

Enhanced PD-implied Ratings by Targeting the Credit Rating Migration Matrix*

Jin-Chuan Duan[†] and Shuping Li[‡]

(May 19, 2021)

Abstract

A high-quality and granular probability of default (PD) model is on many practical dimensions far superior to any categorical credit rating system. Business adoption of a PD model, however, needs to factor in the long-established business/regulatory conventions built around letter-based credit ratings. A mapping methodology that converts granular PDs into letter ratings via referencing the historical default experience of some credit rating agency exists in the literature. This paper improves the PD implied rating (PDiR) methodology by targeting the historical credit migration matrix instead of simply default rates. This enhanced PDiR methodology makes it possible to bypass the reliance on arbitrarily extrapolated target default rates for the AAA and AA⁺ categories, a necessity due to the fact that the historical realized default rates on these two top rating grades are typically zero.

Keywords: default, other-exit, rating stickiness, sequential Monte Carlo

*The authors thank the editor and the two reviewers for their constructive comments to improve this paper.

[†]Asian Institute of Digital Finance and NUS Business School, National University of Singapore, Singapore. E-mail: bizdjc@nus.edu.sg. This research is funded under the grant “Asian Institute of Digital Finance” supported by the National Research Foundation, Singapore under its Industry Alignment Fund - Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

[‡]Asian Institute of Digital Finance, National University of Singapore. E-mail: Li Shuping lishuping@nus.edu.sg

1 Introduction

In recent years and particularly after the global financial crisis in 2008, the push to make credit risk assessments more granular and objective has gained momentum. The reasons in favor of using probability of default (PD) are many folds. The quality issue aside, conventional credit ratings based on letter grades of AAA, AA, etc. are too coarse to separate credit qualities of obligors. Finer separation is often needed for the task of credit underwriting and monitoring. Credit ratings are also difficult to aggregate across obligors, making them less than desirable for various tasks concerning credit portfolio management. In addition, credit ratings by the likes of S&P, Moody's and Fitch (the Big Three) or the classification methods like the Z-score of ? and the O-score of ? only offer relative rankings and their absolute meanings remain unclear above and beyond the historical default rates experienced for each category. The same criticism applies to their many modern variants based on machine learning.

In contrast, a PD model lends itself to rigorous scientific examinations through out-of-sample performance studies. Evidence abounds in the literature to suggest that a PD model can be constructed to dominate in performance the well-established commercial credit rating systems. Scientific inferiority or the introduction of the Dodd-Frank Act, however, has not meaningfully dented, post the global financial crisis, the widespread usage of commercial credit ratings issued by, say, the Big Three.

We reason that the continued wide acceptance of commercial credit ratings has more to do with business familiarity, long-established conventions and regulatory references built around ratings. A credit rating of S&P BBB- (or Moody's Baa3) and above is, for example, known as an investment-grade obligor/instrument meeting certain regulatory and/or fiduciary requirements. Merely offering PDs does not fit well with the common business framework of referencing credit ratings. For a scientifically sound PD model to gain a wider business acceptance, there must be a bridging device which converts PDs to equivalent letter ratings through referencing the default experience of a well-established commercial rating pool offered by a rating agency such as one of the Big Three. This paper aims to achieve this with a new way of generating PD-implied ratings (PDiR) via matching to the average rating migration matrix experienced by the likes of S&P.

The need to convert PDs to letter-based credit ratings is also reflected in the equity-implied-ratings of Moody's, which converts its market value based, PD-like Expected Default Frequency to a conventional credit rating assignment.¹ The PDiR of this paper takes a distinctly different mapping approach and applies on the corporate PDs produced and

¹Interested readers are referred to Market Implied Ratings FAQ issued by Moody's Analytics (June 2010).

maintained under the Credit Research Initiative (CRI) at National University of Singapore (NUS).

The CRI-PD database has a global coverage of 133 economies with over 70,000 exchange-listed firms, and the PDs are daily updated for a range of prediction horizons from one month to five years. The CRI-PDs rely on the forward intensity model of ?, which is part of the long line of literature on assessing likelihood of corporate default by modern statistical/econometric means, for example, ?, ? and ?, to name just a few. The CRI-PD model's implementation has factored in a number of practical and operational considerations required for concurrent applications to many economies around the world. The CRI-PD model undergoes monthly re-calibrations for its six sets of parameters with each unique to a region/type of economies (i.e., China, Emerging Economies, Europe, India, North America, and Other Developed Economies). With its six monthly re-calibrated PD models, the CRI produces daily updated PD values for virtually all exchange-listed firms globally and makes the PDs freely accessible via its web portal.²

The original PDiR methodology introduced in 2011, hereafter PDiR_{old}, allows the CRI team to reference commercial credit ratings in response to its users' request for making easier business applications of the CRI-PDs. A general discussion of PDiR_{old} is available in a CRI white paper (?) whereas its scientific description can be found in ?. PDiR2.0 proposed in this paper significantly improves upon the mapping method of PDiR_{old}, and has already been implemented in the current CRI system. The PDiR2.0 methodology shares the same conceptual objective but significantly improves the mapping method by reflecting a credit rating agency's rating migration experience instead of simply the realized default rates.

