Advances in Default Detection and Early Warning

Authors
Anamaria Pieschacon
Doug Dwyer

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Moody's Analytics recently embarked on a journey to enrich our corporate credit risk model offering. We are extending coverage, now scoring 400+ million firms worldwide, and increasing the accuracy of the existing metrics by combining them with additional risk signals. While preserving comparability, each probability of default (PD) opportunistically takes advantage of all applicable data. Useful information includes financial statements, equity market information, and payment behavior data. The PDs are designed to seamlessly integrate new information into the PDs as the information arrives.

The new framework also supports transforming the signals into actions. For instance, a quadrant-based early warning system identifies firms in need of attention due to recent risk profile changes. A framework for trade credit limits determines a buyer's appropriate limit based on both its size and its credit risk.

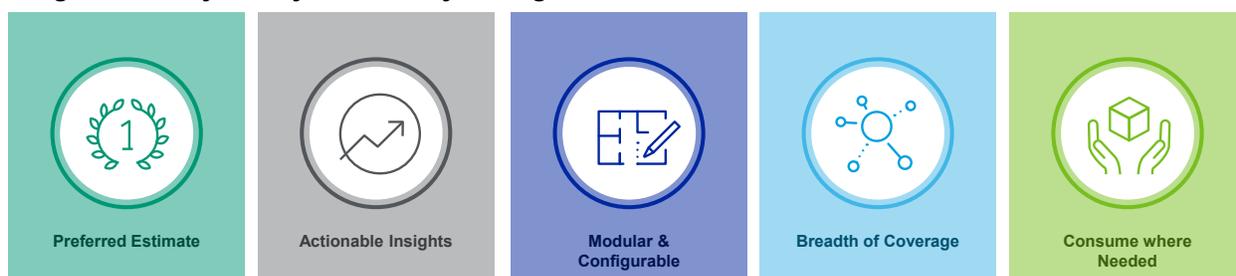## Table of Contents

TABLE OF FIGURES

Introduction

This document discusses the development of Moody's Analytics next-generation corporate credit risk solution. We launched this initiative with the purpose of enhancing our highly regarded models, leveraging innovative data and state-of-the art modeling techniques.

The goal is to not only enrich the existing model suite, but at the same time expand coverage to every public and private company in the world. We can now pre-score the entire BvD Orbis universe (which comprises 400+ million firms worldwide), including companies for which we do not have financial data. Our newly developed Financials-Based Benchmark model delivers a benchmark probability of default (PD) for any company in the world for a given country, industry, and size.

As Figure 1 shows, our redesign has several key features. First, we expand the sources of information considered to provide the best possible PD. We embrace conventional and unconventional data, applying a combination of traditional methods and novel techniques based on new information sources to strengthen the insights provided by our model suite.

Figure 1. Moody's Analytics PD - key redesign features

| Preferred Estimate | Actionable Insights | Modular & Configurable | Breadth of Coverage | Consume where Needed |

The accuracy gains from including additional components in the PD estimation vary across firm types. Heavily-traded public firms enjoy market efficiency, so alternative data provides limited value-add in the calculation of a risk metric. However, some public firms' stocks are thinly traded. For these, financials combined with alternative data may be informative in assessing credit risk. Similarly, for small businesses and medium-sized enterprises (SMEs), we can improve the accuracy of our private PD models by incorporating payment information. For SMEs whose financial data may be unavailable, our new Benchmark model, combined with payment data and news information (when available), provides valuable insights into their probabilities of default. Finally, in some circumstances, not all financial statement information is readily available. For example, sometimes the balance sheet is publicly available, while the income statement is not. In these instances, we augment the financials-based RiskCalc™ framework to produce a PD using limited information.

Second, the new solution provides tools that turn analytics into actionable insights (Figure 2). We provide trade credit limits, which offer guidance for corporates on how much credit to extend based on three risk-appetite levels. Our solution also includes a reimagined Early Warning System, which uses intuitive metrics that signal alerts to help users identify when to worry about deteriorating credit risk. Users can also obtain stressed PDs by projecting credit risks based on alternative macroeconomic scenarios. They can compute scenario-weighted PDs, making use of Moody's Analytics Economic Scenarios. Clients can also access our climate scenarios and project credit risk based on climate scenarios for both public and private firms. They can combine these insights with ESG ratings to obtain a much broader risk assessment than before.

