creditrisk monitor®

supplychain monitor™

Whitepaper

How Crowdsourced Intelligence Enhances the FRISK® Score

The CreditRiskMonitor® platform continues to improve the quality of the information it provides subscribers. This white paper describes how a data source—crowdsourced intelligence—is used to improve the accuracy of the FRISK® Score. To the author's knowledge, CreditRiskMonitor® is the first and only financial risk analysis platform to employ this innovative method.

The crowdsourced intelligence is derived from the aggregate behavior of subscribers as they interact with the CreditRiskMonitor® platform. These subscribers play influential roles in the daily commerce of many of the world's largest corporations, making credit and procurement decisions that affect trillions of dollars in purchase and sale transactions every year. This is a highly sophisticated user base: nearly 40% of the Fortune 1000, along with many of the largest public and private firms, use the platform.

Trade credit is a far more significant source of business financing than is commonly recognized. According to the Federal Reserve, trade credit was the third-largest liability of U.S. non-financial corporations—approaching $2 trillion in early 2016—after bonds (approximately $5 trillion) and loans (approximately $2.6 trillion). As a result, procurement departments play a critical role in business-to-business sales.

By 2025, trade credit, bonds, and loans reached $4.3 trillion, $7.8 trillion, and $5.4 trillion, respectively. Research conducted by CreditRiskMonitor® shows that the usage patterns exhibited by subscribers when evaluating a business on the platform are correlated with their level of concern about that business. When analyzed across a large subscriber base, these aggregate usage patterns provide additional insight into a company's failure risk. Incorporating this unique crowdsourced intelligence into the FRISK® Score algorithm significantly enhances its performance. Overall accuracy is back-tested at 96%, with particularly meaningful improvements in scoring accuracy for the highest-risk companies.

Data in the New FRISK® Score

A large amount of data over many years is required to build effective crowdsourcing-based models. This data trains analytical systems to detect usage patterns that are indicative of strong or weak financial performance. The CreditRiskMonitor® platform is uniquely positioned to leverage the "wisdom of the crowd" for evaluating business financial risk for several key reasons.

Highly structured website

The platform's highly structured website enables subscribers to navigate to specific pages in order to conduct deeper analysis of particular types of information. The crowdsourcing methodology does not rely on the specific financial or descriptive content displayed on these pages—some of which is incorporated into other components of the FRISK® Score—but instead focuses solely on subscriber navigation behavior. In other words, the model evaluates which pages subscribers choose to visit, not what information is contained on those pages. This consistent site structure is essential to enabling reliable crowdsourced analysis.

Detailed user session records

To support ongoing product development and enhancement, detailed records of subscriber user sessions have been maintained. This long-term data retention policy has resulted in a rich historical dataset suitable for advanced modeling. Most importantly, the dataset reflects the behavior of a highly sophisticated subscriber base, making observed usage patterns particularly informative with respect to assessing company risk.

How Website Activity Reveals Information about Credit Risk

Subscribers of the CreditRiskMonitor® platform represent a highly specialized group. Thousands of users, many employed by the world's largest corporations, are responsible for protecting their organizations from credit risk. Other subscribers, also employed by large corporations, monitor the ability of vendors to meet supply obligations. Both roles require a deep understanding of the financial health of the businesses under review.

These professionals rely on the platform to assess the financial risk of publicly listed companies, while also drawing on additional sources of information not available to the general public. Insights derived from industrial relationships, sales and service teams, operations units, and other internal resources often inform their evaluations. As a result, subscribers function as financial "gatekeepers," exerting significant influence over the volume and continuity of business conducted with specific customers or suppliers.

To support recommendations and decision-making, subscribers review information presented across the platform. Unlike many other credit data providers, the platform organizes a broad range of financial and risk-related data in a highly structured format designed to facilitate efficient analysis.

Evaluation typically begins with a company's snapshot page, which provides a concise, high-level view of risk. When additional insight is required, such as understanding the drivers behind a particular FRISK® Score or rating, subscribers navigate to more detailed pages that offer deeper analytical context.

When analyzed in aggregate, distinct usage patterns emerge based on the level of concern regarding a company's financial condition. Low concern is often associated with limited review activity, focused primarily on the snapshot page. Higher levels of concern prompt more extensive exploration across multiple data views to validate conclusions and support internal discussion. These differentiated activity patterns provide meaningful signals that form the foundation of the crowdsourced intelligence incorporated into the FRISK® Score model.