PDiR_{old} deploys twenty PD cutoff values to determine 21 mapped rating categories (including plus/minus modifier in the case of the S&P credit ratings). The mapping target is the conditional average one-year CRI-PDs for the 21 rating categories to their corresponding smoothed realized default rates reported by S&P for its global corporate rating pool. This minimization task cannot be performed individually for each of the 21 rating categories because altering one PD cutoff value affects at least two adjacent rating categories. Moreover, these cutoff values must be increasing with lowering credit qualities in order to be sensible. However, S&P's AAA and AA⁺ categories have not experienced any default over any one-year period, making their realized default rates unsuitable target values for these two categories. Knowing conceptually that the firms in these two rating categories still face default risk regardless of how minute they may be, one is forced to deploy extrapolated values. Thus, the PDiR_{old} mapping system has embedded some undesirable arbitrariness.³

²See ?. The CRI-PDs along with other useful credit risk measures are available at <https://nuscri.org>.

³We view such an extrapolation as arbitrary because moving from, say, AA to AAA, is a change in the ordinal not cardinal ranks.

Instead of targeting historical average default rates across different rating categories, we propose to match the average rating migration matrix. The AAA category in the migration history, for example, exhibits a high level of stickiness to remain at the same rating along with some nontrivial tendency to downgrade to the next rating category. Matching to the realized rating migration matrix allows us to avoid the PDiR_{old} method's reliance on those arbitrarily extrapolated values. Perhaps more significant, PDiR2.0 mimics rating stickiness observed in commercial credit ratings. The mapping method builds in migration buffer zones so that a firm's PD sitting close to a PD cutoff value will not quickly flip back and forth over two adjacent rating categories.

Technique-wise, we continue to deploy the probability-tempered sequential Monte Carlo (SMC) optimization technique along the line of ? to determine the suitable PD cutoff levels for each of many rating categories. Worth noting is the fact that ordinary gradient-based constrained optimization techniques are ill-suitable for the problem at hand, because our target function is discontinuous with respect to the cutoff values. The theoretical migration matrix under the model corresponds to a set of PD cutoff values, and is computed by counting entries from each rating category to another over one year. Therefore, a tiny move on a cutoff value for a very large sample of PDs such as the CRI database of over 70,000 exchange-listed corporations may still cause the theoretical migration matrix to jump, making it impossible to compute the derivatives.

As to the data, we take the CRI-PDs for all exchange-listed firms over the 18-year period from 2000 to 2017 to match the S&P global rating pool's reported average annual rating migration over the same period. Our final PDiR2.0 cutoff values and upgrade/downgrade buffer zones are provided in Table 1, which enables easy PDiR2.0 mapping for business usage.

Figures 1 and 2 illustrate our results on two firms – Lehman Brothers and Apple. The two graphs show that PDiR2.0 indeed meaningfully differs from PDiR_{old} , and both of which are responsive to credit quality changes as compared to the S&P ratings. PDiR2.0 vis-a-vis PDiR_{old} is by design more sticky through our introduction of migration buffer zones. Their levels also differ because the PD cutoff values have been altered, and arguably for the better because PDiR2.0 matches to the rating migration instead of only default rates. Figure 2 for Apple illustrates another point of interest, which in a way motivates our methodological improvement. It has been noted by industry users that the PDiR_{old} assigns too many AAA firms, and they tend not to migrate to other rating categories. The graph shows that Apple being assigned AA^+ under PDiR2.0 in the beginning of 2018 but later revised downward to AA and then AA^- . However, the ratings under the PDiR_{old} and the S&P rating remains at AAA and AA^+ , respectively, throughout the period.