Figure 2. Actionable insights

- Trade Credit Limits
- Early Warning
- Stressed PDs
- Climate-Adjusted PDs
- ESG Ratings

(center: Moody's Analytics PD)
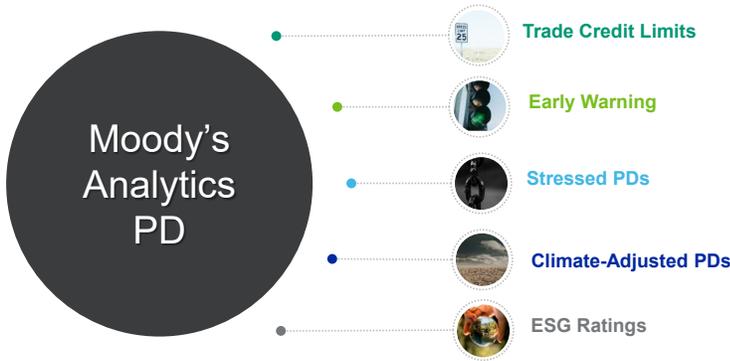
Third, we adopt a modular modeling approach (Figure 3). We build upward, starting with two conceptually coherent anchor models. First, for publicly traded firms, we use Moody's Analytics CreditEdge™ EDF™ (Expected Default Frequency) model, in which default occurs when a firm's liabilities exceed the market value of its assets. The second is the RiskCalc EDF model suite, which relates the financial position and performance of a firm (as measured by its financial statements) to default risk. We enhance these measures with other sources of timely data, such as behavioral payment and alternative data (ongoing research), and we optimally assemble them into our preferred risk metric. This assembled score, with the underlying optimal model weights, constitutes our updated house view. This view is designed to use whatever informative data is available.

Our preferred PD is comparable across different types of firms. Consequently, it is useful to rank-order an entire portfolio by credit risk. Thus, there will be no need to base risk assessment for small firms on behavioral information alone, or for medium sized-enterprises to rely on financial information alone. For all firm types, we make an optimal use of all types of available and relevant signals.

Our modular solution is also configurable; clients can use our best estimate, or they can customize it and use a subset of modules, for instance reverting back to a PD generated from RiskCalc.

Figure 3. Modular model architecture
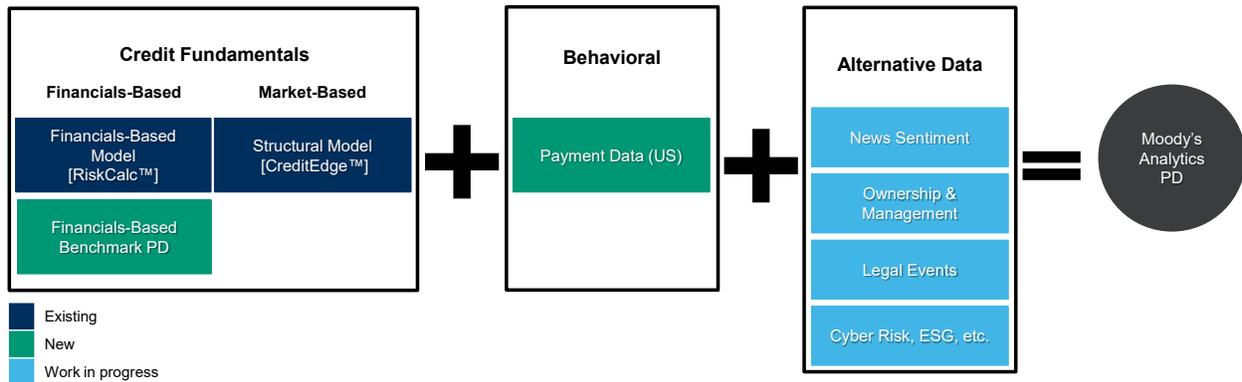
Fourth, our scoring framework covers the entire universe of firms from Orbis, from large financial and non-financial entities to SMEs. We do so with varying degrees of firm-level information. For private firms with financial data, we score them using RiskCalc. For firms without financial information, we have developed the Financials-Based Benchmark PD model that leverages our internal data assets and provides a benchmark risk metric by location, industry, and company size. Using these modules, our solution pre-scores the entire Orbis dataset (400+ million companies and growing), with the ability to score any firm worldwide.

