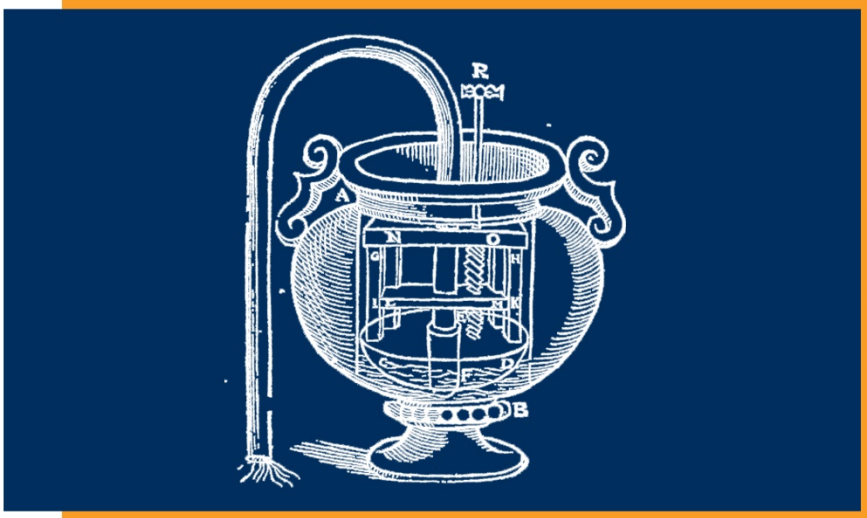
CNR-IRCrES Working Paper

HERA-PRO
Prescriptive Rating Optimization Under Explainable Artificial Intelligence



8/2026

Greta Falavigna

CNR IRCrES
Istituto di Ricerca sulla Crescita Economica Sostenibile

ISSN (online) 2421-7158
ISSN (print) 2421-6798

Direttore Elena Ragazzi

Direzione CNR-IRCrES
Istituto di Ricerca sulla Crescita Economica Sostenibile
Strada delle Cacce 73, 10135 Torino, Italy
Tel. +39 011 3977612
segreteria.ircres@cnr.it
www.ircres.cnr.it

Sede di Roma Via dei Taurini 19, 00185 Roma, Italy
Tel. +39 06 49937809 / Fax +39 06 49937808

Sede di Milano Via Corti 12, 20121 Milano, Italy
Tel. +39 02 23699501 / Fax +39 02 23699530

Sede di Genova Corso Ferdinando Maria Perrone 24, 16152 Genova, Italy
Tel. +39 010 6598798

Comitato di Redazione
Grazia Biorci, Igor Benati, Antonella Emina, Serena Fabrizio, Lucio Morettini, Susanna Paleari, Anna Perin, Emanuela Reale, Secondo Rolfo, Andrea Orazio Spinello, Isabella Maria Zoppi.

https://www.ircres.cnr.it/produzione-scientifica/

I Working Paper del CNR-IRCrES rappresentano il punto di vista degli autori e non riflettono necessariamente quello dell'Istituto nel suo complesso.

The Working Papers published by CNR-IRCrES represent the views of the respective author(s) and not of the Institute as a whole.

CNR-IRCrES Working Paper 8/2026

ottobre 2026 by CNR-IRCrES

HERA-PRO: PRESCRIPTIVE RATING OPTIMIZATION UNDER EXPLAINABLE ARTIFICIAL INTELLIGENCE

GRETA FALAVIGNA

CNR-IRCrES, Consiglio Nazionale delle Ricerche, Istituto di Ricerca sulla Crescita Economica Sostenibile, strada delle Cacce, 73 - Torino 10135, Italy

corresponding author: greta.falavigna@cnr.it

ABSTRACT

This paper introduces HERA-PRO, a prescriptive extension of the Hierarchical Artificial Intelligence Rating Architecture (HERA, Hierarchical Engine for Rating Assessment) designed to support corporate rating improvement strategies through Explainable Artificial Intelligence (XAI). Building on the predictive framework proposed in Falavigna (2026a) and the explainability layer developed in Falavigna (2026b), the proposed approach moves beyond rating forecasting and focuses on identifying actionable interventions capable of increasing the probability of obtaining a more favorable rating class. The methodology combines hierarchical Random Forest classifiers and counterfactual optimization procedures within a unified decision-support framework. Firms are first classified through a network of sector-specific and rating-specific predictive models. Subsequently, a prescriptive optimization layer is introduced to determine the minimum financial adjustments required to achieve a target rating outcome. The proposed procedure explores alternative counterfactual configurations of firms' financial indicators and identifies both single-variable and multi-variable adjustment strategies capable of producing rating upgrades. Unlike conventional rating forecasting systems, which primarily focus on prediction accuracy, HERA-PRO provides decision-makers with actionable recommendations by quantifying the minimum changes necessary to cross rating-classification thresholds. The framework therefore shifts the analytical perspective from "What rating will be assigned?" to "What should change to obtain a better rating?". By integrating prediction, explanation, and prescription within a coherent methodological framework, HERA-PRO contributes to the emerging literature on prescriptive artificial intelligence, counterfactual analytics, and decision-oriented machine learning for corporate credit risk assessment.

KEYWORDS: Prescriptive Analytics; Explainable Artificial Intelligence; Counterfactual Analysis; Corporate Credit Risk; Hierarchical Artificial Intelligence.

DOI: 10.23760/2421-7158.2026.008

How to Cite

Falavigna, G. (2026). *HERA-PRO: Prescriptive Rating Optimization under Explainable Artificial Intelligence* (CNR-IRCrES Working Paper 8/2026). Istituto di Ricerca sulla Crescita Economica Sostenibile. http://dx.doi.org/10.23760/2421-7158.2026.008

SUMMARY

1. INTRODUCTION 3
2. LITERATURE REVIEW 4
2.1. Prescriptive Analytics and Decision-Oriented Machine Learning 4
2.2. Research Gap and Contribution 5
3. THE HERA-PRO FRAMEWORK 5
3.1. From Prediction to Explanation and Prescription 5
3.2. The Prescriptive Rating Optimization Layer 6
4. METHODOLOGY AND RESULTS 7
4.1. Selection of the Reference Firms 8
4.2. Counterfactual Benchmark Construction 9
4.2.1. Individual Indicator Effect Analysis 10
4.3. Minimum Counterfactual Adjustment: Single-Variable (OVAT) Analysis 12
4.4. Multi-Variable Counterfactual Search 15
5. TRANSLATING COUNTERFACTUAL EXPLANATIONS INTO MANAGERIAL ACTIONS 19
5.1. From Financial Indicators to Managerial Levers 19
5.2. Multi-Variable Counterfactual Search: Results 21
5.3. Robustness check: Cross-Firm Evidence 23
6. GLOBAL COUNTERFACTUAL VALIDATION AND WEIGHTED ACTION SCORING 25
7. CONCLUSIONS AND FURTHER STUDIES 31
8. REFERENCES 33
9. APPENDIX A. FINANCIAL INDICATORS 34
10. APPENDIX B. DETAILED COUNTERFACTUAL RESULTS 35
11. APPENDIX C. COUNTERFACTUAL ANALYSES AND MANAGERIAL RECOMMENDATIONS (SECTION 5) 37

1. INTRODUCTION

Corporate credit ratings play a central role in financial markets, influencing firms' access to external financing, funding costs, investment opportunities, and relationships with stakeholders. Consequently, the development of accurate rating forecasting models has attracted significant attention from both academic researchers and practitioners.

Recent advances in machine learning have considerably improved predictive performance in rating assessment tasks. Nevertheless, high predictive accuracy alone is often insufficient for supporting managerial decision-making. Stakeholders increasingly require artificial intelligence systems that not only predict future outcomes but also explain the determinants of these predictions and provide guidance on how such outcomes may be improved.

To address these challenges, Falavigna (2026a) introduced the Hierarchical Artificial Intelligence Rating Architecture (HERA), a framework based on a network of interconnected Random Forest classifiers designed to predict one-year-ahead corporate ratings. The hierarchical structure allowed the decomposition of a complex multi-class rating problem into a sequence of specialized classification tasks, improving both predictive accuracy and interpretability.

Subsequently, Falavigna (2026b) extended the HERA framework through the integration of Explainable Artificial Intelligence techniques. By exploiting SHAP-based local and global explanations, the study identified the financial indicators most strongly associated with rating assignments, thereby increasing model transparency and facilitating the interpretation of predictive outcomes.

Despite these advances, an important question remains unanswered. Once the main rating drivers have been identified, how can this information be transformed into actionable recommendations for firms seeking to improve their future rating position?

This paper addresses this issue by proposing HERA-PRO a prescriptive artificial intelligence framework that extends the predictive and explanatory capabilities of previous HERA versions. Rather than limiting the analysis to forecasting and explanation, the proposed methodology focuses on optimization and decision support. Specifically, it evaluates how targeted modifications in key financial indicators affect the probability of obtaining a more favorable rating classification.

The proposed framework combines machine learning prediction models, explainable artificial intelligence techniques, and counterfactual optimization procedures within a unified architecture. Building upon the predictive and explanatory layers developed in previous versions of the HERA framework, HERA-PRO identifies the minimum set of financially plausible adjustments required to achieve a desired rating outcome. By integrating prediction, explanation, and prescription, the framework transforms rating forecasting into a decision-support tool capable of generating actionable recommendations for firms seeking to improve their future creditworthiness.

The contribution of this study is threefold. First, it bridges the gap between predictive and prescriptive analytics in corporate rating assessment. Second, it provides an interpretable mechanism for translating explainability results into operational recommendations. Third, it offers an AI-based decision support tool capable of quantifying the effectiveness of alternative financial strategies aimed at improving firms' creditworthiness.

The remainder of the paper is organized as follows. Section 2 reviews the literature on prescriptive analytics and decision-oriented machine learning. Section 3 presents the HERA-PRO framework, while Section 4 describes the methodological architecture and the counterfactual optimization procedure. Section 5 discusses the empirical results and illustrates how the proposed approach can support managerial decision-making by identifying financially feasible pathways towards rating improvement, while also demonstrating the scalability of the methodology. Section 6 concludes and discusses the managerial implications, limitations, and future developments of the framework. Three appendices complement the paper by showing financial indicators adopted

and reporting the detailed empirical results and the counterfactual recommendations generated by HERA-PRO[1].

2. Literature Review

2.1. Prescriptive Analytics and Decision-Oriented Machine Learning

Traditional machine learning approaches are mainly designed to generate accurate predictions. In recent years, however, increasing attention has been devoted to prescriptive analytics, a research stream that aims not only to predict future outcomes but also to identify actions capable of improving those outcomes. Within this perspective, predictive models become components of broader decision-support systems that assist users in selecting alternative courses of action (Bertsimas & Kallus, 2020).

The transition from predictive to prescriptive analytics represents a natural evolution of data-driven decision making. While predictive models estimate what is likely to occur, prescriptive approaches seek to determine which decisions should be adopted to achieve a desired objective. Bertsimas and Kallus (2020) argue that the ultimate value of machine learning lies not only in forecasting future events but also in supporting optimal decision-making processes through data-driven prescriptions.

