carey and Hrycay (2001) find that agency ratings are significantly less volatile than alternative ratings. Kealhofer, Kwok and Weng estimate default probabilities based on the default model of Merton (1974) and categorize borrowers according to these probabilities. Carey and Hrycay use a logit model to assign borrowers to rating grades. Typically, 40% to 50% of these ratings remain stable over a one-year horizon, compared to 80% to 90% in the case of agency ratings. The ratings constructed in Kealhofer, Kwok and Weng as well as in Carey and Hrycay are based on seven and five categories, respectively. Since I use 17 grades in this paper, rating stability will generally be lower. However, the simulation results and the empirical evidence can still be compared with respect to relative differences in rating stability.

Table 2 summarizes simulated one-year transition probabilities for various assumptions about the acceptable reversal probability p^* and the time horizon m used to compute this probability. If there is no rating bounce avoidance, that is, p^* is equal to the maximum value of 0.5, rating stability is equal to the 35% that were used to calibrate the model. For $p<0.5$, rating stability increases. It ranges from 40% to 68%. With a tolerated reversal probability of 0.2 and a time horizon of six months, the rating stability is 56%. Rating bounce avoidance can thus lead to a considerable increase in the stability of credit ratings. The simulated figures largely mirror the empirical differences between agency ratings and rating systems that are known not to be influenced by ratings management. Since the precise rating policy of the agencies is not known, it is

difficult to judge whether rating bounce avoidance completely explains the empirical evidence. There are two observations that indicate that it may not. Both in Kealhofer, Kwok and Weng (1998) and in Carey and Hrycay (2001), the stability of agency ratings is up to twice the one of alternative rating systems. Such an increase is not observed in the simulations, nor is the stability close to the maximum ones documented for Standard & Poor's modified grades. In Standard & Poor's (2001), the maximum empirical stability for modified grades is 90%, the median across the 17 grades is 78%. This does not leave us with a puzzle. The agencies' policy of rating through the cycle, which is not modeled here, can also lead to a significant increase in rating stability (Löffler, 2001). Together, the two peculiarities of the agencies' rating approach could well explain the empirical facts.

3.2 Rating changes are serially dependent

Empirical studies of rating changes have documented significant positive serial dependence (Altman and Kao, 1992, and Lando and Skødeberg, 2001). Such a dependence can arise even if ratings are continuous and rating analysts efficiently use available information. An analyst who learns that the default probability of a firm will decrease over time will not completely incorporate this information into the current rating if the rating horizon is shorter than the time span in which the firm's restructuring is completed. Partial responses to new information, however, will create positive serial dependence. Since I model credit quality as a random walk, this explanation can be ruled out for the simulation experiments conducted here.

Altman and Kao (1992) examine the rating dynamics of 1970-1985 new bond issues. They measure serial dependence through a statistic defined as the frequency of subsequent rating changes in one direction divided by the frequency of subsequent rating changes in opposite directions. If ratings exhibit positive drift, the statistic is

larger than one. If an upgrade is more likely to be followed by a downgrade, and vice versa, the statistic is smaller than one. I compute the statistic within the simulated samples. In each run, which spreads over three years, I take only the first two rating changes to compute the statistic. If a simulation run contains less than two rating changes, it does not enter the calculation of the test statistic.

The results are reported in Table 3. If raters do not try to avoid rating reversals ($p*$=0.5), the statistic is 0.42, meaning that the probability of observing rating changes in opposite directions is more than twice the one of observing rating changes in identical directions. This is the reflection of the rating bounce. If ratings are set to avoid this bounce, the statistics range from 1.02 to 3.30. For the parameter values chosen here, rating bounce avoidance thus leads to positive serial dependence in rating changes. The values are broadly in line with the ones reported by Altman and Kao (1992) separately for issuer groups, and for up- and downgrades. The mean (median) of their statistics is 1.752 (1.475), with a range of 0.2 to 3.83.