Fifth, we aim to improve users' workflows, not only by pre-scoring the Orbis dataset, but also by helping them consume the content where it is needed. The framework is available via Moody's Analytics solutions, API, or data service (Figure 4).

Clients can use the API to connect our data, analytics, and tools to their data and then deliver results to their users, along with other insights. They can also obtain portfolio views, metrics, and other tools directly via our other products designed for loan origination, portfolio monitoring, transfer pricing and trade credit. We are also exploring additional channels to deliver content in the future.

**Figure 4. Putting insights in the right place, at the right time.**

| Directly within your existing tools | Through our Products | Via third party channels |

Taking Advantage of Available Data

A key component of this initiative is that the risk metric utilizes any and all available relevant information. If a firm's financial information is missing, for instance, we estimate it such that it neither penalizes nor rewards the firm for the missing information.

Figure 5 illustrates how we deal with missing data within the financial statement-based models. When we have enough data to construct all the financial ratios needed to estimate the full-financials model, we can score the company using the standard RiskCalc model suite. But if the financials are incomplete, we use either a limited financials adjustment (which relaxes the input requirements of RiskCalc to a minimum standard) or the Financials-Based Benchmark model that makes use of the data available for similar firms. The Benchmark model is already available, and the Near-Neighbors adjustment within RiskCalc will be released soon. In the meantime, those firms that fall under the Near-Neighbors adjustment category can already be scored using the Benchmark model.

Relaxing the input restrictions for RiskCalc using the Near-Neighbors approach allows us to score an additional 12 million firms in Orbis, and the Benchmark model provides estimates for 350+ million firms.

Figure 5. Refining and expanding coverage of financials-based PD models
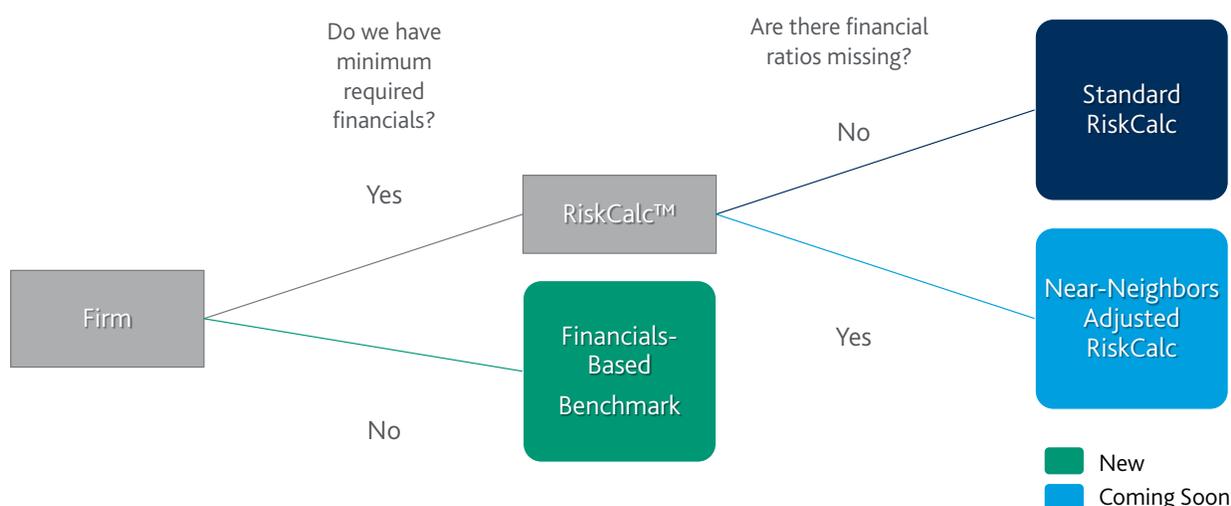
Coming Soon: Near-Neighbors Adjusted RiskCalc

When financial statements are incomplete, and we cannot construct every financial ratio needed using the standard RiskCalc models, we can utilize a forthcoming adjustment using the Near-Neighbors approach. As mentioned, this improvement allows us to relax constraints and score millions of additional firms in the Orbis dataset.