How the FRISK® Score Achieves High Accuracy

A key reason why the FRISK® Score achieves strong predictive accuracy is its integration of multiple data sources. Since its introduction in 2007, the model has incorporated three primary components: stock market information, financial statement data and agency ratings.

The model is designed to operate using any available subset of these components. Each component independently contains information related to failure risk and can support a standalone credit model. While these data sources often move together, they do not do so perfectly. The model's ability to reconcile divergence among components is central to the FRISK® Score's superior performance.

Stock market data, financial statements, and agency ratings each offer distinct perspectives on bankruptcy risk. Combining these perspectives improves overall predictive power while mitigating the limitations of individual data sources.

Incorporating Crowdsourced Website Usage Information

Although the crowdsourced usage data differs substantially from traditional financial and market data, its integration into the FRISK® Score follows a similar modeling framework. A primary requirement is the availability of extensive historical data for training. The hierarchical structure of the CreditRiskMonitor® platform has provided many years of subscriber activity data.

Model development involves a trade-off between emphasizing recent information and incorporating longer historical periods that include a broader range of economic conditions. In addition, ongoing enhancements to the platform, including the introduction of new data sources and changes to page content, must be considered. For these reasons, crowdsourced usage data is incorporated beginning in 2010. Other components of the FRISK® Score are back-tested using data starting in 2003.

Aggregate behavior is central to the effectiveness of crowdsourced intelligence. Only usage patterns consistently exhibited by a large number of subscribers provide meaningful signals. As a result, this information is most valuable for companies that attract sustained attention from a significant portion of the subscriber base. Through this enhancement, subscribers effectively contribute anonymously to improved risk assessment for each other.

Why aggregate behavior is essential for reliable crowdsourcing

As with the previous versions of the FRISK® Score, the enhanced model was trained using a historical dataset of bankruptcies. Model parameters were adjusted using standard statistical techniques to optimize predictive performance. The methodology employed was static pool analysis, in which the full historical period used to build the model is divided into a series of 12-month pools. All relevant information available at the beginning of each pool is collected for a set of active companies—those that have not previously filed for bankruptcy. Companies that filed for bankruptcy during the corresponding 12-month period are then identified and flagged.

Using data from all the pools, the training process identifies website usage patterns that are most predictive of bankruptcy. Analysis indicates that predictive value is derived not from the absolute number of user "clicks," but from the patterns of activity observed. Specifically, the relative frequency with which subscribers visit certain pages compared with others provides meaningful signals of financial stress when evaluated in aggregate.

Evaluating model performance

Once the model was developed, static pool analysis was again used to evaluate its performance. In this validation phase, the model's initial score for each active company at the start of a static pool was recorded, and outcomes were then assessed based on which companies filed for bankruptcy during the pool period. As with previous versions of the FRISK® Score, two validation tests were conducted. One involved generating a set of ROC[1] curves, and the other evaluated the model's ability to correctly classify companies that ultimately filed for bankruptcy.

In their simplest form, credit models can be used in a binary manner by selecting a score threshold that classifies companies as either "low-risk" or "high-risk." For any given threshold two performance measures can be calculated: the false positive[2] rate and the true positive rate. In this context, the false positive rate represents the proportion of all scored companies that did not file for bankruptcy during the relevant pool period but were classified as high risk. With any model, minimizing this rate is important, as false positives represent viable businesses that may be unnecessarily rejected or subjected to additional scrutiny. Operational considerations also favor a low false positive rate, as elevated risk scores often trigger additional analysis and delay decision-making.

The true positive rate represents the proportion of companies that filed for bankruptcy during the static pool period and were correctly classified as high risk at the beginning of the pool. For obvious reasons, this rate should be maximized in any model.

Finding the optimal risk threshhold

Varying the score threshold produces different combinations of false positive and true positive rates. A very low threshold yields a low false positive rate, but an unacceptably low true positive rate, resulting in missed bankruptcy warnings. Conversely, a very high threshold increases the true positive rate but also produces an excessive false positive rate, leading to the rejection of many financially sound companies. At either extreme, the economic value of the model is limited.

By systematically varying the score threshold or operating point (the "O" in ROC), a range of true and false positive rate pairs is generated, each corresponding to a point on the ROC curve. This process enables identification of an optimal operating point for a model. The selection of FRISK® Score= 5 as the threshold separating high- and low-risk companies was informed by this kind of analysis.