More recently, the development of Explainable Artificial Intelligence (XAI) has further stimulated interest in decision-oriented machine learning. Among the emerging approaches, counterfactual explanations have attracted considerable attention because they provide actionable recommendations by identifying the minimum changes required to obtain a different prediction (Verma et al., 2024; Guidotti, 2024). Rather than simply explaining why a model generated a given outcome, counterfactual methods answer a more operational question: what should be changed to achieve a preferred outcome.

The growing interest in prescriptive machine learning has also stimulated the development of counterfactual explanations and algorithmic recourse, two closely related research streams that explicitly focus on actionable decision support. Counterfactual explanations aim to identify the smallest changes required in an input instance to obtain a different prediction. Rather than describing why a prediction was generated, they answer the question: *"What should be different for a preferred outcome to occur?"* (Guidotti, 2024; Verma et al., 2024). In high-stakes domains such as finance, healthcare, and credit assessment, counterfactual explanations are particularly attractive because they transform model outputs into actionable recommendations that can be interpreted by end users.

Closely related to this perspective, the literature on algorithmic recourse investigates how individuals or organizations can modify their characteristics in order to obtain a desired decision from an algorithmic system. While counterfactual explanations focus on identifying alternative scenarios, algorithmic recourse emphasizes the feasibility and actionability of such interventions, seeking recommendations that can realistically be implemented by decision makers (Ustun et al., 2019; Karimi et al., 2022). Both approaches move beyond prediction and explanation towards decision-oriented artificial intelligence, thereby providing a conceptual foundation for prescriptive analytics. HERA-PRO follows this direction by identifying the minimum financial adjustments required to achieve a target rating outcome within a hierarchical rating forecasting framework.

This perspective is particularly relevant in financial applications, where managers and analysts are often less interested in understanding a prediction than in identifying feasible strategies capable of improving future performance. Recent studies have therefore started to combine machine learning, causal reasoning, and prescriptive frameworks to support policy design and decision-making in credit-risk management and related financial domains (Huh, 2025).

[1] All empirical analyses were developed and executed using MATLAB R2024a and R2024b

2.2. Research Gap and Contribution

Despite the rapid growth of research on explainable artificial intelligence and prescriptive analytics, their application to corporate rating forecasting remains limited. Existing studies largely focus either on improving predictive accuracy or on explaining model outputs through feature importance measures, partial dependence analysis, and SHAP-based interpretations. Comparatively less attention has been devoted to transforming rating forecasting models into instruments capable of generating actionable recommendations for firms seeking to improve their future rating positions (Verma et al., 2024).

This gap is particularly evident within hierarchical rating forecasting systems. The HERA framework introduced by Falavigna (2026a) contributes to the prediction of future corporate ratings through a sector-specific hierarchical architecture, while HERA-XAI (Falavigna, 2026b) extends the framework by providing global and local explanations of individual rating forecasts. However, neither framework directly addresses the question of how firms can modify their financial characteristics to achieve a desired rating outcome.

To fill this gap, the present study introduces HERA-PRO, a prescriptive extension of the HERA family. Building upon the explanations generated by HERA-XAI, the proposed framework employs explainable machine learning and counterfactual reasoning to identify potential financial adjustments capable of increasing the probability of achieving a target rating. In doing so, HERA-PRO represents a transition from predictive and explanatory analytics toward prescriptive decision support, transforming rating forecasting models into tools that can assist managers in evaluating strategic financial actions and their potential consequences.

3. THE HERA-PRO FRAMEWORK

3.1. From Prediction to Explanation and Prescription

The HERA research programme has evolved through three successive stages, each corresponding to a different level of analytical maturity. The first stage, represented by the original HERA framework (Falavigna, 2026a), focused on the prediction of one-year-ahead corporate ratings through a sector-specific hierarchical architecture composed of interconnected Random Forest classifiers. The objective of HERA was to improve forecasting performance by decomposing the original multiclass rating problem into a network of specialized binary classification tasks. In this first stage, the framework primarily addressed the question:

What is the most likely future rating of a firm?

The second stage, HERA-XAI (Falavigna, 2026b), introduced an explainability layer designed to open the black box of the forecasting system. Through the combined use of Permutation Importance, Partial Dependence Plots, and SHAP values, HERA-XAI provided both global and local explanations of rating forecasts. The framework therefore moved beyond prediction and addressed a second question:

Why did the model generate a specific rating forecast? Is it possible to identify the financial indicators that most strongly influence future rating outcomes?

Within the HERA-XAI framework, both PDPs and SHAP values proved effective in enhancing the interpretability of the forecasting architecture, although they served different purposes. Partial Dependence Plots contributed to understanding how variations in individual financial indicators affect predicted rating probabilities, thereby revealing the direction and non-linear nature of the relationships learned by the model. However, such information is primarily

descriptive and does not necessarily indicate which actions should be undertaken by managers to improve the rating profile of a specific firm.

Similarly, both global and local SHAP analyses identified the financial indicators exerting the greatest influence on rating assignments and rating-transition probabilities. These techniques made it possible to quantify the contribution of each variable to the model's predictions and substantially improved the transparency of the forecasting process. Nevertheless, SHAP values are designed to explain a prediction rather than prescribe an intervention. In other words, they can identify why a firm is associated with a particular rating outcome, but they do not directly indicate how that outcome can be modified.

Consequently, an important gap remains between explainability and decision support. Managers are typically less interested in understanding which variables influence a rating forecast than in identifying the specific actions capable of improving the firm's future creditworthiness. This limitation motivates the development of HERA-PRO, which extends the HERA family from predictive and explanatory analytics toward prescriptive analytics. By exploiting counterfactual reasoning, HERA-PRO seeks to determine which modifications in the firm's financial profile are necessary to achieve a desired rating outcome, thereby transforming model explanations into actionable managerial recommendations. Therefore, the central question is no longer what will happen or why it will happen, but rather:

What should be changed to achieve a desired rating outcome?

Conceptually, the HERA family can be interpreted as a layered architecture composed of three complementary components:

1. HERA: Hierarchical Classification Architecture;
2. HERA-XAI: Explainability Layer;
3. HERA-PRO: Prescriptive Optimization Layer.

The first layer generates firm-level rating forecasts. The second layer identifies and explains the financial determinants underlying those forecasts. The third layer leverages this explanatory information to generate actionable recommendations aimed at improving a firm's future creditworthiness. Together, the three layers define a unified framework that progressively evolves from predictive analytics to explainable analytics and, ultimately, to prescriptive analytics.

3.2. The Prescriptive Rating Optimization Layer

The Prescriptive Rating Optimization Layer represents the final stage in the evolution of the HERA framework. While HERA focuses on forecasting future rating outcomes and HERA-XAI explains the determinants underlying such predictions, HERA-PRO extends the analysis towards prescriptive decision support by identifying the financial adjustments required to achieve a desired rating outcome.

The methodological rationale of HERA-PRO is based on a counterfactual perspective. Rather than investigating why a particular rating was assigned or evaluating the effects of predefined changes in individual financial indicators, the framework reverses the analytical process. The desired outcome is specified ex ante, and the model searches for the minimum financial modifications required to obtain that outcome. Therefore, the central question is no longer how a change affects the prediction, but rather which changes are necessary to alter the prediction itself.

In practical terms, the procedure starts from an observed firm classified into a given rating category and identifies a benchmark firm belonging to a more favourable rating class. The financial distance between the two observations defines a feasible counterfactual direction connecting the current financial profile to a target profile associated with a higher rating level.

The framework then explores this counterfactual space by generating a sequence of intermediate financial configurations located between the original observed firm and the selected benchmark. Each configuration is evaluated through the HERA hierarchical classification architecture and compared with the corresponding optimal Youden threshold. The search

continues until the smallest financial adjustment capable of crossing the classification boundary is identified.

The resulting counterfactual solution represents the minimum set of modifications required to transform the original rating prediction into the desired target rating. Consequently, HERA-PRO shifts the focus from prediction and explanation to prescription, providing managers, analysts, and financial institutions with actionable recommendations that are directly linked to the underlying classification model.

Unlike traditional scenario analyses, where the researcher specifies ex ante the magnitude and direction of the shocks applied to financial indicators, the counterfactual approach adopted in HERA-PRO specifies the desired outcome and allows the algorithm to identify the corresponding financial adjustments. This inversion of the analytical perspective constitutes the key methodological innovation of the prescriptive layer.

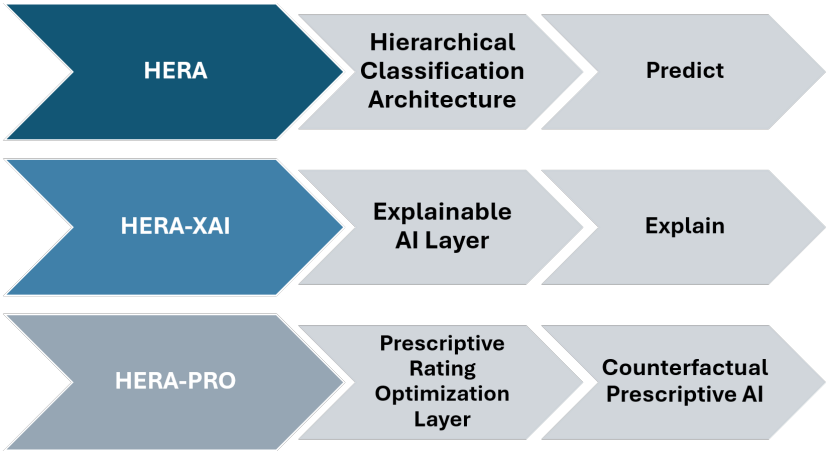
The overall framework can therefore be summarized as follows:

1. HERA (Prediction Layer): Rating Forecast;
2. HERA-XAI (Explainability AI Layer): Understanding the Determinants of the Forecast;
3. HERA-PRO (Prescriptive Rating Optimization Layer): Counterfactual Reasoning and Optimal Financial Adjustments.

Unlike traditional explainable machine-learning approaches, which are primarily descriptive, HERA-PRO translates predictive and explanatory information into operational recommendations. In doing so, the framework bridges the gap between rating forecasting and strategic decision-making, transforming artificial intelligence from a predictive instrument into a prescriptive decision-support system.

Figure 1 illustrates the progressive development of the HERA architecture. HERA (Falavigna, 2026a) focuses on one-year-ahead corporate rating forecasting through a hierarchical network of Random Forest classifiers (Predictive AI). HERA-XAI (Falavigna, 2026b) extends the framework by incorporating explainability methods to identify the determinants of rating assignments (Explainable AI). HERA-PRO builds upon these foundations by introducing a prescriptive layer that translates model explanations into actionable recommendations aimed at improving future rating outcomes (Prescriptive AI).