Figure 1: Lehman Brothers' PDs and PDiRs

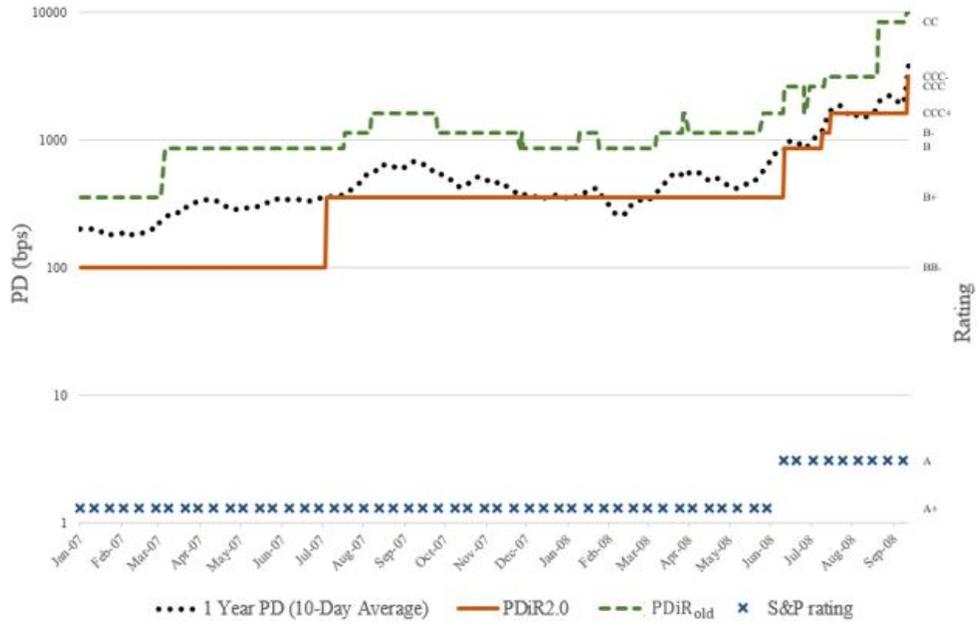


Figure 2: Apple Inc.'s PDs and PDiRs

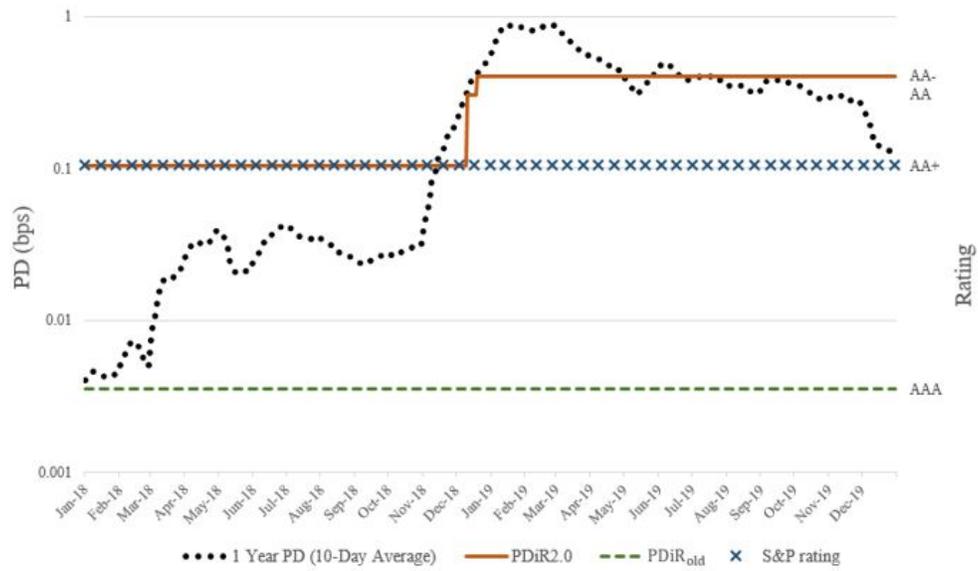


Table 1: The CRI one-year PD mapping with buffer zones to the PDiR by referencing the S&P global rating pool’s migration history. The cutoff values apply to the 10-day moving average of one-year PDs.

	Initial Assignment		Upgrade To		Downgrade To	
	lb (bps)	ub (bps)	lb (bps)	ub (bps)	lb (bps)	ub (bps)
AAA	0	0.0035	0	0.0027	-	-
AA ⁺	0.0035	0.1044	0.0027	0.0035	0.1044	0.3060
AA	0.1044	0.3060	0.0035	0.1044	0.3060	0.4069
AA ⁻	0.3060	0.4069	0.1044	0.3060	0.4069	1.2928
A ⁺	0.4069	1.2928	0.3060	0.4069	1.2928	3.0646
A	1.2928	3.0646	0.4069	1.2928	3.0646	3.9506
A ⁻	3.0646	3.9506	1.2928	3.0646	3.9506	9.9936
BBB ⁺	3.9506	9.9936	3.0646	3.9506	9.9936	22.0796
BBB	9.9936	22.0796	3.9506	9.9936	22.0796	28.1227
BBB ⁻	22.0796	28.1227	9.9936	22.0796	28.1227	46.2056
BB ⁺	28.1227	46.2056	22.0796	28.1227	46.2056	82.3715
BB	46.2056	82.3715	28.1227	46.2056	82.3715	100.4544
BB ⁻	82.3715	100.4544	46.2056	82.3715	100.4544	357.0556
B ⁺	100.4544	357.0556	82.3715	100.4544	357.0556	870.2578
B	357.0556	870.2578	100.4544	357.0556	870.2578	1126.8589
B ⁻	870.2578	1126.8589	357.0556	870.2578	1126.8589	1630.8764
CCC ⁺	1126.8589	1630.8764	870.2578	1126.8589	1630.8764	2638.9113
CCC	1630.8764	2638.9113	1126.8589	1630.8764	2638.9113	3142.9287
CCC ⁻	2638.9113	3142.9287	1630.8764	2638.9113	3142.9287	4449.8571
CC	3142.9287	8370.6423	2638.9113	7063.7139	4449.8571	8777.9817
C	8370.6423	10000	-	-	8777.9817	10000

2 PDiR mapping methodology

The first task is to identify the target migration matrix and determine whether we in light of the sample size should combine some finer rating categories together. In this paper, our target is the realized credit migration matrix of the S&P global corporate rating pool. Instead of dealing with 21 rating categories, we consider migration over nine consolidated categories (e.g., combining A⁺, A and A⁻ into A). In addition, we factor in the default and withdrawal rates. Our specific treatments and reasons for will be detailed later.

Mapping the CRI one-year PDs to letter ratings depends on finding appropriate PD cutoff values, i.e., upper and lower bounds for each rating category. Obviously, the upper bound

of one rating category is the lower bound for another rating category of lesser credit quality. Therefore, all cutoff values are tightly linked and need to be monotonically decreasing with improving credit quality. In addition to these cutoff values, we also need to define migration buffer zones in order to build in rating stickiness. Without these buffer zones, ratings are prone to flip back and forth for firms with PDs close to some cutoff values. In short, the cutoff values along with the buffer zones define credit migration for firms in the sample and form a critical component of the mapping methodology. For the balance of the paper, we will refer to this core of our mapping method as the rating assignment operator.