Where the FRISK® Score makes the biggest difference

In most cases, the enhanced FRISK® Score remains unchanged from the previous version. Several factors contribute to this outcome. First, companies that attract meaningful crowdsourced data represent only a subset of all companies scored. Second, the original FRISK® Score based on stock market data, financial statements, and agency ratings already demonstrated strong predictive performance. In many instances, website usage patterns reflect the same underlying information captured by traditional inputs, resulting in no change to the score when adding crowdsourced intelligence.

However, for a significant subset of companies that are of particular interest to subscribers, crowdsourced usage data provides additional insight not fully reflected in market prices, financial disclosures, or ratings data. For this group, the enhanced model produces meaningful score differentiation.

---

[1]ROC stands for "Receiver Operating Characteristics" and the name comes from the age of analog communications where this technique was used to evaluate a receiver's ability to discriminate signal from noise.

[2]Ironically, a "positive" here is a bankruptcy; positive just means that the model correctly predicted the event (i.e., bankruptcy) you are looking for.

Figure 1 shows the ROC curve for companies whose FRISK® Scores differ between the enhanced and previous models. For this population, the enhanced model demonstrates a significant improvement in performance.

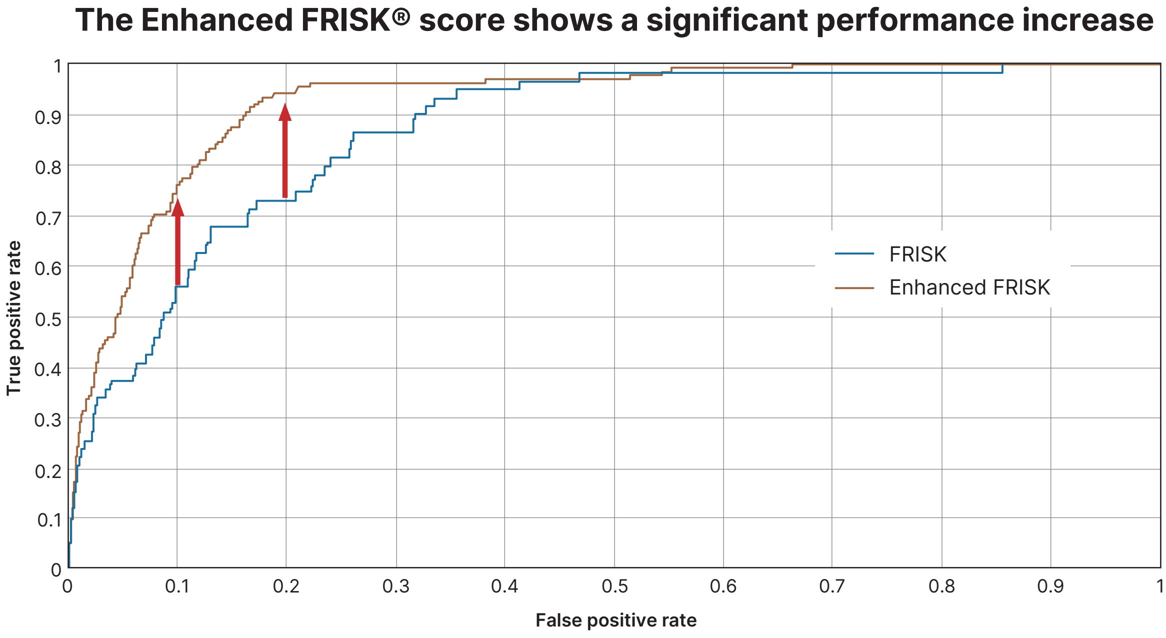

Figure 1: ROC curve for businesses for which the enhanced FRISK® score changes

Further insight into the value of crowdsourced data is provided by examining the distribution of bankrupt companies across FRISK® Score buckets. Figure 2 displays the cumulative FRISK® Score distribution for all bankrupt companies in the validation dataset. Each bar represents the percentage of bankrupt companies that received a given FRISK® Score or lower at the start of each static pool.

Capturing more failtures earlier

For FRISK® Score= 5, the cumulative percentage is 96%, indicating that 96% of all bankruptcies in the validation dataset received a score of 5 or lower. In other words, the enhanced model provided early warning for 96% of public company bankruptcies in the United States. While this represents only a one-percentage-point improvement over the previous model's 95% result, it corresponds to a 20% reduction in the number of bankruptcies previously missed.