Figure 1. Evolution of the HERA Framework: From Predictive to Prescriptive Artificial Intelligence.



4. Methodology and Results

The objective of this section is not to generate additional explanations, but to transform the explanatory information provided by HERA-XAI into a prescriptive decision-support framework. More specifically, the analysis adopts a counterfactual perspective, aiming to identify how firms'

financial characteristics should change in order to achieve a desired rating outcome. Rather than asking why a particular rating forecast was generated, the proposed approach addresses the complementary question of what should be different for an alternative and more favourable rating scenario to emerge.

Consistently with HERA (Falavigna, 2026a) and HERA-XAI (Falavigna, 2026b), the empirical analysis is conducted using the same corporate rating dataset, the same hierarchical classification architecture, and the same set of financial indicators. The novelty of HERA-PRO does not therefore reside in the forecasting model itself, but in the way the information extracted from the Explainability AI Layer is exploited to construct counterfactual scenarios and derive firm-specific prescriptive recommendations.

Notice that HERA was originally designed to support rating prediction also for small firms characterized by limited financial disclosure. Consequently, the framework relies on a parsimonious set of eight accounting-based indicators, allowing broad applicability even when detailed financial information is unavailable. This parsimonious structure is particularly advantageous in a prescriptive context, since the search for feasible counterfactual solutions is restricted to a limited number of financially interpretable and managerially actionable variables.

The methodological procedure therefore follows a three-stage logic. First, the Hierarchical Classification Architecture generates the baseline rating forecast. Second, the Explainability Layer identifies the financial indicators exerting the greatest influence on the prediction. Third, the Prescriptive Optimization Layer explores alternative counterfactual configurations of these indicators in order to identify the minimum set of financially plausible adjustments capable of increasing the probability of achieving a target rating level. Unlike the illustrative policy simulations presented in HERA-XAI, which evaluate the effects of predefined changes in financial indicators, HERA-PRO adopts a counterfactual perspective and searches for the minimum adjustments required to achieve a desired rating outcome.

4.1. Selection of the Reference Firms

The first step of the prescriptive optimization procedure consisted in identifying an appropriate pair of reference firms for the counterfactual analysis. The objective was not to select two arbitrary firms, but rather to identify a pair of observations located as close as possible to the decision boundary separating the two rating classes considered within the HERA framework.

To illustrate the proposed methodology, the analysis focuses on firms belonging to the Production sector[2]. However, the same procedure can be applied without modification to firms operating in the Trade or Services sectors, since HERA adopts sector-specific hierarchical classification architectures (Falavigna, 2026a; Falavigna, 2026b).

The empirical application refers to Node 6 of the HERA architecture, which is responsible for distinguishing firms belonging to rating classes CCC and B. As described in Falavigna (2026a), classification decisions are obtained by comparing the probability score generated by the Random Forest classifier with the optimal threshold identified through the Youden Index optimization procedure. Let s_i denote the score assigned to firm i and let τ_Y represent the corresponding Youden threshold. The classification rule can be expressed as:

$$\text{Rating}_i = \begin{cases} \text{CCC} & \text{if } s_i < \tau_Y \\ \text{B} & \text{if } s_i \geq \tau_Y \end{cases}$$

In order to identify the reference firm belonging to the CCC class, all Production-sector firms classified as CCC at Node 6 were considered. Among them, the selected firm was the one

[2] In line with HERA (Falavigna, 2026a) and HERA-XAI (Falavigna, 2026b), the sectoral classification adopted in HERA-PRO is based on a three-sector aggregation of NACE Rev. 2 two-digit activities. The Production sector includes NACE codes 01-43; the Trade sector includes NACE codes 45-53; and the Services sector includes NACE codes 55-82. This classification is maintained throughout the HERA Family to ensure methodological consistency and comparability of results across the forecasting, explainability, and prescriptive modules.

exhibiting the score closest to the Youden threshold from below. In other words, the procedure identified the CCC firm located nearest to the decision boundary separating the two classes.

Subsequently, a corresponding peer firm belonging to class B was identified. Following a symmetric criterion, the selected peer was the Production-sector firm classified as B with the smallest score exceeding the same Youden threshold. Consequently, the two firms were located on opposite sides of the classification boundary while remaining as close as possible in terms of predicted rating score.

This selection strategy was adopted because it minimizes the distance between the original observation and the desired target class, thereby providing a particularly suitable setting for counterfactual analysis. Rather than comparing firms characterized by very different financial profiles, the procedure focuses on observations that are already close to the upgrade boundary, allowing the prescriptive framework to identify the minimum set of financial adjustments potentially required to achieve a transition from CCC to B.

At this stage, no SHAP-based analysis is performed. Although HERA-XAI (Falavigna, 2026b) demonstrated the usefulness of local explainability techniques for identifying the determinants of rating predictions, the same study also showed that explanatory relevance does not necessarily coincide with prescriptive effectiveness. A financial indicator may contribute substantially to the prediction while having only a limited impact on the probability of crossing a rating boundary. For this reason, HERA-PRO adopts a different perspective and directly investigates which financial modifications are actually capable of generating a rating improvement. The financial characteristics of the two selected firms are presented in the following section.

The calibration parameters used throughout this analysis, including the Youden threshold associated with Node 6, were obtained from the estimation procedures described in HERA (Falavigna, 2026a) and HERA-XAI (Falavigna, 2026b). Therefore, HERA-PRO does not require additional model training, but rather exploits the predictive and explanatory layers previously developed. The search procedure identified the CCC observation closest to the decision boundary, characterized by a score of 0.60, and the nearest B peer, characterized by a score of 0.62, exactly equal to the optimal Youden threshold. Since the two firms are separated exclusively by the classification boundary, they constitute an ideal benchmark pair for investigating the minimum financial adjustments required to achieve a rating upgrade.

4.2. Counterfactual Benchmark Construction

The selected CCC firm and its corresponding B peer were subsequently compared in order to identify the financial characteristics that differentiate the two observations. Since the two firms are located on opposite sides of the same decision boundary and belong to the same economic sector, their comparison provides a suitable benchmark for the construction of a counterfactual trajectory from the lower to the higher rating class.

Table 1 reports the values of the eight financial indicators employed by the HERA framework for the two selected firms. In addition to the original values, the table reports the absolute difference (AbsDiff), the percentage difference (PercDiff), the direction of the variation, and the corresponding managerial action required to align the CCC profile with the B benchmark.

The absolute difference was computed as:

$$AbsDiff_i = X_i^B - X_i^{CCC}$$

where X_i^B and X_i^{CCC} denote the values of indicator $i$for the B firm and the CCC firm, respectively.

Similarly, the percentage difference was calculated as:

$$PercDiff_i = 100 \times \frac{X_i^B - X_i^{CCC}}{\mid X_i^{CCC} \mid}$$

which measures the relative adjustment required to move from the initial CCC profile toward the benchmark B profile.

The variables *Direction* and *Action* provide an intuitive interpretation of the required modification. In particular, a positive difference suggests that the indicator should increase, whereas a negative difference indicates that a reduction would be required. Therefore, the comparison between the two firms generates a first set of candidate financial actions potentially associated with a rating improvement. In Appendix A table A.1 shows a summary of financial indicators.

Table 1. Counterfactual Benchmark Comparison between the Selected CCC Firm and the Reference B Peer

Indicator	CCC	B	AbsDiff	PercDiff	Direction	Action
DEBTA	0.9489	0.6919	-0.2570	-27.0854	-	Decrease
DEBPROD	0.9547	0.5301	-0.4247	-44.4813	-	Decrease
ROS	0.0288	-0.0437	-0.0725	-251.7564	-	Decrease
ROA	0.0286	-0.0571	-0.0857	-299.3071	-	Decrease
DEBCOV	0.0302	-0.0825	-0.1126	-373.3432	-	Decrease
FAPROD	0.0103	0.1246	0.1143	1111.0164	+	Increase
FCCOV	0.6429	-0.0125	-0.6554	-101.9444	-	Decrease
FAEQ	0.2000	0.2705	0.0705	35.2313	+	Increase

However, the presence of a difference between two observations does not necessarily imply that the corresponding indicator is able to produce a rating upgrade when modified in isolation. The benchmark comparison only identifies how the two firms differ; it does not quantify the actual effectiveness of each variable in altering the classification outcome. For this reason, an additional analysis was carried out to evaluate the marginal contribution of each indicator to the transition from class CCC to class B.

4.2.1. Individual Indicator Effect Analysis

Although the benchmark comparison identifies the financial adjustments required to replicate the profile of the reference B firm, it does not reveal whether all indicators contribute equally to the rating upgrade process. In fact, HERA-XAI (Falavigna, 2026b) showed that explanatory relevance and prescriptive effectiveness do not necessarily coincide. Financial indicators that strongly influence a prediction may not be the most effective levers for changing a rating outcome.

To address this issue, each indicator was evaluated individually. Starting from the financial profile of the selected CCC firm, one indicator at a time was replaced with the corresponding value observed for the benchmark B firm, while all remaining variables were kept unchanged. The modified profile was then submitted to the Node 6 Random Forest classifier and the resulting probability of belonging to class B was recorded.

Let X^{CCC} denote the original financial profile of the CCC firm and X^{B} denote the profile of the benchmark B firm. For each indicator j, a modified observation was constructed as:

$$X^{(j)} = \left(X_1^{CCC}, \dots, X_j^{B}, \dots, X_8^{CCC}\right)$$

where only the j-th indicator was replaced by the corresponding value observed in the B firm.

The resulting probability variation was computed as:

$$Gain_j = P\left(B \mid X^{(j)}\right) - P(B \mid X^{CCC})$$

where $P(B \mid X^{CCC})$ represents the baseline probability associated with the original CCC firm, whole $P\left(B \mid X^{(j)}\right)$ represents the new probability of CCC firm with the replaced indicator.

A positive value of $Gain_j$ indicates that the modification moves the firm closer to the B class, whereas a negative value implies that the adjustment reduces the likelihood of an upgrade. The analysis therefore provides a ranking of indicators according to their prescriptive effectiveness, allowing the identification of the most influential financial levers for the subsequent counterfactual optimization procedure.

The analysis therefore measures the marginal contribution of each indicator to the rating transition process and allows the identification of those financial variables that are truly capable of moving the firm across the classification boundary.

Table 2 summarizes the results of this procedure by comparing the baseline probability associated with the original CCC firm with the probability obtained after each individual intervention. The table also reports whether the resulting score exceeds the Youden threshold and therefore leads to an effective upgrade from class CCC to class B.