In previous RiskCalc versions, when inputs were missing, we would impute them, based on region and sector means. However, we would not compute an EDF measure in the absence net income (a key income statement item). We were reluctant to relax the constraint on the minimum required inputs, as mean imputation would systematically underestimate default probabilities. The new Near-Neighbors framework utilizes machine learning techniques to generate an unbiased EDF estimate of what the EDF would be if all of the information was actually provided. This innovation expands the scope of RiskCalc models so that they produce an estimated PD using any available financial statement information (e.g., includes balance sheet but not an income statement). Importantly, the resulting PD will neither systematically understate nor systematically overstate the exposure's risk.

Financials-Based Benchmark Model

We have also developed a model that estimates a benchmark PD when no financial statement information is available. With this model, we obtain a proxy PD for 350+ million firms, using characteristics of similar firms.

We used data covering more than 100 countries, 24 million firms, and 100 million financial statements across the world for model development. While we apply a consistent model framework globally, the estimation is split into two modules: U.S. and non-U.S. The U.S. module is estimated on the U.S. data from Moody's Analytics Credit Research Database (CRD), while the non-U.S. module leverages the breadth and depth of the Orbis dataset universe.

The target variable of both modules is the one-year through-the-cycle PD (TTC PD), known as Financial Statement Only RiskCalc EDF (FSO EDF) measure (see Figure 6). We tested different possible drivers for each module. When selecting a model, we look for specifications that perform well when we have data, and also on the simulations where we set some of the data as missing. This practice allows us to assess model performance as if the data were actually missing. We considered several methodologies, including complex machine learning models. In the end, we opted for a linear framework with non-linear transformations of the features, to maximize interpretability and transparency while minimizing the loss of power.

As Figure 6 shows, in addition to sector and region, we selected size and leverage as explanatory variables for both datasets. In many instances, we may not know a firm's total assets or leverage. Imputation based on sector and region alone does not facilitate discrimination across firms. Therefore, we developed two sub-models to estimate total assets and leverage using information on sales, number of employees, and profitability, making use of the fact that at least one of these metrics is present across virtually all statements. By doing so, we refine the model inputs and provide an accuracy boost to the Benchmark model.

Figure 6. Benchmark model schematic

$$FSO\ EDF_{1yr} \sim \underbrace{1 + f(size) + f(leverage)}_{TTC\ PD_{1yr}} = \begin{matrix} Financials\ Based \\ Benchmark\ PD \end{matrix}$$
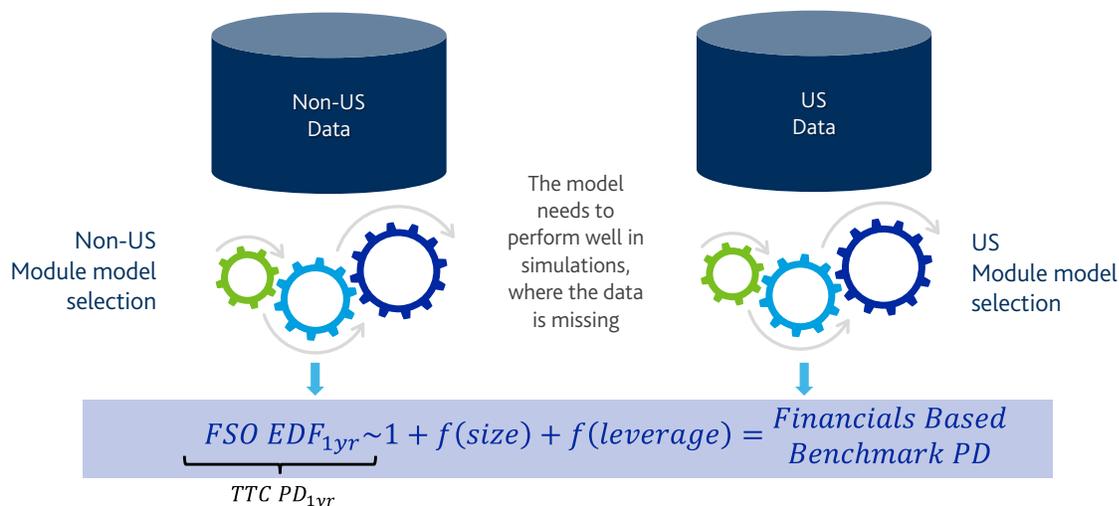
For the U.S. module, we model PD by census region and industry clusters. For the non-U.S. module, we model PD by country and industry clusters. To deal with clusters of limited number of observations, we also designed tiers of granularity, so that small country and industry clusters fall back on country group and industry, and ultimately global industry.