The rating policy modeled here could thus account for the existing evidence. This is important as an another peculiarity of agency ratings, the through-the-cycle approach, cannot (see Löffler, 2001). Lando and Skødeberg (2001) document that the rating drift is especially pronounced for downgrades. This is sometimes explained by noting that 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 13,1997, p. 70-71). But it could also be explained through avoidance of rating reversals. It seems likely that rating changes entail cost for the issuers, and that these costs are larger for downgrades. After a downgrade, investors may be forced to sell bonds, and covenants may restrict the flexibility of the borrower. If rating agencies act, at least partly, in the interest of their clients, they might try to avoid rating reversals particularly for downgrades. In the framework of this paper, such a policy is not the same as "doling out bad news in small

doses" because rating changes are delayed rather than handed out piecemeal. Rating drift arises because rating changes are only made when the credit quality is relatively close to the boundary triggering a further rating change.

3.3 Ratings lag changes in default probabilities

Based on the option-theoretic models of Merton (1974) and Geske (1977), Delianedis and Geske (1999) use balance sheet data, equity values and equity volatilities to compute risk-neutral default probabilities for borrowers rated by Standard & Poor's. They examine how these default probabilities evolve before a rating change, and find that they rise (fall) several months before a downgrade (upgrade).

Within the simulated samples, I examine the default probabilities one month before the first downgrade, regardless of the magnitude of the downgrade. Recall from Table 1 that the initial default probability is 0.58%; the median default probability of the next lower rating class is 1.01%. The critical value for credit quality z that separates the two grades is 5.02. If the rating agency does not aim at avoiding rating reversals, a downgrade occurs as soon as the default probability exceeds $T(-5.02, 3) = 0.76\%$. Due to the discrete nature of the rating system, the default probability one month before a downgrade will not be equal to the initial one. Downgrades are more likely to be observed if the credit quality has declined within the range associated with the initial rating. This effect is documented in Table 4. Even if raters are not concerned about reversals ($p^*=0.5$), the median default probability one month before a downgrade has increased relative to the initial one, from 0.58% to 0.69% The effect is relatively small, however, which is due to the rating system being relatively fine. With rating bounce avoidance ($p^*<0.5$), rating changes lag much more behind changes in default probabilities. Depending on the parameters, the median default probability one month before a rating change can be up to 1.18%, twice as large as the initial default

probability of 0.58%.

How do these figures compare to the results in Delianedis and Geske (1999)? For investment-grade bonds, the median risk-neutral default probability one month before any downgrade is 1.1%, while the median default probability of a benchmark sample which does not contain downgrades is 0.7%. Both in absolute and in percentage terms, the effects documented in Delianedis and Geske are thus smaller than the ones that can be generated through rating bounce avoidance.

In a related experiment, I compare the effects of the rating policy to the ones that would arise from infrequent rating reviews. Assume that agencies monitor ratings only in six month intervals. In the simulation, monitoring dates are thus t = 0, 6, 12 and so forth. At a monitoring date, the rating is set according to the credit quality, that is, there is no rating bounce avoidance. With such an infrequent monitoring, the simulated median default probability before a downgrade is 0.90%, which is lower than some of the values that obtain with rating bounce avoidance. With $p^*=0.2$ and $m=0.5$, for example, the median default probability is 0.98% (see Table 4). This shows that rating bounce avoidance can be more harmful to the quality of a rating system than restricting the frequency of reviews to only two per year.

4 Concluding remarks

The paper has shown that the wish to avoid frequent reversals of credit ratings could account for the stylized facts of agency ratings. Empirically, rating changes occur relatively seldom, they are serially dependent, and predictable using borrower fundamentals. Numerical simulations reveal that rating bounce avoidance can explain these peculiarities very well. Moreover, predictability cannot be explained by another

characteristic of the agencies' rating system, the through-the-cycle approach (Löffler, 2001). Rating bounce avoidance thus is an important candidate for explaining the stylized facts of agency ratings. Another candidate is informational inefficiency. If rating agencies are slow to react to new information, for instance because ratings are reviewed only infrequently, stability will increase, and rating changes will become predictable. Differentiating between these alternative explanations is difficult. The analysis has shown, however, that rating bounce avoidance can reduce the informational content of ratings by more than a rating system which monitors credit quality only twice per year.