Table 2. Marginal Prescriptive Effect of Individual Financial Indicators on the Probability of Transition from CCC to B

Indicator	ProbBaseline	YudenTh	ProbNew	Gain	ObservedChange	Upgrade	Efficiency
FCCOV	0.6	0.62	0.92	0.32	0.655	True	0.488
FAPROD	0.6	0.62	0.78	0.18	0.114	True	1.575
FAEQ	0.6	0.62	0.72	0.12	0.070	True	1.703
DEBPROD	0.6	0.62	0.7	0.1	0.425	True	0.235
DEBTA	0.6	0.62	0.64	0.04	0.257	True	0.156
ROS	0.6	0.62	0.58	-0.02	0.073	False	-0.276
DEBCOV	0.6	0.62	0.52	-0.1	0.113	False	-0.710
ROA	0.6	0.62	0.42	-0.2	0.086	False	-2.101

In Table 2 columns represent:

- **Indicator:** Financial indicator considered in the simulation.
- **ProbBaseline:** Baseline probability of belonging to class B associated with the original CCC firm before any modification is applied:

$$ProbBaseline = P(B \mid X^{CCC})$$

- **YoudenTh:** Optimal classification threshold identified through the Youden Index optimization procedure for Node 6 of the HERA framework. A firm is classified as B whenever:

$$P(B) \geq \tau_Y$$

where τ_Y denotes the Youden threshold.

- **ProbNew:** Probability of belonging to class B after replacing the value of the selected indicator with the corresponding value observed for the benchmark B firm:

$$ProbNew_j = P\left(B \mid X^{(j)}\right)$$

- **Gain:** Marginal improvement generated by the modification of indicator j:

$$Gain_j = ProbNew_j - ProbBaseline$$

Positive values indicate that the intervention increases the probability of obtaining rating B, whereas negative values indicate a deterioration of the rating profile.

- **ObservedChange:** Size of the variation of the indicator observed when moving from the CCC firm to the benchmark B firm. This value is obtained as the difference in absolute value between the value of index for CCC firm and the value of the same index for B firm.

- **Upgrade:** Binary outcome indicating whether the modified observation crossed the Node 6 decision boundary:

$$Upgrade = \begin{cases} 1 & \text{if } ProbNew \geq \tau_Y \\ 0 & \text{otherwise} \end{cases}$$

An upgrade therefore occurs when the modified firm is classified as B.

- **Efficiency:** Measure of the effectiveness of the intervention in generating an increase in the probability of belonging to class B. The variable identifies which indicators provide the greatest rating improvement relative to the observed adjustment and therefore represent the most effective prescriptive levers within the HERA-PRO framework. This variable is calculated es the ratio between Gain and ObservedChange. This information is particularly relevant because represents the shift from the explanatory importance (HERA-XAI) to the prescriptive effectiveness (HERA-PRO).

Although the previous analysis provides valuable insights into the marginal effectiveness of individual financial indicators, it remains primarily a stylized exercise. In particular, the procedure assumes that a single indicator can be directly replaced with the corresponding value observed for the benchmark B firm while all remaining variables remain unchanged. From a managerial perspective, however, such an assumption may not always be realistic. Financial ratios are not decision variables per se, but rather synthetic indicators derived from one or more accounting items. Consequently, actions undertaken to modify a specific ratio are likely to affect several balance-sheet components simultaneously, generating indirect effects on other indicators included in the model. Furthermore, the magnitude of the observed adjustment may not always be economically feasible within a reasonable planning horizon.

For these reasons, the individual-indicator analysis should be interpreted as a preliminary assessment of prescriptive effectiveness rather than as an operational recommendation. A more realistic approach requires the identification of counterfactual configurations in which multiple financial indicators evolve simultaneously, allowing the framework to determine the minimum set of financially plausible adjustments necessary to achieve a rating upgrade.

Therefore, the analysis is extended beyond individual indicator substitutions to a counterfactual optimization framework. Rather than asking whether a single variable can move the firm across the classification boundary, the next step seeks to identify the minimum combination of financially feasible changes capable of transforming the original CCC profile into a profile classified as B by the HERA framework.

4.3. Minimum Counterfactual Adjustment: Single-Variable (OVAT) Analysis

The analysis presented in Section 4.2 evaluated how the rating prediction of a CCC-rated firm would change if each financial indicator were individually replaced by the corresponding value observed for the benchmark B-rated firm. Although this comparison provides useful information

on the direction of the changes associated with the rating transition, it does not necessarily imply that such large adjustments are required to obtain an upgrade. In practice, a firm may achieve a more favorable rating outcome through substantially smaller modifications of its financial profile.

From a managerial perspective, identifying the minimum adjustment required to alter the forecasted rating is particularly relevant. While large financial restructurings may be costly, time-consuming, or even infeasible, smaller interventions on selected indicators may already be sufficient to move a firm across the classification boundary.

Before introducing the multidimensional optimization procedure, a preliminary one-variable-at-a-time (OVAT) analysis is performed. In this exercise, each financial indicator is modified individually while all remaining indicators are held constant. The objective is not to identify the optimal counterfactual profile, but rather to assess the sensitivity of rating assignments to changes in individual financial dimensions and to quantify the effort required by each indicator, taken in isolation, to achieve a rating upgrade.

Although informative, this approach has an important limitation. Firms typically improve their financial conditions through simultaneous changes in multiple dimensions rather than through the modification of a single indicator. Consequently, the most relevant managerial question is not which individual variable is capable of generating an upgrade on its own, but rather what combination of financial adjustments minimizes the overall effort required to achieve the target rating. For this reason, the OVAT analysis should be interpreted as a sensitivity analysis, whereas the multidimensional optimization procedure presented in the subsequent sections represents the core contribution of HERA-PRO.

Consequently, the objective of this first prescriptive analysis is to determine the minimum variation required for each financial indicator, taken individually, to obtain the desired upgrade from the lower rating class (CCC) to the higher rating class (B).

More specifically, the analysis addresses the following questions:

What is the minimum adjustment required for each individual financial indicator to achieve the rating upgrade? Is it possible to obtain the upgrade by acting exclusively on a single variable, such as FCCOV, while leaving all the remaining indicators unchanged?

To answer these questions, a single-variable counterfactual search procedure is implemented. For each financial indicator, all remaining variables are maintained at their observed values, while the selected indicator is progressively modified within the interval defined by the value observed for the reference CCC-rated firm and the corresponding value observed for the benchmark B-rated firm.

The search procedure was based on the generation of 100 equally spaced values between the observed value of the selected financial indicator for the CCC firm and the corresponding value observed for the benchmark B firm. The interval was therefore discretized into 99 equal segments, producing a sequence of candidate counterfactual values according to:

$$x_j^{(k)} = x_j^{CCC} + \frac{k-1}{99}\left(x_j^B - x_j^{CCC}\right), \qquad k = 1, \dots, 100$$

where x_j^{CCC} and x_j^B denote the values of indicator j for the CCC firm and the benchmark B firm, respectively.

Since 100 points generate 99 intervals, the step size is not predetermined but depends on the distance between the initial and target values of the financial indicator under consideration. Consequently, larger differences between the two firms produce wider increments, whereas smaller differences generate finer adjustments.

Each candidate value generated by the procedure is sequentially evaluated through the HERA forecasting architecture, while all remaining variables are kept constant. The search stops when the predicted probability associated with the target class exceeds the classification threshold determined by the Youden optimization criterion. The corresponding value represents the minimum univariate adjustment required to obtain the desired upgrade.

Although only one variable is modified at a time, the procedure already exhibits a counterfactual nature. Rather than evaluating the consequences of arbitrarily chosen scenarios, the analysis seeks to identify the smallest financially plausible modification capable of producing an alternative outcome. In this sense, the methodology moves beyond prediction and explanation and begins to answer a more operational question:

What should be different for the firm to obtain a better rating?

Table 3 reports the results of the single-variable counterfactual search. For each financial indicator, the table compares the baseline probability associated with the original CCC-rated firm with the probability obtained after applying the minimum adjustment capable of generating an upgrade. In addition, the table reports the observed variation required for each indicator, the corresponding gain in classification probability, and a measure of efficiency that relates the rating improvement to the magnitude of the required adjustment.

Table 3. Minimum Univariate Counterfactual Adjustments Required for the Upgrade from CCC to B

Indicator	Start Value	Target Value	Required Value	Required Step	Prob Reached	Success	Min Change
ROS	0.029	-0.044	0.028	2	0.66	True	0.001
ROA	0.029	-0.057	0.028	2	0.64	True	0.001
COVDEB	0.030	-0.082	0.029	2	0.62	True	0.001
FAPROD	0.010	0.125	0.011	2	0.68	True	0.001
DEBTA	0.949	0.692	0.944	3	0.64	True	0.005
FCCOV	0.643	-0.013	0.636	2	0.66	True	0.007
DEBPROD	0.955	0.530	0.946	3	0.62	True	0.009
FAEQ	0.200	0.270	0.211	17	0.62	True	0.011

The results indicate that only limited adjustments to individual financial indicators are required to obtain a more favorable rating classification in the subsequent year. For example, a decrease of only 0.1% in the ROS ratio is sufficient to raise the probability of belonging to class B to 0.66, a value exceeding the corresponding Youden threshold and therefore leading to an upgrade from class CCC to class B. Conversely, the largest adjustment is required for the FAEQ indicator, which needs to increase by 1.1% in order to achieve the same objective.

These changes are relatively small because, as previously discussed, the selected CCC and B firms were deliberately chosen as the observations located closest to the Node 6 decision boundary. Consequently, even modest variations in the financial indicators are sufficient to move the firm across the classification threshold.

However, this evidence should not be generalized to all firms belonging to class CCC.

At first sight, some counterfactual recommendations may appear economically counterintuitive, particularly when they involve decreases in profitability indicators such as ROS or ROA. Nevertheless, HERA-PRO generates firm-specific rather than universal prescriptions. The objective of the optimization is not to maximize each individual financial indicator, but to identify the minimum multidimensional adjustment required to achieve the target rating. As a result, certain indicators may move in directions that would not normally be considered desirable when viewed in isolation, while remaining part of a broader financial configuration associated with an improved rating outcome. In addition, if the same procedure were applied to a company located further away from the Youden threshold, substantially larger adjustments might be

required, potentially exceeding what managers could realistically achieve within a reasonable planning horizon.

For this reason, the analysis is extended in the next step to a multi-variable counterfactual framework. Rather than modifying one indicator at a time, the proposed procedure simultaneously considers changes in multiple financial dimensions, allowing the identification of more realistic and economically feasible pathways towards a rating upgrade.

4.4. Multi-Variable Counterfactual Search

While the single-variable search described in the previous section evaluates the impact of modifying one financial indicator at a time, firms rarely improve their financial condition through isolated changes. Rating transitions are typically the result of simultaneous adjustments involving several dimensions of corporate performance, including profitability, leverage, debt sustainability, and asset structure. Consequently, a multivariate counterfactual procedure is introduced to identify the minimum joint modification of the firm's financial profile required to obtain the desired rating upgrade.

The procedure starts from two firms belonging to adjacent rating classes. The first firm ($Firm_0$) represents the reference CCC-rated company selected for the analysis, while the second ($Firm_1$) corresponds to the benchmark firm belonging to class B. Rather than modifying the indicators independently, the method progressively moves the financial profile of $Firm_0$ toward the profile observed for $Firm_1$.