Coming Soon: Behavioral Payment PD Model

For the U.S., we are developing a behavioral PD model that exclusively uses trade payment data. With our recent acquisition of Cortera, we can now look into the payment behavior of 3.2+ million firms with trade information and generate a payment-based PD measure.

This model does not include financial information. Instead, it uses trade payment data, which provides meaningful information about credit risk for small- and medium-sized firms, for whom we may or may not have financial statements.

In absence of the financial statements, a model derived exclusively from trade behavioral data provides good predictive power in terms of credit risk. The Payment PD Model also adds predictiveness to the PD estimated from RiskCalc and to the Benchmark PD, once combined with one or the other in an ensemble (Figure 7). The additional predictive power is higher for smaller companies, where financial information is either scarce or of lower quality, relative to larger companies.

Unlike annual financial statements, payment information is updated monthly. Therefore, it can help detect early credit quality deterioration. In fact, the payment information becomes more useful in assessing credit risk as the financial statement information ages throughout the year.

Coming Soon: Ensemble

The Moody's Analytics PD is an aggregation of all the sources of credit risk assessable for a given borrower. The aggregation is customized to each firm. For example, the weight on financial information relative to behavioral information will be lower for a small private business than for a medium-sized enterprise. For larger private firms, the optimal ensemble will still be a mixture of the financials-based PD from RiskCalc and the Payment PD, but with a decreasing weight on the latter.

The ensemble weight on the Public Firm CreditEdge EDF metric is likely close to 100% for non-financial public firms in markets with largely efficient financial markets. For less liquid ones, other signals may provide additional predictive power. That being said, in the first version of the Moody's Analytics PD, we assign a 100% weight to the CreditEdge EDF metric for all public firms. We may revisit this in the future, depending on the results of ongoing research.

Figure 7. Combining financials-based and payment-based signals
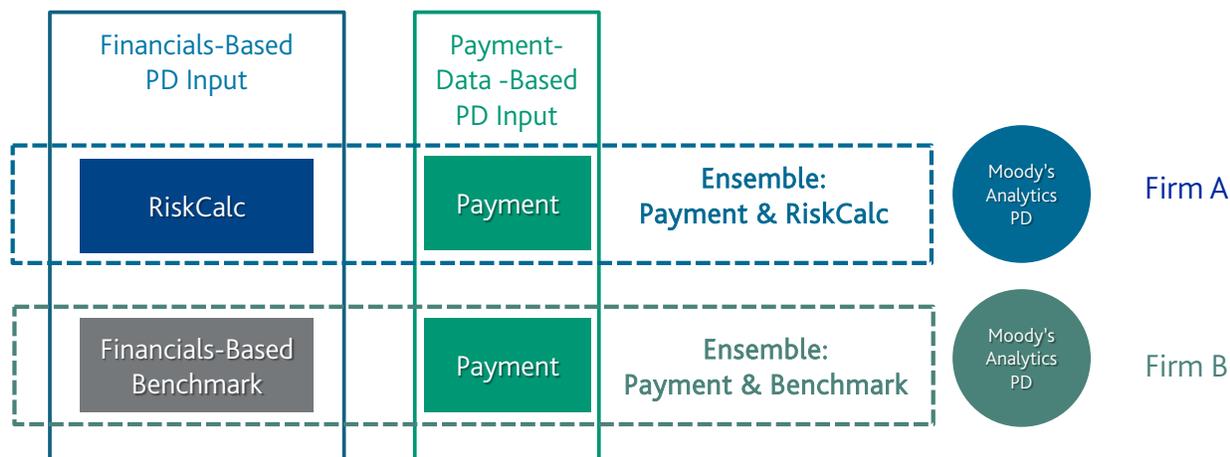
Figure 7 provides an example of how we obtain our preferred PD metric for two private firms: A and B. For firm A, we can obtain a financials-based PD using RiskCalc. For Firm B, on the other hand, we do not have financials, so we use the Benchmark PD. In each case, we combine the financials-based PD input with the payment data-based PD model, and we obtain a unique preferred PD for each of them. We achieve this result using a different set of weights for each firm.