Concentrating risk in the highest risk buckets

An equally important improvement is the enhanced model's ability to more accurately assign companies to new higher-risk score buckets. As shown in **Figure 2**, the most pronounced changes occur in the lowest (highest-risk) buckets. Although the total proportion of bankrupt companies classified as high risk increases by only one percentage point, a larger share of bankruptcies is concentrated in the most severe risk categories.

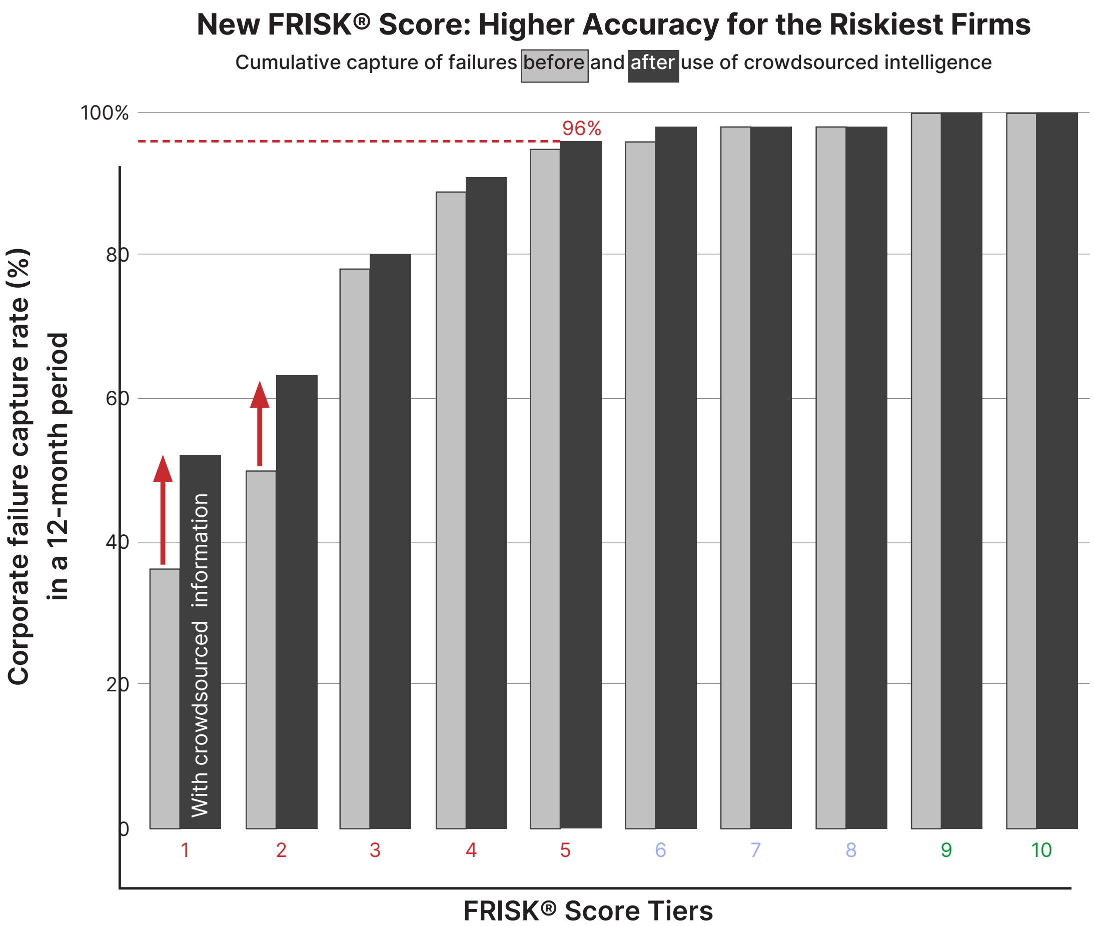

Figure 2: enhanced FRISK® Score bucket distribution of bankrupt

Specifically, 50% of all bankruptcies are captured in the FRISK® Score = 1 bucket, compared with 35% under the previous model. This represents a substantial improvement in risk stratification, enabling subscribers to more effectively prioritize the highest risk companies.

Why Crowdsourced Intelligence Strengthens Credit Risk Decisions

The enhanced FRISK® Score has been in operation since June 11, 2016. It leverages the aggregate website usage activity of our subscribers to provide a better estimate of bankruptcy risk for a small group of companies that are of keen interest to our subscribers. The enhancement impacts the scores of U.S. businesses that are of interest to a significant fraction of our subscribers. The scores produced have the same relationship between each FRISK® Score bucket and probability of bankruptcy.