This movement is governed by a parameter α, defined on the interval $[0, 1]$, according to:

$$x(\alpha) = x_0 + \alpha(x_1 - x_0)$$

where:

- x_0 denotes the vector of financial indicators observed for the CCC-rated firm;
- x_1 denotes the vector of financial indicators observed for the benchmark B-rated firm;
- α represents the proportion of the distance travelled between the two financial profiles.

When $\alpha = 0$, the firm's financial characteristics coincide with the original CCC-rated profile. When $\alpha = 1$, all indicators assume the values observed for the benchmark B-rated firm. Intermediate values of αgenerate alternative financial configurations located between these two extremes.

Notice that the α^* represents the minimum effort required to obtain the upgrade.

To identify the minimum counterfactual adjustment required for the upgrade, the parameter α is progressively increased over a finely discretized grid ranging from 0 to 1. For each candidate value of α, a new counterfactual financial profile is generated and submitted to the HERA forecasting architecture. The corresponding probability of belonging to class B is then evaluated against the Youden threshold adopted by the classification system.

The search stops when the first value of α capable of producing a rating upgrade is found. This value, denoted as α^*, identifies the minimum proportion of the overall adjustment path necessary to move the firm from the lower rating class to the higher one. Consequently, α^* provides a synthetic measure of the firm's distance from the target rating class: lower values indicate that only limited modifications are required, whereas higher values suggest that the firm is structurally farther from the benchmark profile.

From a methodological perspective, this procedure represents a genuine counterfactual search. Instead of evaluating arbitrarily defined scenarios, the algorithm seeks the smallest simultaneous adjustment of all financial indicators capable of generating a different classification outcome. The resulting counterfactual profile therefore answers the question:

What is the minimum combination of changes required for this firm to be classified as B rather than CCC?

The final output of the procedure consists of the counterfactual values associated with α^*, together with the absolute and relative adjustments required for each financial indicator. These values constitute the prescriptive recommendation generated by the HERA-PRO framework.

The counterfactual search procedure described in the previous section identified the minimum financial adjustment required to move the selected CCC firm across the Node 6 classification boundary and obtain a rating upgrade to class B. The resulting counterfactual solution corresponds to the first point along the path connecting the original CCC observation and the benchmark B peer that satisfies the Youden-based classification criterion.

Table 4 reports the characteristics of the minimum counterfactual solution. For each financial indicator, the table compares the original value observed for the CCC firm with the value associated with the counterfactual configuration identified by the peer weaker B. The table also reports the absolute and relative adjustments required to achieve the upgrade, together with the direction of the change. Finally, the parameters α^* and $Prob^*$ summarize the position of the counterfactual solution along the CCC-B trajectory and the corresponding probability of being classified as a B-rated firm.

The results provide a direct interpretation of the prescriptive recommendations generated by HERA-PRO. Unlike the individual-indicator analysis presented in the previous section, the counterfactual solution identifies a coordinated set of simultaneous financial adjustments that is sufficient to cross the classification boundary while minimizing the distance from the original financial profile.

Table 4. Minimum Counterfactual Solution Required to Achieve the Transition from CCC to B

Indicator	Start Value	Counterfactual Value	Absolute Change	Relative Change	Direction	Alpha Star	Prob Star
DEBTA	0.9489	0.9486	-0.0003	-0.0271	Decrease	0.001	0.76
DEBPROD	0.9547	0.9543	-0.0004	-0.0445	Decrease	0.001	0.76
ROS	0.0288	0.0287	-0.0001	-0.2518	Decrease	0.001	0.76
ROA	0.0286	0.0285	-0.0001	-0.2993	Decrease	0.001	0.76
DEBCOV	0.0302	0.0301	-0.0001	-0.3733	Decrease	0.001	0.76
FAPROD	0.0103	0.0104	0.0001	1.1110	Increase	0.001	0.76
FCCOV	0.6429	0.6422	-0.0007	-0.1019	Decrease	0.001	0.76
FAEQ	0.2000	0.2001	0.0001	0.0352	Increase	0.001	0.76

The parameter α represents the minimum proportional adjustment simultaneously applied to all actionable financial indicators that is sufficient to alter the predicted rating outcome. More specifically, α does not correspond to the value assigned to a particular variable, but rather to a common percentage variation affecting the entire financial profile of the firm according to the economic direction associated with each indicator. Positive contributors are increased by α, whereas indicators negatively associated with rating quality are reduced by the same proportion. The resulting value of α therefore provides a direct measure of the minimum financial effort required to induce a rating transition.

In the illustrative example presented above, the optimal value was found to be α = 0.001. Although this result indicates that only a very limited adjustment of the firm's financial profile is required to obtain an upgrade from CCC to B, the finding is largely attributable to the specific characteristics of the selected case. In particular, the firm was intentionally chosen among the observations located very close to the decision boundary of Node 6, implying that only marginal modifications were necessary to alter the classification outcome.

Indeed, the recommendations presented in this section should be interpreted as conditional on the specific firm and benchmark considered. Since HERA-PRO operates at the individual-firm level, different firms may require different counterfactual trajectories even when targeting the same rating category. Moreover, although the weakest B-rated firm above the Youden threshold represents a theoretically appealing reference point, alternative benchmark firms may generate different adjustment paths.

For this reason, to assess the robustness and general applicability of the proposed prescriptive framework, the same procedure was subsequently replicated on a sample of ten randomly selected CCC firms from the Production sector. This additional analysis allows the evaluation of the algorithm under more heterogeneous conditions and demonstrates its ability to identify rating-improving financial adjustments even when firms are located farther from the classification threshold.

Table 5 summarizes the minimum proportional adjustment (AlphaStar) required to achieve the target rating together with the corresponding predicted probability (ProbStar). The results confirm that the proposed algorithm remains operational even for firms located further away from the decision threshold, although substantially larger interventions may be required. Notice that Youden thresohold is equal to 0.62, then all firms are classified as B (ProbStar equal or higher than Youden threshold).

The detailed counterfactual results are presented in Appendix B, Table B.1.

Table 5. Counterfactual Optimization Outcomes Across Ten Production Firms

FirmID	AlphaStar	ProbStar
1	0.034	0.62
2	0.399	0.62
3	0.752	0.64
4	0.025	0.64
5	0.004	0.64
6	0.78	0.64
7	0.006	0.64
8	0.764	0.62
9	0.03	0.62
10	0.02	0.64

As a final exercise, the proposed procedure was further tested under a more demanding target configuration. Up to this point, the counterfactual benchmark adopted for the best-performing CCC firm corresponded to the weakest firm belonging to rating class B, namely the observation located immediately above the Youden threshold. However, a manager may not simply be interested in crossing the classification boundary and entering the higher rating category. In many cases, the objective may be to strengthen the firm's financial position well beyond the minimum admissible level required for the upgrade.

Under this perspective, the target is no longer represented by the mere crossing of the Youden threshold. Instead, the firm may aim at reaching a stronger position within rating class B. Since the probability associated with class B ranges from the Youden threshold (0.62) to 1, a higher target probability can be selected according to managerial objectives.

To illustrate this possibility, a randomly selected CCC firm different from the benchmark firm considered in the previous analyses was chosen. As counterfactual benchmark, the algorithm selected the first B-rated firm with a predicted probability at least 0.15 points higher than the

Youden threshold. Consequently, the selected peer firm exhibited a probability score greater than or equal to 0.77. In details, the selected peer B shows a score equal to 0.77.

The randomly selected CCC firm presented an initial probability of belonging to class B equal to 0.24. Since this value was below the Youden threshold, the firm was correctly classified as CCC. The optimization procedure was then applied to determine the proportion of the distance separating the CCC firm from the stronger B benchmark according to the established measure over the Youden threshold.

The results are reported in Table 6. Unlike the previous exercise, where only marginal changes were required because both firms were located very close to the classification boundary, the estimated value of α was found to be close to one. This result indicates that substantial modifications in the financial variables underlying the rating indicators would be required to achieve the desired target profile. Nevertheless, the resulting financial configuration produced a predicted probability of 0.78 of belonging to class B. Since this value is substantially higher than the Youden threshold, the firm would not only obtain the upgrade but would also reinforce its position within the target rating class. The results indicate that the firm must implement adjustments corresponding to 66.7% of the distance separating its current financial profile from the selected benchmark profile in order to achieve a predicted probability of 0.78 of belonging to rating class B.

This experiment further demonstrates the flexibility and scalability of the HERA-PRO framework. Depending on managerial objectives, the benchmark profile may correspond either to the weakest firm in the target class, when the objective is to identify the minimum effort required for an upgrade, or to stronger benchmark firms when the objective is to achieve a more substantial improvement in financial quality.

Table 6: Counterfactual Financial Adjustments Required to Reach a Stronger B-Rated Benchmark Profile (starting from a random CCC)

Indicator	Start Value	Counterfactual Value	Absolute Change	Relative Change	Direction	Alpha Star	Prob Star
DEBTA	0.995	0.828	-0.167	-16.822	Decrease	0.666	0.78
DEBPROD	1.036	10.503	9.466	913.531	Increase	0.666	0.78
ROS	0.016	0.338	0.323	2058.356	Increase	0.666	0.78
ROA	0.015	0.021	0.006	41.320	Increase	0.666	0.78
DEBCOV	0.015	0.027	0.012	77.791	Increase	0.666	0.78
FAPROD	0.640	0.380	-0.259	-40.567	Decrease	0.666	0.78
FCCOV	0.688	1.229	0.541	78.709	Increase	0.666	0.78
FAEQ	653	218.135	-434.865	-66.595	Decrease	0.666	0.78

We now consider the new firm, denoted as *FirmStar*, generated by the counterfactual optimization procedure described above, and evaluate its rating classification within the entire HERA framework rather than only at the local level of Node 6. The previous analysis demonstrated that the counterfactual financial profile is sufficient for the firm to cross the Node 6 decision boundary and obtain a B classification. However, the resulting improvements may be substantial enough to allow the firm to achieve an even higher rating category. To assess this possibility, *FirmStar* is submitted to the complete HERA hierarchical architecture, and its final rating assignment is computed. According to Table 6, the values of variables of *FirmStar* are as follows:

FirmStar = [0.828 10.503 0.338 0.021 0.027 0.380 1.229 218.135]

The simulation performed through the complete HERA framework (Falavigna, 2026a) confirms the assignment of the firm to rating class B.

It is important to note that the counterfactual optimization procedure was originally conducted at the local level of Node 6, which is responsible for distinguishing between classes CCC and B. Therefore, the obtained result guarantees that the firm successfully crosses the Node 6 decision boundary. However, once the counterfactual profile is evaluated within the entire hierarchical HERA architecture, the final rating assignment may differ from the local prediction because of the additional classification stages. In this respect, the global classification may confirm the B rating or even lead to a higher rating category, depending on the firm's position within the subsequent nodes of the hierarchy.