Actionable Insights

We are currently developing multiple actionable insights from our preferred PD. As shown in Figure 2, these include trade credit limits, an updated early warning system, climate-adjusted PDs, and stressed PDs, among others. We next explain our approach to calculate trade credit limits and the new early warning system. Other applications will be part of future releases.

Trade Credit Limits

Many wholesale marketer/traders have traditionally measured and monitored credit risk by assigning credit limits to counterparties. These credit limit assignments are made based on pre-approved tables, qualitatively based on past experience. Those tables show the maximum suggested limit for customers of a given size (sales or net worth) and creditworthiness (internal or external scoring). Those tables typically suggest higher limits for larger customers and lower limits for heightened credit risk customers, sometimes allocating zero unsecured credit to firms in or near financial distress.

Our Trade Credit Limit Model provides a quantitative approach for obtaining a maximum credit limit, which can be extended to a buyer under a given set of assumptions. The model uses a framework to calculate a credit limit based on the net value added in a transaction. The model considers appropriate assumptions of the lending parameters depending on the country of the buyer. The financials-based PD from RiskCalc, as well as firm sales, are key parameters for the Trade Credit Limits Model. For firms with few or no financial statement coverage, we use PD and Sales Benchmark models to estimate the credit limit levels, depending on the appropriate country, region, and industry, as covered by BvD Orbis across 260 countries and 61 industries globally.

Our quantitative approach calculates maximum economic credit limits based on expected profit. The intuition behind the economic limits is straightforward: The margin on the trade flow must cover the expected loss due to non-payment and cover the cost of capital required to withstand those losses. That cost is not just the cost of financing; it also includes a premium for undertaking credit risk. At the maximum credit limit, profits created are offset by the expected default losses and charges for using firm capital, setting the net economic value equal to zero.

Figure 8. Credit limit calculation.
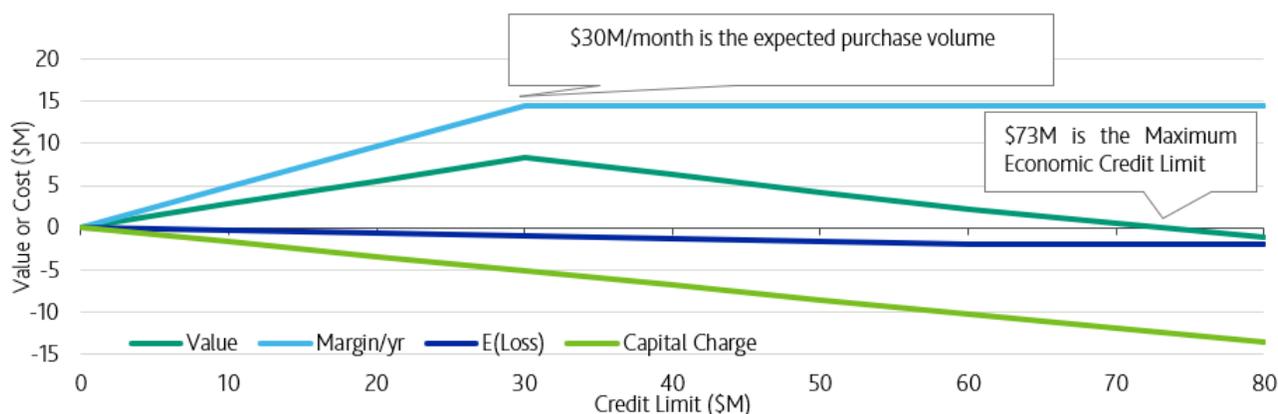
Figure 8 describes this intuition. The customer in this example is expected to purchase $30M/month worth of goods/services for their business needs. The annual margin increases until the customer buys all they need. Beyond that point, additional credit limit does not provide any added value. The net value to the seller also includes the costs—expected credit loss, capital charges, and risk premium—which increase with the credit limit used. Our quantitative approach calculates maximum economic credit limits based on expected profit. The intuition behind the economic limits is straightforward: The margin on the trade flow must cover the expected loss due to non-payment and cover the cost of capital required to withstand those losses. That cost is not just the cost of financing; it also includes a premium for undertaking credit risk. At the maximum credit limit, profits created are offset by the expected default losses and charges for using firm capital, setting the net economic value equal to zero.