Next, we need to define an optimization objective function that reflects the goodness of fit between the target credit migration matrix and the one deduced by the rating assignment operator which is functionally defined over a set of PD cutoff values. Finally, a workable optimization algorithm is required to solve this very complex optimization problem.

2.1 S&P credit migration matrix

Our S&P's realized one-year migration matrices (rating migration from the beginning of a year to the end of a year) for 18 years, from 2000 to 2017, are taken from European Securities and Markets Authority (ESMA).⁴ The reported numbers are for the nine natural consolidated categories (i.e., AAA, AA, A,..., CC, C).

Beyond these nine consolidated rating categories, the default rates over the subsequent one year for different rating categories are naturally important. In addition, it is imperative for us to factor in the fact that the withdrawal rate from the S&P ratings likely differs from the CRI sample's exit rate for reasons other than default. The reason is fairly obvious. A corporate obligor may choose not to be rated by S&P for many possible reasons; for example, a CC rated firm is likely reluctant to pay for a CC rating because it is far worse than having no rating at all. In contrast, the other-exit rate for the CRI sample is a result of stock exchange de-listings for reasons other than bankruptcies. The data reveal that the driving force behind the other-exit rate is mergers/acquisitions.

The above consideration takes us to a critical adjustment; that is to gross up the migration matrix with the withdrawal/other-exit rate specific to each of the nine rating categories. Specifically, we divide all entries of each row (including default rate) of the S&P credit migration matrix by one minus the withdrawal rate unique to that rating category. For the model's implied migration matrix based on tracking the firms in the CRI sample, we divide the entries of each row instead by the other-exit rate specific to that category. Such an adjustment results in a 9×10 matrix (for the target or model migration matrix) where every row sums up to 1. The grossed-up credit migration matrix for the S&P global rating

⁴ESMA's Central Repository at <https://cerep.esma.europa.eu/cerep-web/>

pool, denoted by \hat{M} , and the withdrawal rates for different rating categories are displayed in Table 2.

Table 2: The S&P average realized one-year migration matrix (grossed up by the withdrawal rates in the bottom row)

	AAA	AA	A	BBB	BB	B	CCC	CC	C	Default
AAA	0.82228	0.16552	0.00849	0	0	0.00265	0.00106	0	0	0
AA	0.00332	0.90381	0.08639	0.00582	0.00022	0.00022	0	0.00007	0	0.00015
A	0.00013	0.01932	0.92681	0.04951	0.00229	0.00071	0.00032	0.00035	0.00003	0.00052
BBB	0.00014	0.00092	0.03741	0.91678	0.03749	0.00421	0.00068	0.00055	0.00014	0.00168
BB	0.00012	0.00058	0.00082	0.05136	0.86076	0.07292	0.00485	0.00111	0	0.00748
B	0	0.00025	0.00075	0.00191	0.059962	0.84023	0.05458	0.00483	0	0.03748
CCC	0	0	0	0.00042	0.00333	0.17194	0.52498	0.02290	0	0.27644
CC	0	0	0.00727	0	0.01091	0.06901	0.13455	0.27273	0	0.50545
C ⁵	0	0	0	0	0	0.25	0	0	0.25	0.5
Withdrawal Rate	0.02835	0.06150	0.04972	0.06794	0.09808	0.11927	0.23673	0.30380	0.55556	

2.2 Rating assignment operator and its implied migration matrix

Since PD only takes on values between 0 and 10,000 bps, the lower (upper) bound for the best (worst) credit quality category is naturally set. With the nine consolidated rating categories, eight PD cutoff values are needed. We denote them by $\{U_{AAA}, U_{AA}, U_A, U_{BBB}, U_{BB}, U_B, U_{CCC}, U_{CC}\}$. Because of the migration buffer zones and the eventual need to accommodate the plus and minus rating modifiers, we need additional thresholds to set their respective ranges. Instead of treating the thresholds as free parameters, we divide each PD segment defined under the nine consolidated rating categories into four equal subsegments.⁶ With the higher (lower) 25% PD piece classified as for the minus (plus) subcategory, and through which we also define the migration buffer zones. Although AAA, CC and C do not have plus or minus modifier, the 25% or 75% subsegment is still needed for setting the migration buffer zones.

The rating assignment operator runs on two modes – initial assignment and migration determination. We use the 10-business day moving average PDs for initial assignment and migration determination. Adopting a 10-day moving average is to smooth out PD estimates

⁵For the sample period, the ESMA database only reports nine cases for the C category with one migration to B, one stay in C, two defaults and five withdrawals.

⁶Forcing a buffer zone to correspond to an entire rating notch may appear arbitrary. We have actually experimented with introducing free parameters to define the buffer zones. Optimization ran smoothly without experiencing any difficulty, and the fitting performance naturally improved. But these buffer zone boundaries are quite variable across different rating categories, indicating that different rating groups comprise significantly varying numbers of firms. Our final modeling choice reflects trading-off the added goodness-of-fit for the model’s parsimony and stability.

that inevitably inherit trading and/or other measurement noises. Initial rating assignment is straightforward. When a firm is assigned a PDiR2.0 for the first time, its rating is determined by the 10-day moving average PD falling into a particular segment along with the top and bottom 25% sub-segments for setting a rating modifier if applicable.