Moody's claims that it manages ratings in order to "balance the market's need for timely updates on issuer risk profiles, with its conflicting expectation for stable ratings" (Cantor, 2001, p.175). It is beyond the scope of this paper to evaluate what the market really wants, and whether rating agencies act in response to these preferences, or use this to cover any deficiencies of their ratings. It seems obvious, however, that the market's preferences are not homogeneous. Rating management cannot serve all market participants alike. In addition, even if rating management exactly meets an investor's expectation for stability, there may be situations where this particular investor might want to know the precise credit quality, not the one obscured by rating management.

There seem to be two ways of reducing informational losses due to rating bounce avoidance. One is to communicate the precise rating policy, the other is to change the rating system. A move towards greater transparency would be to state how wide the tolerance regions are in terms of rating grades. The analysis has shown that rating management can blur differences between adjacent rating categories. In effect, rating management can offset the increase in accuracy achieved through the rating modification (+ and − in the case of Standard & Poor's) introduced by the rating agencies in the early 1980s. Rating agencies could try to elicit market feedback on

whether such an inaccuracy is indeed what the market wants.[12]

The problem of rating bounces could be reduced by moving from a discrete rating system to a continuous one. This does not imply that the rating is equated with default probabilities; it could still reflect other dimensions of credit risk, e.g. recovery risk, or be based on a combination of default probabilities for various time horizons. There are various possible arguments against continuous ratings. For cognitive reasons, rating analysts might find it easier to aggregate their information into discrete categories; rating agencies might be tempted to introduce random variation into ratings to pretend continuous monitoring activities; market participants might overestimate the accuracy of such a continuous rating. These arguments are appealing, but it has been shown in other contexts that continuously measured expectations can provide better results than qualitative ones. Batchelor (1986), for example, recommends to ask for continuous expectations of consumer price inflation rather than for qualitative responses.

The upcoming reform of capital adequacy requirements (Basel Committee on Banking Supervision, 2001) has spurred a discussion on the design of rating systems.[13] Since rating bounce avoidance appears to be a driving factor behind rating dynamics, it should receive more attention in this discussion. The issue is not confined to external rating agencies. If banks manage their internal ratings, too, the effects might be even more pronounced because many banks use less rating grades than the established rating agencies.[14]

---

[12] In the aftermath of the Enron default, Moody's has initiated a dialogue on the quality and timeliness of ratings (cf. Moody's Investors Service, 2002). So far, Moody's policy of avoiding rating reversals has not been questioned.

[13] See, for example, Krahnen and Weber (2001).

[14] See Krahnen and Weber (2001).

References

Altman E. and D. L. Kao (1992): Rating drift in high-yield bonds. Journal of Fixed Income, 15-20.

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

Basel Committee on Banking Supervision (2001): Overview of the new Basel capital accord. Basel.

Batchelor, R.A. (1986): Quantitative v. qualitative measures of inflation expectations. Oxford Bulletin of Economics and Statistics 48, 99-120.

Campbell, J., A. Lo and C. MacKinlay (1997): The econometrics of financial markets. Princeton: Princeton University Press.

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.

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

Crouhy, M., D. Galai, and R. Mark (2001): Prototype risk rating system. Journal of Banking and Finance 25, 47-95.

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

Economist (1997): Risks beyond measure. Dec 13, 70-71.

Ederington, L., and J. Goh (1998): Bond rating agencies and stock analysts: Who knows what when? Journal of Financial and Quantitative Analysis 33, 569-585.

Edwards, W. (1968): Conservatism in human information processing. In B. Kleinmuntz (Ed.), Formal representation of human judgment, New York, 17-52.

Frey, R., and A.J. McNeil (2001): Modeling dependent defaults. Journal of Banking and Finance, forthcoming.

Geske, R. (1977): The valuation of corporate liabilities as compound options. Journal of Financial and Quantitative Analysis 12, 541-552.

Gordy, M. and E. Heitfeld (2001): Of Moody's and Merton: a structural model of bond rating transitions. Working Paper, Board of Governors of the Federal Reserve System.

Gupton, G., C. Finger, and M. Bhatia (1997): CreditMetrics − Technical document. New York.