5. Translating Counterfactual Explanations into Managerial Actions

5.1. From Financial Indicators to Managerial Levers

The counterfactual optimization procedure developed in HERA-PRO identifies the financial indicators that should be modified in order to achieve a target rating outcome. However, managers do not directly control financial ratios; they act upon underlying business and financial decisions that subsequently affect those indicators. Consequently, a translation mechanism is required to convert counterfactual indicator configurations into actionable managerial recommendations.

To this end, a managerial action dictionary was developed linking each financial indicator to the firm-level decisions capable of influencing its value. The dictionary is constructed on the basis of the accounting relationships underlying the HERA indicators and associates each managerial action with the indicators affected and the expected direction of change.

The construction of the Action Dictionary started by defining a set of managerial actions potentially available to firms and linked to the financial indicators considered in the analysis. Subsequently, each action was evaluated in terms of its expected impact on the selected indicators, with particular attention to the direction of the relationship (i.e., whether the action was expected to increase or decrease a given indicator).

In detail, seven managerial levers were identified. The first lever, Increase Production, affects a broad set of indicators including DEBPROD, ROS, ROA, DEBCOV, FAPROD and FCCOV. Higher production levels generally improve profitability and debt coverage measures while reducing the relative weight of indebtedness and asset intensity indicators. The second lever, Reduce Operating Costs, primarily influences profitability-related indicators such as ROS, ROA and DEBCOV, while also contributing to a reduction in financial pressure measured through FCCOV.

A third lever is represented by Reduce Total Debt, which directly lowers leverage indicators such as DEBTA and DEBPROD and simultaneously improves the DEBCOV ratio. The fourth lever, Increase Equity, operates through the strengthening of the firm's capital structure and affects the FAEQ indicator. Similarly, Reduce Financial Charges targets the FCCOV ratio by decreasing the burden of interest expenses.

The sixth managerial lever, Reduce Fixed Assets, influences both asset-intensity indicators (FAPROD and FAEQ) and profitability measures such as ROA. Finally, Increase Current Assets affects the liquidity and financial structure of the firm, contributing to lower values of DEBTA while also modifying the ROA indicator.

Table 7 summarizes the managerial action dictionary developed within HERA-PRO. For each managerial lever, the table reports the financial indicators affected by the action and the expected direction of change. This mapping represents the final translation layer through which counterfactual indicator configurations are converted into actionable managerial recommendations.

Notice that the action associated with each financial indicator was determined endogenously from the benchmark comparison. Let x_0 denote the value observed for the reference firm and $x_1$the corresponding value observed for the benchmark firm. If $x_1 > x_0$, the counterfactual search considered increasing values of the indicator; conversely, if $x_1 < x_0$, the search considered decreasing values. The direction of the adjustment therefore emerged directly from the observed difference between the two firms and was not predefined by the researcher.

Table 7. HERA-PRO Action Dictionary: Managerial Levers, Affected Indicators, and Expected Direction of Change

Managerial Action	Affected Indicators	Expected Effect
Increase Production	DEBPROD	−
	ROS	+
	ROA	+
	DEBCOV	+
	FAPROD	−
	FCCOV	−
Reduce Operating Costs	ROS	+
	ROA	+
	DEBCOV	+
	FCCOV	−
Reduce Total Debt	DEBTA	−
	DEBPROD	−
	DEBCOV	+
Increase Equity	FAEQ	−
Reduce Financial Charges	FCCOV	−
Reduce Fixed Assets	FAPROD	−
	FAEQ	−
	ROA	+
Increase Current Assets	DEBTA	−
	ROA	−

Note: "+" indicates that the managerial action is expected to increase the value of the corresponding financial indicator, whereas "–" indicates an expected decrease.

The proposed dictionary constitutes the final link between the Explainability Layer and the Prescriptive Optimization Layer. While the counterfactual analysis identifies which indicators should change in order to improve the rating forecast, the managerial action dictionary translates those statistical recommendations into concrete business actions. Consequently, HERA-PRO is able not only to determine the financial profile associated with a more favourable rating outcome but also to suggest the managerial strategies that may realistically lead to such a profile.

5.2. Multi-Variable Counterfactual Search: Results

The multi-variable counterfactual search described in the previous section (Section 4) identified the minimum financial adjustment required to move the selected firm from rating class CCC to class B. The resulting counterfactual solution is reported in Table 4 and provides, for each financial indicator, the direction and magnitude of the variation necessary to cross the Node 6 classification boundary.

More specifically, the procedure generated a counterfactual profile characterized by $\alpha^* = 0.001$ and an associated probability of belonging to class B equal to 0.76, exceeding the optimal Youden threshold. The table therefore identifies, for each indicator, whether an increase or a decrease would be required to obtain the rating upgrade.

However, financial indicators do not constitute managerial actions. Ratios such as DEBTA, DEBPROD, ROA, or FAEQ are synthetic measures derived from several balance-sheet items and cannot be directly modified by corporate decision makers. Consequently, an additional translation step is required to convert the counterfactual solution into economically meaningful managerial recommendations.

As explained in Section 5.1 to achieve this objective, a managerial action dictionary was introduced. Each action contained in the dictionary is associated with a predefined set of expected effects on one or more financial indicators. For example, a debt reduction strategy is expected to decrease leverage-related indicators, whereas an increase in production is expected to improve profitability and operating efficiency measures.

The counterfactual solution obtained in Table 4 was then compared with the expected effects associated with each managerial action. For every indicator included in the counterfactual solution, the sign of the required modification was identified:

$$Direction_i = \begin{cases} + & \text{if the indicator must increase} \\ - & \text{if the indicator must decrease} \end{cases}$$

The same indicator-specific effects were retrieved from the managerial action dictionary and compared with the counterfactual directions.

For each managerial action, a Match Score was computed as:

$$Matches_a = \sum_{i=1}^{n} I\left(Effect_{ai} = Direction_i\right)$$

where $Effect_{ai}$ denotes the expected effect of managerial action a on indicator i, $Direction_i$ represents the sign required by the counterfactual solution, and $I(\cdot)$ is an indicator function equal to one when the two signs coincide and zero otherwise.

Consequently, the variable Matches measures the degree of consistency between a managerial action and the financial modifications required by the counterfactual optimization procedure. Higher values indicate actions whose expected financial effects are more closely aligned with the minimum counterfactual solution identified by HERA-PRO.

To illustrate the matching procedure, consider a counterfactual solution generated by HERA-PRO for a firm initially rated CCC. Suppose that the optimization process produces the following required changes in financial indicators:

- ROS ↑
- ROA ↑
- DEBPROD ↓
- FAPROD ↓
- FCCOV ↓

Consider the managerial action “Increase Production”, it is associated with the following expected effects:

- ROS +
- ROA +
- DEBPROD -
- FAPROD -
- FCCOV -

In this case, all five indicators exhibit the same direction of change as predicted by the dictionary. Therefore, the action receives five matches and a Match Score of 100%.

Conversely, consider the action "Reduce Debt", which may be associated with:

- DEBTA -
- DEBPROD -
- COVDEB +

If the counterfactual solution only requires a decrease in DEBPROD while the remaining indicators move in different directions, the action will receive only one match out of three possible matches, resulting in a substantially lower Match Score.

This example highlights the underlying logic of the Action Dictionary. Rather than prescribing actions directly, HERA-PRO first identifies the financial changes required to achieve the target rating and subsequently evaluates which managerial actions are most consistent with those changes. Actions exhibiting a higher Match Score are interpreted as being more coherent with the counterfactual pathway identified by the optimization process.

Table 8 reports the resulting managerial action ranking. The results show that *Increase Production*, *Reduce Total Debt*, and *Increase Current Assets* achieve the highest Match Score, with two consistent modifications each. These actions therefore appear to be the most coherent managerial strategies with respect to the counterfactual pathway identified for the selected CCC firm. Conversely, actions such as *Increase Equity* or *Reduce Fixed Assets* do not exhibit any correspondence with the required counterfactual adjustments and therefore receive a Match Score equal to zero.

Importantly, the Match Score should not be interpreted as a measure of effectiveness. Rather, it represents a measure of consistency between an observed counterfactual solution and the expected effects of a managerial action. In this sense, the procedure provides an economically interpretable translation of the counterfactual optimization results, allowing the prescriptions generated by HERA-PRO to be expressed in a language directly understandable by managers and decision makers.

Table 8. Managerial Actions Consistent with the Counterfactual Upgrade from CCC to B

Action	Matches
Increase Production	2
Reduce Total Debt	2
Increase Current Assets	2
Reduce Operating Costs	1
Reduce Financial Charges	1
Increase Equity	0
Reduce Fixed Assets	0

Therefore, the evidence presented in Table 8 should not be interpreted as a general characterization of the transition from CCC to B. Rather, it represents a firm-specific counterfactual solution obtained for two observations located near the classification boundary. While this setting is particularly useful for illustrating the mechanics of the HERA-PRO framework, it may underestimate the magnitude and diversity of the financial adjustments required for firms located farther from the target class.

In addition, notice that the direction of the adjustment is case-specific and depends on the selected counterfactual benchmark. The procedure itself is general and can be applied to any pair of adjacent rating classes and any sector-specific model. Indeed, since the benchmark firm represents the weakest observation still classified as B, some of the resulting managerial recommendations may be economically counterintuitive and should be interpreted as rating-boundary adjustments rather than as general managerial best practices.

To assess the robustness and generalizability of the proposed methodology, the next section extends the analysis to a sample of randomly selected CCC-rated firms. The same counterfactual optimization procedure is replicated for each observation, allowing the identification of recurring financial levers and providing a broader assessment of the prescriptive recommendations generated by HERA-PRO.

However, the aggregated results should not be interpreted as identifying universally optimal managerial actions. Rather, they provide preliminary evidence on the financial adjustments most frequently associated with successful counterfactual upgrades within the specific sample analyzed in this study. Since the purpose of HERA-PRO is methodological rather than normative, the results should be viewed as an exploratory assessment of the framework's ability to generate firm-specific prescriptive recommendations. Further analyses on larger samples and alternative rating transitions will be required to evaluate the stability and external validity of the observed patterns.

5.3. Robustness check: Cross-Firm Evidence

The analyses presented in the previous sections were based on individual case studies and demonstrated the ability of HERA-PRO to identify counterfactual financial configurations associated with improved rating outcomes. While such examples provide useful illustrations of the prescriptive optimization process, they do not allow an assessment of whether the derived recommendations are specific to a particular firm or reflect more general patterns.

To evaluate the robustness of the proposed framework, the complete prescriptive optimization procedure was replicated on a sample of ten randomly selected firms belonging to the Production sector and classified as CCC at Node 6. For each company, HERA-PRO generated a counterfactual financial profile capable of achieving the target class B and subsequently translated the resulting indicator adjustments into managerial actions through the action dictionary introduced in Section 5.1.