Figure 8 shows the maximum credit limit suggested is $73M, which makes the net value to the seller equal to zero. At this point, the seller remains indifferent between walking away from the relationship and maintaining this credit limit. The seller prefers that the buyer keeps their credit limit at $30M, which is one month outstanding. This level maximizes the value of the relationship to the seller, as seen in the peak of the dark green line. But the seller is willing to extend this limit (up to $73M), to maintain the relationship with the buyer.

When all inputs are available, the computation of this credit limit is straightforward. But for a sizeable portion of the dataset, this calculation cannot be executed directly because financial statements are missing. To overcome this challenge, we developed a sales benchmark, which allows us to produce a proxy invoice amount.

We compute credit limits for all non-financial firms in the Orbis dataset with Active or Unknown status. This implies that we do not calculate trade credit limits for firms that are inactive due to bankruptcy, liquidation, or dissolution. We exclude financial firms from the sample, as the trade credit limit concept does not apply to banks or insurance companies. Our model implementation produces conservative, standard, and aggressive estimations to utilize under different scenarios. We produce these estimates by varying the model's margin assumptions.

Figure 9 shows the three trade credit limit levels, splitting up the data by deciles. On the top row, we show the credit limit by sales and PD decile, and, on the bottom row, we show the credit limits-to-sales ratio, once again by sales and PD decile. As expected, credit limits are higher for larger firms, as identified by revenue. The sales/revenue deciles are based on the revenue reported in USD. The PD buckets are based on RiskCalc PDs. Higher risk firms are offered a lower credit limit, which is intuitive. The credit limit relative to the overall sales of the firm also decreases with increasing credit risk of the firm. All else equal, larger firms also receive a lower credit limit relative to their size, as they tend to work with a larger number of suppliers than smaller firms.

Figure 9. Average trade credit limit by bucket (global).
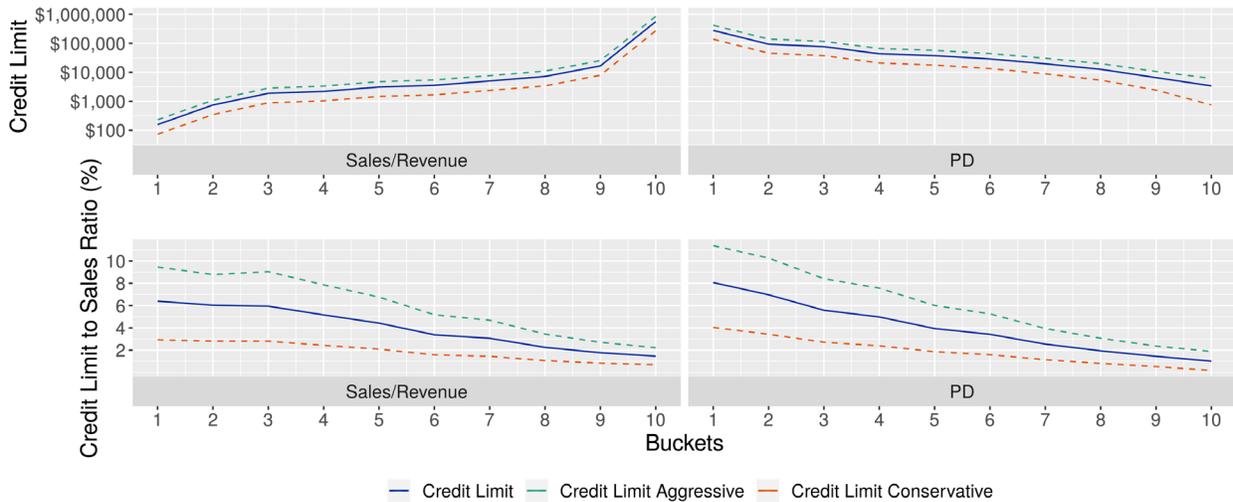
Early Warning System