Migration determination is more complicated and follows the following steps.

- Migrating to any of AA⁺, AA, AA⁻,..., CCC⁺, CCC, CCC⁻ requires the firm's 10-day moving average PD to cross beyond one whole subcategory. For example, an A obligor is downgraded to BBB⁺ only if its PD moves into the interval defined by BBB. Likewise, to upgrade a BBB⁺ obligor to A, its PD must move into the interval defined by A⁺.
- To upgrade a firm to AAA (or CC), its 10-day moving average PD must be lower than the PD level corresponding to 75% of the AAA (or CC) interval.
- To downgrade a firm to CC (or C), its 10-day moving average PD must be larger than the PD level corresponding to 25% of the CC (or C) interval.

Running the migration assignment operator on the firms in the CRI sample over the 18-year sample period produces their time series of ratings falling into some of the 21 rating categories. Finally, we put all firms into the nine consolidated rating categories and compute the model's implied 9×10 migration matrix. We further adjust the entries of each row by grossing up with the other-exit rate for each of the nine rating category computed with the CRI model's probability of other exit formula for the firms falling in that category. Denote the model's final implied migration matrix for the CRI sample by $\mathbf{M}(\boldsymbol{\theta})$ where $\boldsymbol{\theta}$ stands for the set of PD cutoff values, i.e., $\{U_{AAA}, U_{AA}, U_A, U_{BBB}, U_{BB}, U_B, U_{CCC}, U_{CC}\}$.

2.3 Model calibration

The eight model parameters must satisfy the constraint: $0 < U_{AAA} < U_{AA} < U_A < U_{BBB} < U_{BB} < U_B < U_{CCC} < U_{CC} < 1$. Furthermore, we impose an additional constraint that the resulting proportion of AAA firms over the 18-year period should be at least as high as the proportion of AAA firms in the S&P global rating pool, which is 1.53% over the 18-year period. In the later implementation, we round it to 1.5% instead.

Our optimization problem is thus to find $\boldsymbol{\theta}$ to minimize the sum of squared differences between the S&P credit migration matrix and the model's implied migration matrix. First, we need to define an index set for the ease of further exposition. This index set corresponding to a row of the migration matrix allows us to limit the fitting to the diagonal, two off-diagonals

(one in each direction) and the default column of the migration matrix; that is,

$$\mathbf{A}_i = \begin{cases} \{1, 2, 10\} & i = 1 \\ \{i - 1, i, i + 1, 10\} & 2 \leq i \leq 8 \\ \{8, 9, 10\} & i = 9 \end{cases}$$

Moreover, we denote by $P_{AAA}(\boldsymbol{\theta})$ the percentage of firms classified as AAA according to the PD cutoff values. Using the observed percentage of AAA firms under the S&P global rating pool, i.e., $P_{AAA}(\text{S\&P}) = 1.5\%$, our optimization explicitly matches the CRI sample’s average realized one-year migration matrix under the PDiR2.0 methodology to the average one-year rating migration matrix experienced by the S&P global rating pool; that is,

$$\begin{aligned} \min_{\boldsymbol{\theta}} \quad & \sum_{i=1}^9 \sum_{j \in \mathbf{A}_i} (\mathbf{M}_{i,j}(\boldsymbol{\theta}) - \hat{\mathbf{M}}_{i,j})^2 \\ \text{subject to} \quad & \begin{cases} \boldsymbol{\theta} \text{ satisfies } 0 < U_{AAA} < U_{AA} < \dots < U_{CCC} < U_{CC} < 1 \\ P_{AAA}(\boldsymbol{\theta}) \geq P_{AAA}(\text{S\&P}) \end{cases} \end{aligned} \quad (1)$$

Limiting our fitting to the diagonal, two off-diagonals (one in each direction) and the default column of the migration matrix is to ignore multi-category moves which are rare and noisy.

Note that the above objective function is discontinuous with respect to $\boldsymbol{\theta}$, because any cutoff value falls between two adjacent PDs in the sample will yield the same functional value, whereas changing a cutoff value that happens to equal a PD value of a firm over the sample period will result in a jump in the functional value. Thus, usual gradient-based optimization algorithms cannot be directly applied to solve this problem. We thus deploy the probability-tempered sequential Monte Carlo (SMC) method along the line of ? to obtain the optimal solution. The technical detail is provided in Appendix.

3 Results

Applying the PDiR2.0 methodology on the data produces the PD cutoff values. Table 1 reports the results which define ratings and migration. The calibrated model gives rise to the CRI sample’s average realized one-year migration matrix in Table 3 which is tallied over the 18-year sample period after being grossed up by the other exit rates corresponding to different rating categories. This CRI sample’s realized migration matrix under the PDiR2.0 methodology is designed to mimic to the extent possible the S&P rating migration matrix reported in Table 2.