To assess the consistency of the managerial recommendations generated by HERA-PRO across different firms, a frequency-based measure denoted as mean_Matches was computed for each managerial action. For a given action, the variable *Matches* counts the number of counterfactual indicator adjustments associated with that specific managerial lever for each firm. The value of mean_Matches is then obtained by averaging these counts across the ten randomly selected CCC firms:

$$\text{mean_Matches}_a = \frac{1}{N}\sum_{i=1}^{N} \text{Matches}_{ai}$$

where $a$denotes the managerial action, $i$identifies the firm, and $N = 10$ is the number of firms included in the robustness analysis.

Consequently, mean_Matches measures the average number of correspondences between a managerial action and the counterfactual indicator modifications required by the analyzed firms. Higher values indicate actions that are more frequently associated with rating-improvement counterfactuals, whereas lower values identify interventions that emerge less systematically and are therefore more firm-specific.

Table 9 reports the cumulative frequency (mean_Matches) with which each managerial action was recommended across the ten firms.

This aggregated analysis allows the identification of the managerial levers most frequently associated with rating upgrades and provides a first assessment of the generalizability of the prescriptive recommendations generated by HERA-PRO.

Table 9. Cumulative Frequency of Recommended Managerial Actions Across Ten CCC Firms

Action	mean_Matches
Increase Production	2.6
Reduce Total Debt	1.7
Increase Current Assets	1.6
Reduce Fixed Assets	1.5
Reduce Operating Costs	1
Increase Equity	0.8
Reduce Financial Charges	0.8

The results show that, among ten firms randomly selected from the CCC group, the managerial action most frequently recommended by the algorithm to achieve an upgrade to the next rating class is Increase Production, followed by Reduce Total Debt and Increase Current Assets. Conversely, Increase Equity and Reduce Financial Charges appear less frequently among the recommended interventions.

While the action dictionary remains constant across firms, the counterfactual pathways identified by HERA-PRO are inherently firm-specific. Consequently, firms belonging to the same rating class may require substantially different financial adjustments to achieve an identical rating transition. The results presented in this section should therefore be interpreted as indicative of recurring managerial patterns rather than as a unique prescription valid for all firms.

Detailed firm-level results are reported in Appendix C. Table C.1 presents the complete counterfactual indicator configurations obtained for each of the ten randomly selected CCC firms, while Table C.2 reports the corresponding managerial actions identified through the HERA-PRO action dictionary. These supplementary results provide a more granular view of the prescriptive optimization process and support the robustness analysis discussed in the main text.

For the final exercise presented in the previous section, the managerial actions most strongly associated with the counterfactual solution are reported in Table 10. The ranking is obtained by comparing the signs and directions of the counterfactual indicator adjustments reported in Table 6 with the expected effects associated with each managerial action contained in the HERA-PRO action dictionary.

The results reported in Table 10 suggest that, for the randomly selected CCC-rated firm, achieving a probability of 0.77 and consequently obtaining an upgrade to the next rating class is primarily associated with increasing production, followed by reducing operating costs and reducing fixed assets. Overall, the findings indicate that enhancing operational efficiency is the most effective pathway toward the target rating outcome. At the same time, the counterfactual solution points to the importance of improving the firm's financial and capital structure through debt reduction and stronger capitalization, both of which contribute to the improvement of the predicted rating.

It is important to note that these recommendations are specific to the firm under analysis and derive from the particular counterfactual configuration identified by HERA-PRO. Therefore, they should be interpreted as firm-level prescriptive recommendations rather than as universally applicable managerial strategies.

Table 10. Managerial Actions Associated with the Counterfactual Upgrade of the Randomly Selected CCC Firm

Action	Matches
Increase Production	4
Reduce Operating Costs	3
Reduce Fixed Assets	3
Reduce Total Debt	2
Increase Equity	1
Increase Current Assets	1
Reduce Financial Charges	0

6. Global Counterfactual Validation and Weighted Action Scoring

The counterfactual optimization procedure described in the previous sections is performed at the level of a single hierarchical classifier. Specifically, the current version of HERA-PRO identifies the minimum financial adjustment required for a firm to cross a specific decision boundary and obtain a one-notch rating improvement. For example, within RF6 the optimization process determines the minimum value of α necessary for a CCC-rated firm to be reclassified as B.

While this local approach provides an economically meaningful measure of the effort required to improve a firm's creditworthiness, it naturally raises a broader question: can the same methodology be extended to identify the minimum changes required to achieve larger rating upgrades?

To address this issue, the HERA-PRO framework can be generalized from a local optimization problem to a global optimization problem defined on the entire HERA architecture. Rather than targeting the crossing of a single decision boundary, the objective becomes the achievement of a desired final rating class.

Let $Firm_0$ denote a firm belonging to rating class CCC and let $Firm_t$ represent a benchmark firm belonging to a higher rating category, such as BB or BBB. As in the original version of HERA-PRO, a counterfactual trajectory is constructed as:

$$Firm(\alpha)=Firm_0+\alpha(Firm_t-Firm_0)$$

where $\alpha\in[0,1]$ measures the proportion of the financial distance between the original firm and the selected benchmark that has been covered.

In the local version of HERA-PRO, the optimal value α^* is identified when the probability generated by the relevant hierarchical classifier exceeds the associated threshold. In the global version, however, each counterfactual configuration $Firm(\alpha)$ is evaluated through the complete HERA architecture, producing a final rating assignment rather than a node-specific classification.

The optimization problem can therefore be expressed as:

$$\alpha*=\min\{\alpha:Rating(Firm(\alpha))\geq Rating_{Target}\}$$

where $Rating_{Target}$ denotes the desired rating category.

Consequently, the framework is no longer restricted to one-notch upgrades. A CCC-rated firm may be evaluated against alternative rating targets such as: CCC→BB or CCC→BBB, allowing the identification of progressively more ambitious improvement paths.

To illustrate the concept, consider a randomly selected CCC-rated firm and two benchmark firms respectively belonging to classes BB and BBB. For each benchmark, the optimization procedure searches for the minimum α required for the counterfactual profile to be assigned at least the corresponding target rating when processed through the entire HERA framework. The resulting values of α provide an intuitive measure of the financial effort associated with alternative rating trajectories.

From a managerial perspective, this extension significantly increases the practical relevance of HERA-PRO. Whereas the local optimization framework answers the question:

What is the minimum adjustment required to obtain the next rating class?

The global version addresses a more strategic question:

What is the minimum adjustment required to achieve a specific long-term rating objective?

As a result, Global HERA-PRO transforms the counterfactual framework from a local rating-improvement tool into a comprehensive rating optimization system capable of supporting medium-term financial planning and strategic decision-making.

In Table 11, results of simulation are presented. The first column reports the financial indicators of the CCC-rated firm, whereas the second column presents the corresponding values of the target firm (BB or BBB). The third column measures the gap between the two profiles, while the fourth column reports the estimated value of α. The fifth column shows the percentage adjustment required from the initial CCC profile, and the final column reports the counterfactual values that the firm should attain to achieve the target rating class.

Table 11. Simulated counterfactual values for Firm CCC

					BB	
Indicator	**CCC_ Value**	**Target_ Value**	**Gap**	**Alpha**	**Alpha_x_Gap**	**Counterfactual_ Value**
DEBTA	0.816	0.778	-0.038	0.470	-0.018	0.798
DEBPROD	1.722	10.371	8.648	0.470	4.065	5.787
ROS	-0.056	0.441	0.497	0.470	0.233	0.178
ROA	-0.026	0.033	0.059	0.470	0.028	0.002
COVDEB	-0.032	0.043	0.075	0.470	0.035	0.003
FAPROD	1.333	12.729	11.396	0.470	5.356	6.689
CFCCOV	-2.000	0.013	2.013	0.470	0.946	-1.054
FAEQ	2.000	4.294	2.294	0.470	1.078	3.078
					BBB	
Indicator	**CCC_ Value**	**Target_ Value**	**Gap**	**Alpha**	**Alpha_x_Gap**	**Counterfactual_ Value**
DEBTA	0.816	0.338	-0.477	0.580	-0.277	0.539
DEBPROD	1.722	0.255	-1.468	0.580	-0.851	0.871
ROS	-0.056	0.011	0.066	0.580	0.039	-0.017
ROA	-0.026	0.014	0.041	0.580	0.024	-0.003
COVDEB	-0.032	0.043	0.075	0.580	0.044	0.011
FAPROD	1.333	0.178	-1.155	0.580	-0.670	0.663
CFCCOV	-2.000	0.125	2.125	0.580	1.233	-0.768

FAEQ	2.000	0.348	-1.652	0.580	-0.958	1.042

In addition, in this global framework we propose managerial actions based on Weighted Match Score. While the Match Score measures the proportion of indicator directions that are consistent with a given managerial action, it does not account for the magnitude of the financial adjustments required by the counterfactual solution. Consequently, a minor variation and a substantial variation contribute equally to the score as long as their signs coincide.

To overcome this limitation, HERA-PRO introduces a Weighted Match Score, which incorporates the intensity of the changes identified by the counterfactual optimization procedure. Let Δ_j denote the percentage variation required for indicator j to move from the original firm profile to the counterfactual profile $Firm^*$. For a given managerial action a, only the indicators included in the Action Dictionary are considered. The Weighted Match Score is then computed as the ratio between the total absolute variation associated with indicators whose direction is consistent with the action and the total absolute variation of all indicators affected by that action:

$$WeightedMatchScore_a = 100 \times \frac{\sum_{j \in M_a} | \Delta_j |}{\sum_{j \in A_a} | \Delta_j |}$$

where:

- A_a is the set of indicators associated with action a in the Action Dictionary;
- M_a is the subset of indicators whose counterfactual variation exhibits the same sign as the effect predicted by action a;
- $| \Delta_j |$ denotes the absolute percentage variation of indicator j.

The resulting score ranges from 0 to 100. A value of 100 indicates that all the relevant changes required by the counterfactual solution are fully consistent with the managerial action under consideration, whereas a value of 0 indicates that none of the required variations are coherent with the expected effects of that action.

Table 12 reports the Match Scores and Weighted Match Scores associated with the managerial actions contained in the Action Dictionary for two alternative target ratings, namely BB and BBB. While the Match Score measures the proportion of indicator directions consistent with a given action, the Weighted Match Score additionally accounts for the magnitude of the financial adjustments required by the counterfactual solution.

For the BB target, the results suggest that reducing operating costs represents the most coherent managerial strategy. Although the action achieves a Match Score of 75%, its Weighted Match Score reaches 93.07%, indicating that the largest counterfactual adjustments are concentrated in financial indicators directly affected by cost reductions. Conversely, increasing current assets exhibits a Match Score of 50% but a Weighted Match Score of only 2.03%, suggesting that the matched indicators correspond to relatively small adjustments and therefore play a limited role in the counterfactual trajectory.