Holthausen, R. and R. Leftwich (1986): The effect of bond rating changes on common stock prices. Journal of Financial Economics 17, 57-89.

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

Krahnen, J.P. and M. Weber (2001): Generally accepted rating principles: a primer. Journal of Banking and Finance 25, 3-24.

Lando, D. and T. M. Skødeberg (2002): Analyzing rating transitions and rating drift with continuous observations. Journal of Banking and Finance, forthcoming.

Löffler, G. (2001): An anatomy of rating through the cycle. Working Paper, University of Frankfurt.

Löffler, G. (2002): Implied asset value distributions. 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.

Moody's Investors Service (2001): Default and recovery of corporate bond issuers: 2001.

Moody's Investors Service (2002): The Bond rating process: a progress report.

Standard & Poor's (2001): Ratings performance 2000. Default, Transition, Recovery, and Spreads.

Zhou, C. (2001): An analysis of default correlations and multiple defaults. Review of Financial Studies 14, 555-576.

Table 1

Rating boundaries and one-year default probabilities within the model

Rating	Rating Boundaries		Default probability at midpoint of rating boundaries	Historical default rates	
				S&P	Moody's
1 ~ AAA	16.07			0%	0%
2 ~ AA+	15.07	16.07	0.03%	0%	0%
3 ~ AA	14.06	15.07	0.04%	0%	0%
4 ~ AA-	13.06	14.06	0.04%	0.03%	0.06%
5 ~ A+	12.06	13.06	0.05%	0.02%	0%
6 ~ A	11.05	12.06	0.07%	0.05%	0%
7 ~ A-	10.05	11.05	0.09%	0.05%	0%
8 ~ BBB+	9.04	10.05	0.12%	0.12%	0.07%
9 ~ BBB	8.04	9.04	0.17%	0.22%	0.06%
10 ~ BBB-	7.03	8.04	0.24%	0.35%	0.39%
11 ~ BB+	6.03	7.03	0.36%	0.44%	0.64%
12 ~ BB	5.02	6.03	0.58%	0.94%	0.54%
13 ~ BB-	4.02	5.02	1.01%	1.33%	2.47%
14 ~ B+	3.01	4.02	1.95%	2.91%	3.48%
15 ~ B	2.01	3.01	4.34%	8.38%	6.23%
16 ~ B-	1.00	2.01	11.45%	10.32%	11.88%
17 ~ CCC	0.00	1.00	32.50%	21.94%	18.85%

The model default probabilities follow from assuming credit quality to follow a *t* distribution with 3 degrees of freedom, and assuming defaults (the default barrier is 0) to occur only at the end of horizon. Default rates are from Standard & Poor's (2001) and Moody's Investors Service (2001).

Table 2
Simulated one-year stability of credit ratings for different rating policies

p^* (tolerated reversal probability)	m (time horizon for reversal probability)		
	0.25	0.5	1
0.1	0.60	0.67	0.68
0.2	0.47	0.56	0.66
0.3	0.40	0.44	0.51
0.5	0.35	0.35	0.35

Table 3
Simulated serial dependence statistics for different rating policies

p^* (tolerated reversal probability)	m (time horizon for reversal probability)		
	0.25	0.5	1
0.1	2.84	3.17	3.30
0.2	1.47	2.10	3.15
0.3	1.02	1.27	1.76
0.5	0.42	0.42	0.42

The statistic is defined as the frequency of observing subsequent rating changes in the same direction divided by the frequency of observing subsequent rating changes in opposite directions. It is greater than one for positive serial dependence.

Table 4
Simulated median default probabilities in the month before a downgrade from grade 12 (default probability = 0.58%) for different rating policies

p^* (tolerated reversal probability)	m (time horizon for reversal probability)		
	0.25	0.5	1
0.1	1.04%	1.17%	1.18%
0.2	0.88%	0.98%	1.15%
0.3	0.79%	0.84%	0.92%
0.5	0.69%	0.69%	0.69%

Figure 1
Schematic representation of the rating policy for a one-period change in credit quality of a borrower rated i in $t = 0$