We are reimagining the way we define our early warning system by adopting a quadrant approach with intuitive and flexible settings, that will allow users to customize it based on their needs and define the criteria to get an accurate signal of when to worry. The quadrant classifies risk in absolute and relative terms. Conceptually, the x-axis represents the assessment of absolute risk, while the y-axis represents the change in the relative risk level. As shown in Figure 10, on the x-axis, we measure the firm's PD distance to trigger, and on the y-axis, we rank firms in terms of the 12m change in their PD static implied rating. When the PD of the company is below its trigger, the firm will be classified as either low (green) or medium (yellow) risk. If the trigger threshold is crossed, the firm is classified as either high or severe risk (orange or red). On the y-axis, an implied rating improvement translates into a positive change, moving upwards. Conversely, an implied rating deterioration renders a downward shift.

Figure 10. Early Warning Quadrant - A simplified schematic
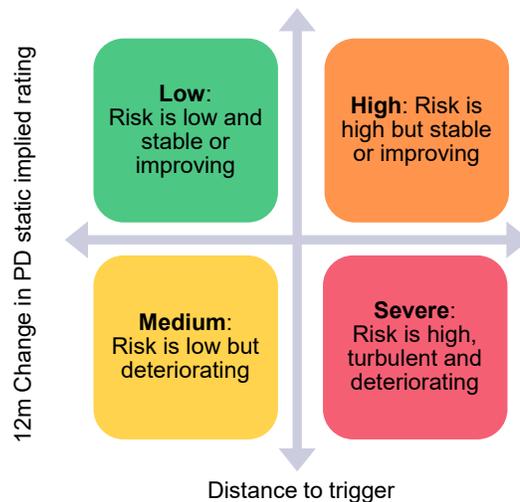
The trigger is a PD threshold constructed to identify the riskiest firms. A higher trigger leads to fewer firms being flagged and consequently a lower false positive rate (type 1 error). On the other hand, a lower trigger increases the number of flagged firms and thus reduces the false negative rate (type 2 error). Therefore, the trigger design needs to deal with this

error tradeoff and reach an optimal balance between the two types. Ideally, it should also provide flexibility to reflect the user's risk appetite.

The new EWS trigger methodology offers an intuitive and transparent framework. We are making the trigger a function of a PD percentile, based on peer group. Ideally, we would construct triggers at the most granular peer group level. In practice, however, at a given time, a peer group defined at great granularity may not have enough firms that could represent the "true" underlying PD distribution for this peer group. To resolve this issue, we employ a Bayesian-adjusted trigger to leverage more information from a more broadly defined peer group when data for a more granular peer group are not sufficiently representative. We classify peer groups by a combination of country and industry.

We introduce ample flexibility into the trigger setting, so that users can adjust the hit rate (positive signal rate) as needed. To accommodate clients' needs for flexibility in the hit rate (Positive Signal Rate), the Next Generation EWS' trigger methodology offers a set of triggers, with different through-the-cycle hit rates, ranging from 60% to 95%. In our default setting, the hit rate will have a cyclical adjustment that makes it procyclical, so that users can flag more names during bad times and fewer names in quiet times.

Conclusion

We recently launched an initiative to enrich our corporate credit risk model offering. We have extended coverage and increased the accuracy of existing metrics by combining them with additional risk signals. In the coming months, we will continue to improve and expand our solution, by implementing novel approaches and exploring additional conventional and unconventional sources of data. We will also continue to roll out tools that turn analytics into meaningful actionable insights.