Table 3: The PDiR2.0 average one-year migration matrix (grossed up by the other-exit rates in the bottom row)

	AAA	AA	A	BBB	BB	B	CCC	CC	C	Default
AAA	0.74597	0.14508	0.06703	0.02620	0.01283	0.00221	0.00014	0	0	0.00055
AA	0.05248	0.68795	0.18582	0.05394	0.01800	0.00141	0.00013	0	0	0.00028
A	0.00251	0.11907	0.66472	0.16292	0.04624	0.00372	0.00039	0.00002	0	0.00040
BBB	0.00029	0.01034	0.13100	0.69974	0.14855	0.00786	0.00083	0.00006	0	0.00132
BB	0.00014	0.00128	0.01361	0.20471	0.73579	0.03321	0.00286	0.00041	0.00001	0.00797
B	0.00002	0.00048	0.00311	0.04339	0.18840	0.69900	0.01649	0.00280	0.00022	0.04609
CCC	0	0.00059	0.00118	0.02309	0.06809	0.33807	0.27413	0.02605	0.00059	0.26821
CC	0	0	0	0.00426	0.02553	0.15319	0.13617	0.11064	0.00426	0.56596
C	0	0	0	0	0	0.08333	0.08333	0.08333	0.08333	0.66667
Other-Exit Rate ⁷	0.04567	0.03395	0.03469	0.03817	0.05308	0.11553	0.35584	0.48690	0.5	

A few observations are in order. The CRI sample’s realized migration matrix under PDiR2.0 is diagonal-heavy just like the S&P migration matrix, but discrepancies exist; for example, AAA in the CRI sample under the calibrated model has about 75% of chance to remain at AAA as compared to 82% in the S&P migration matrix. Note that both have been grossed up by the other-exit/withdrawal rates. Were the constraint on the proportion of AAA firms in the CRI sample to exceed 1.5% as experienced by the S&P global corporate rating pool not enforced, the gap could be further narrowed. In short, this is a price that one pays to ensure a reasonable number of AAA firms under the PDiR2.0 methodology. Of course, this constraint induces a ripple effect down the rating ladder and manifests in other diagonal elements; for example, AA firms have the 69% probability to remain as AA in contrast to the S&P experience of 90%.

By the CRI sample’s realized migration matrix, AAA firms face a one-year realized default rate at 5.5 basis points in contrast to 0% experienced by the S&P global corporate rating pool. This can be understood as a chance happening or simply put, a sampling error. According to Table 1, the one-year PD for the AAA category cannot exceed 0.0035 basis points. Factoring in migration stickiness, an AAA firm’s one-year PD will still be capped at 0.306 basis points. In short, an extremely low probability event had occurred to some AAA firms in the CRI sample. Alternatively, one could argue that the quality of the PDs for this small set of firms is not up to par, but that would be an issue concerning the CRI-PD model as opposed to the PDiR2.0 methodology. As to the default rates for other rating categories, they seem to match the realized rate to a reasonable degree.

To get a sense of time variations, we present in Table 4 the diagonal entries of the migration matrix in year 2008 and 2017 for both the CRI sample under the PDiR2.0 methodology

⁷For the sample period, the calibrated model assigns 24 firms in the CRI database to the C category. Among them, 12 firms experienced other exits, eight firms defaulted, and the remaining ones either migrated or stayed.

and the S&P global corporate rating pool. Discrepancies are expected because the PDiR2.0 methodology attempts to match the average credit migration not the individual migration matrices. In fact, matching individual migration matrices makes little sense because the CRI-PD model and S&P deploy entirely different concepts and methodologies in assessing credit risks.

Table 4 nevertheless shows that a significantly higher proportion of the CRI sample firms being downgraded from their investment grades under PDiR2.0 during the 2008 global financial crisis as compared to the S&P sample firms, reflecting the fact that the CRI-PDs are more responsive to market conditions. Perhaps also at play is related to the common belief that credit rating agencies are reluctant to downgrade obligors until they have to face up to the reality.

Table 4: Staying-put rates for rating categories in 2008 and 2017

	2008		2017	
	CRI Sample	S&P	CRI Sample	S&P
AAA	33.03%	58.74%	84.67%	78.57%
AA	25.18%	79.30%	76.19%	98.46%
A	29.99%	92.93%	69.30%	95.66%
BBB	44.10%	92.02%	71.95%	94.63%
BB	80.90%	82.39%	74.70%	89.43%
B	81.18%	82.08%	72.65%	87.60%
CCC	18.18%	47.44%	48.25%	51.93%
CC	22.22%	20.00%	36.36%	30.00%
C	0.00%	0.00%	0.00%	0.00%

Table 5: Proportions of the CRI sample firms in years around 2008

	2006	2007	2008	2009	2010	2011
AAA	1.794%	1.682%	1.466%	0.621%	0.464%	0.999%
AA	11.348%	9.870%	7.935%	2.707%	3.383%	6.813%
A	24.116%	22.647%	20.116%	10.183%	13.642%	19.057%
BBB	33.998%	35.904%	36.426%	28.053%	39.773%	38.501%
BB	22.168%	24.133%	27.746%	45.340%	35.234%	28.577%
B	6.389%	5.561%	6.056%	12.053%	7.365%	5.902%
CCC	0.171%	0.186%	0.198%	0.864%	0.117%	0.139%
CC	0.011%	0.017%	0.057%	0.166%	0.022%	0.012%
C	0.004%	0%	0%	0.012%	0%	0%

To get a sense of the number of firms in the CRI sample of over 70,000 exchange-listed corporations in 133 economies falling into each of the nine consolidated rating categories, we present in Table 5 their proportions in individual years from 2006 to 2011, for which we purposely cover 2008, the year of the global financial crisis. It is comforting to see that the PDiR2.0 methodology has assigned significantly smaller numbers of highly-rated (A and above) firms in 2008 and two years after.

Had we removed the condition that requires AAA firms to be no less than 1.5% of the whole sample, the proportion of AAA firms under a re-calibrated model would have dropped significantly to around 0.06%. The impact of the AAA percentage requirement can be seen in Table 6 where the migration has been estimated without the constraint.

Table 6: The PDiR2.0 average one-year migration matrix without the AAA percentage constraint (grossed up by the other-exit rates in the bottom row)

	AAA	AA	A	BBB	BB	B	CCC	CC	C	Default
AAA	0.78629	0.10887	0.05242	0.0121	0.02016	0.02016	0	0	0	0
AA	0.01015	0.81111	0.11396	0.04496	0.01452	0.00453	0.00031	0	0	0.00047
A	0.00005	0.01988	0.77656	0.16058	0.03693	0.00539	0.00033	0.00001	0	0.00026
BBB	0.00001	0.00044	0.11865	0.72987	0.13496	0.01439	0.00089	0.00002	0	0.00077
BB	0	0.00011	0.00884	0.14787	0.77678	0.05909	0.00271	0.00017	0.00001	0.00442
B	0	0.00003	0.00169	0.01967	0.2262	0.70225	0.01507	0.00138	0.00015	0.03356
CCC	0	0	0.00021	0.00916	0.06415	0.26494	0.48115	0.00958	0.00062	0.17017
CC	0	0	0	0	0.02548	0.08917	0.15924	0.08917	0	0.63694
C	0	0	0	0	0	0.08333	0.08333	0.08333	0.08333	0.66667
Other-Exit Rate	0.15358	0.04217	0.03425	0.03652	0.04475	0.09961	0.26184	0.50629	0.5	

Comparing Tables 3 and 6, removing the AAA constraint has the anticipated effect of driving the staying-put rates for all rating categories closer to the S&P’s historical experience. Moreover, the default rate for AAA firms also drops to 0%, same as the S&P records. It is debatable, however, whether placing the AAA percentage constraint is desirable from the vantage point of financial practitioners. A PDiR system assigns far fewer AAA firms than the long-established benchmark may cause operational difficulties in credit portfolio management and thus hinders the acceptance of a scientifically superior credit assessment system.

4 Conclusion

We propose and illustrate an enhanced PDiR methodology that matches the model’s implied average migration matrix deduced from 70,000 exchange-listed firms of 133 economies in the NUS-CRI database to the historical credit migration experience of the S&P global corporate rating pool. This PDiR2.0 methodology can be likewise executed by referencing other corporate rating pools, for example, Moody’s and Fitch. It can also be applied on completely

different types of credit pool, say, consumer credits, through benchmarking a PD model on consumers against an existing scoring system.

With the help of PDiR2.0 methodology, businesses are more likely to migrate to a superior granular PD system while ensuring a high degree of operational continuity and compatibility with their existing management infrastructure (credit approval, credit limits, credit monitoring, etc) built around credit ratings/scores over many years.

The percentage of AAA firms (being a cap, floor or both) in our modeling set-up can be easily altered to meet a mapping system designer’s preference. Similar constraints can also be placed on other rating categories with minor technical adjustments. In summary, the PDiR2.0 methodology is robust and flexible enough to come up with bespoke mapping systems catering to specific user demands.

Appendix: A probability-tempered sequential Monte Carlo algorithm

The general description of the algorithm closely follows that of ? which is in turn based on ?. The idea is to turn our minimization problem into a sampling problem with a simple transformation of the objective function, i.e., $f(\boldsymbol{\theta}) \propto \exp[-L(\boldsymbol{\theta})]$ with $L(\boldsymbol{\theta})$ denoting the objective function in (1). Note that $f(\boldsymbol{\theta})$ can be viewed as a probability distribution function up to a norming constant. It is not a density function because $L(\boldsymbol{\theta})$ is as discussed in the main text a discontinuous function with respect to $\boldsymbol{\theta}$.

Sequential Monte Carlo (SMC) is a powerful way to sample $\boldsymbol{\theta}$ under $f(\boldsymbol{\theta})$ without having to know the norming constant. The point in the sample that corresponds to the highest transformed functional value will be our solution to the original minimization problem. This SMC technique is particularly useful in the current context because $L(\boldsymbol{\theta})$ is, due to the nature of our rating migration problem, discontinuous with respect to $\boldsymbol{\theta}$. Theoretically, there are an infinite number of points that attain the same optimal functional value, but this non-singleton set is extremely small in volume (i.e., the Lebesgue measure) in the space for $\boldsymbol{\theta}$ due to the fact that the CRI sample is a very large set of firm-year observations. In short, non-uniqueness does not present a practical issue.

Our SMC algorithm is executed with an initialization sampler, $I(\boldsymbol{\theta})$, which is not a prior distribution in the sense of the Bayesian statistics. It is because the SMC optimizer is rather generic and need not be for statistical analyses. Apart from computational efficiency, the only technical condition imposed on $I(\boldsymbol{\theta})$ is the necessity of having its support to contain the target function. Sequential sampling in a sequence of self-adaptive n steps is a must, because completing good sampling in a single step is practically impossible. The evolving sample through sequential sampling will after a while reaches its terminal state, which is

$f(\boldsymbol{\theta})$. The intermediate target probability distribution is by design a tempered value of $f(\boldsymbol{\theta})$ where γ_n falls between 0 and 1; that is,

$$f_n(\boldsymbol{\theta}^{(n)}) \propto \left(\frac{\exp[-L(\boldsymbol{\theta}^{(n)})]}{I(\boldsymbol{\theta}^{(n)})} \right)^{\gamma_n} \times I(\boldsymbol{\theta}^{(n)}). \quad (2)$$

Our probability-tempered sequential Monte Carlo algorithm is a five-step procedure. For the following discussion, we set the initial SMC sample size to 1,000.

- **Step 1: Initialization**

Draw an initial sample of 1,000 particles, and denote them by $\boldsymbol{\theta}^{(0)}$. Each dimension of $\boldsymbol{\theta}$ is based on a normal distribution fixed at some sensible mean and variance, and it is made to be independent of others. Since the elements of $\boldsymbol{\theta}$ must be increasing in value and its implied proportion of AAA firms must be at least 1.5%, we will reject those sampled particles failing to meet these requirements.

- **Step 2: Reweighting and resampling**

At each tempering step, we perform reweighting to prepare for the advancement to the next tempering step in the self-adaptive sequence. Let $\boldsymbol{w}^{(n)}$ denote the weights vector corresponding to the 1,000 particles and “.” be the dot-product operator. Thus,

$$\boldsymbol{w}^{(n)} = \boldsymbol{w}^{(n-1)} \cdot \left(\frac{\exp[-L(\boldsymbol{\theta}^{(n-1)})]}{I(\boldsymbol{\theta}^{(n-1)})} \right)^{\gamma_n - \gamma_{n-1}} \quad (3)$$

where the next γ_n is determined to maintain a minimum effective sample size (ESS) of 500, i.e., 50% of the intended sample size. The ESS is per usual defined as $\frac{(\sum_{i=1}^{1000} w_i)^2}{\sum_{i=1}^{1000} w_i^2}$. It should be clear that if a few particles carry heavier weights vis-a-vis others, the ESS will become lower, meaning that the sample at hand is effectively small. Such a γ_n always exists if the previous $\boldsymbol{w}^{(n-1)}$ exceeds the ESS threshold. The self-adaptive advancement of γ_n to reach 1 is typically quite fast.

Resampling $\boldsymbol{\theta}^{(n-1)}$ according to the weights $\boldsymbol{w}^{(n)}$ to produce an equally-weighted sample, $\boldsymbol{\theta}^{(n)}$. After resampling has been performed, $\boldsymbol{w}^{(n)}$ has to be reset to the vector of 1's.

- **Step 3: Support boosting**

The empirical support has shrunk due to resampling and needs to apply the Metropolis-Hastings (MH) move to boost the support, which is conducted as follows:

- Propose new θ^* based on some proposal distribution, $\mathcal{Q}(\theta^*|\theta^{(n)})$. Directly deploying normal kernels tends to generate particles that violate the constraints stated earlier. Attempting to replace the entire vector in one go is also too aggressive leading to a low acceptance rate. An efficient and intuitive random-segment proposal is actually possible in our case. First of all, we sample random starting and ending indices between 1 and 8 of the target 8-dimensional vector with an intention to only replace this segment. Use a regression proposal constructed on the 1000-particle sample where each element for replacement progressively deploys a regression of that element on two other anchoring elements defined by the starting index minus 1 and the ending index plus one. The starting index is revised upward by one sequentially as one progresses through the segment for replacement. If the starting (or ending) index turns out to be 1 (or 8), no anchoring at that position is possible, and in that case regression will be conducted on one fewer regressor. This random-segment regression sampler can yield a very high acceptance rate with only a few sampled particles being rejected due to their violation of the constraints stated earlier.
- Compute the MH acceptance probability, α_i , for each of the 1,000 particles:

$$\alpha_i = \min \left(1, \frac{f_n(\theta_i^*) \mathcal{Q}(\theta_i^{(n)}|\theta^*)}{f_n(\theta_i^{(n)}) \mathcal{Q}(\theta_i^*|\theta^{(n)})} \right). \quad (4)$$

- With probability α_i , set $\theta_i^{(n)} = \theta_i^*$, otherwise keep the old particle.

Repeatedly perform the above MH move until the cumulative acceptance rate reaches 200%.

- **Step 4: Repeat Steps 2 and 3 until reaching $\gamma = 1$**
- **Step 5: k -fold duplication**

We further apply the k -fold duplication proposed in ? to increase the SMC sample size which allows us to bypass the tempering steps. Specifically, the original sample of size 1,000 is duplicated to a sample of size $1000k$ after adding $k-1$ identical copies. Then, perform the support-boosting step as in Step 3 to remove duplicated particles and turn the sample into a truly representative sample of size $1000k$. In the implementation, we perform at most two rounds of 2-fold duplication and allow for early termination if no meaningful improvement (i.e., increase in the objective functional value) is detected.

References