The results become even more pronounced for the BBB target. Three actions, namely Reduce Total Debt, Increase Equity, and Reduce Fixed Assets, achieve perfect scores both in terms of Match Score and Weighted Match Score. This finding suggests that the transition towards investment-grade ratings requires a substantial restructuring of firms' financial conditions, primarily through improvements in capital structure and leverage-related indicators. Moreover, while Increase Production and Reduce Operating Costs remain relevant actions, their Weighted Match Scores indicate that debt reduction and capitalization strategies become increasingly important as more ambitious rating targets are considered.

More generally, the comparison between Match Scores and Weighted Match Scores highlights the importance of considering not only the direction of the required financial changes but also their magnitude. Actions with similar Match Scores may differ substantially in terms of their effective contribution to the counterfactual solution, as evidenced by the large discrepancies observed for some managerial actions. Therefore, the Weighted Match Score provides a more

economically informative ranking of candidate actions and strengthens the prescriptive component of the HERA-PRO framework.

Table 12. Action Dictionary Results: Comparison between Match Scores and Weighted Match Scores for Alternative Rating Targets

		BB
Action	**MatchScore**	**WeightedScore**
Reduce Operating Costs	75	93.07
Increase Production	50	48.12
Reduce Total Debt	66.67	32.02
Reduce Fixed Assets	33.33	18.89
Increase Current Assets	50	2.03
Increase Equity	0	0
Reduce Financial Charges	0	0
		BBB
Action	**MatchScore**	**WeightedScore**
Reduce Total Debt	100	100
Increase Equity	100	100
Reduce Fixed Assets	100	100
Increase Production	83.33	86.47
Reduce Operating Costs	75	82.68
Increase Current Assets	50	27.4
Reduce Financial Charges	0	0

Unlike the traditional Match Score, the Weighted Match Score assigns greater importance to indicators requiring larger adjustments. As a result, managerial actions are ranked not only according to the number of matching indicators but also according to the economic relevance of the financial changes associated with the counterfactual transition. This provides a more informative and realistic assessment of the actions that are most consistent with the rating improvement path identified by HERA-PRO. The Match Score evaluates directional consistency, whereas the Weighted Match Score evaluates both directional consistency and economic relevance.

The final exercise extends the previous analysis by comparing the randomly selected CCC-rated firm with ten firms randomly drawn from the B, BB, and BBB rating classes. Figure 2 reports the distribution of the corresponding α^* values and highlights the median effort required to achieve each target rating category.

The results provide further support for the evidence emerging from the illustrative examples discussed in this section. Specifically, the required counterfactual effort increases as firms pursue more ambitious rating objectives. While relatively limited financial adjustments may be sufficient to achieve a transition from CCC to B (median value of $\alpha^*=0.2$), substantially larger changes are generally required to reach higher rating classes such as BB and BBB (i.e., respectively: median value of $\alpha^*=0.6$ for BB class and 0.7 for BBB class).

Although these findings seem to suggest that α^* can serve as a synthetic measure of the financial effort associated with alternative rating trajectories, confirming the intuitive expectation that more ambitious rating targets require increasingly substantial financial adjustments, the following simulation reveals a more nuanced picture. Using the same CCC-rated firm and multiple benchmark firms randomly selected from the target rating classes, markedly different α^* values emerge. As shown in Figure 3, the effort required to achieve a rating upgrade (i.e., B, BB, BBB, and A) is influenced not only by the target rating category but also by the benchmark firm selected and by the non-linear geometry of the HERA classification architecture. Figure 3

provides an interesting illustration of the logic underlying Global HERA-PRO. For the selected firm, the minimum effort required to move from CCC to B is associated with a median value of α* equal to 0.45. The value increases to 0.50 when the target rating becomes BB and to 0.55 when the objective is BBB. Surprisingly, reaching class A requires an α equal to 0.50, the same value observed for the transition from CCC to BB. At first sight, this result may appear counterintuitive, since higher rating targets would normally be expected to require larger counterfactual adjustments. However, the non-linear and hierarchical structure of HERA implies that the effort associated with a rating upgrade is not determined solely by the target class, but also by the geometry of the decision boundaries and by the financial distance between the original firm and the selected benchmark. Consequently, similar values of α* may correspond to substantially different counterfactual transformations. What changes is not necessarily the overall effort measure itself, but the composition of that effort, namely the financial indicators involved, and the magnitude of the underlying economic adjustments required to achieve the target rating.

Figure 2. Counterfactual effort required to achieve alternative rating targets: distribution of α^* values across randomly selected benchmark firms (from CCC to B, BB, BBB)

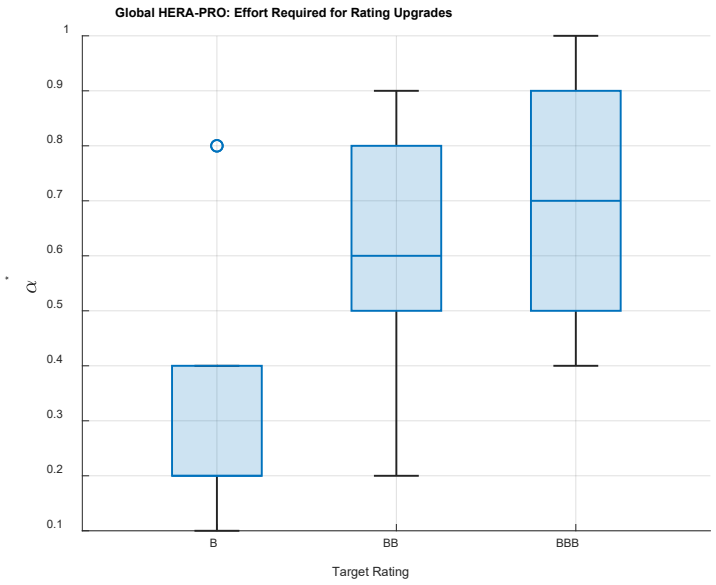
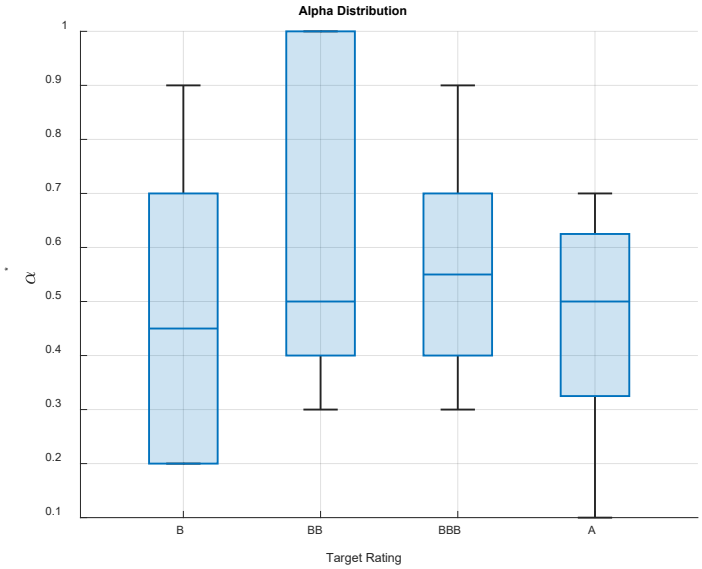


Figure 3. Counterfactual effort required to achieve alternative rating targets: distribution of α^*values across randomly selected benchmark firms (from CCC to B, BB, BBB, A)



Indeed, while α^* measures the minimum proportion of the counterfactual path required to achieve a target rating, it does not capture the economic magnitude of the financial adjustments involved. For this reason, a Weighted Effort indicator is introduced to quantify the overall financial effort associated with a rating upgrade.

For each financial indicator j, the counterfactual adjustment is first computed as:

$$\Delta_j = \alpha^*\left(X_{Target,j} - X_{0,j}\right)$$

where $X_{0,j}$ denotes the firm's original value and $X_{Target,j}$ represents the corresponding value observed in the benchmark firm belonging to the target rating class.

The adjustment is then expressed as a percentage change relative to the initial firm profile:

$$\Delta_j^{\%} = 100 \cdot \frac{\Delta_j}{\mid X_{0,j} \mid}$$

Finally, the Weighted Effort measure is calculated as the sum of the absolute percentage changes across all financial indicators:

$$WeightedEffort = \sum_{j=1}^{p} \mid \Delta_j^{\%} \mid$$

Unlike α^*, which represents a relative position along the counterfactual path, Weighted Effort captures the overall magnitude of the financial adjustments required to achieve the target rating. Higher values indicate that more substantial financial changes are needed and therefore correspond to more demanding rating improvement trajectories.

Table 13 summarizes the values of α and the corresponding weighted financial effort required to achieve different target rating classes within the Global HERA-PRO framework. The results reveal that the relationship between rating ambitions and financial effort is not strictly monotonic. While the median value of α increases from 0.45 for target B to 0.55 for target BBB, the median

α* associated with target A is equal to 0.50, matching the value observed for BB. This finding reflects the hierarchical and non-linear nature of the HERA architecture, where the effort required to reach a given rating class depends not only on the target rating itself but also on the position of the benchmark firm and the geometry of the decision boundaries.

A different picture emerges when considering the Weighted Effort indicator. Although some target classes exhibit similar values of α, the associated financial effort may vary substantially. In particular, the BB target displays the highest average weighted effort, suggesting that the corresponding counterfactual trajectories require more substantial adjustments in the underlying financial indicators. Conversely, the BBB target is associated with a larger median α but a considerably lower average weighted effort, indicating that the magnitude and composition of the required financial adjustments can differ significantly across targets.

These findings highlight the importance of complementing α with a measure of weighted financial effort. While α captures the proportion of financial distance that must be covered to reach a target rating, the weighted effort indicator provides additional information on the intensity of the underlying changes required in firms' financial profiles. Consequently, similar values of α* may correspond to substantially different managerial and financial implications.

Table 13. Global HERA-PRO Results: Distribution of α and Weighted Financial Effort Across Alternative Rating Targets

Target	mean_ Alpha	median_ Alpha	min_ Alpha	max_ Alpha	mean_ Weighted Effort	median_ Weighted Effort	min_ Weighted Effort	max_ Weighted Effort
B	0.47	0.45	0.2	0.9	1,144.42	494.152	199.056	3,782.826
BB	0.63	0.5	0.3	1	4,180.30	851.286	440.302	26,290.814
BBB	0.58	0.55	0.3	0.9	928.74	872.588	681.966	1,309.974
A	0.44	0.5	0.1	0.7	2,684.42	1,159.173	434.577	9,497.820

7. Conclusions and Further Studies

This paper introduced HERA-PRO, a prescriptive extension of the Hierarchical Artificial Intelligence Rating Architecture originally proposed in Falavigna (2026a) and subsequently enriched through the explainability layer developed in HERA-XAI (Falavigna, 2026b). While the original HERA framework focused on rating prediction and HERA-XAI concentrated on explaining the determinants of rating forecasts, HERA-PRO advances the framework further by transforming predictive and explanatory information into actionable managerial recommendations.

The main contribution of the proposed approach lies in the integration of hierarchical artificial intelligence, explainable machine learning, and counterfactual reasoning within a unified decision-support framework. Unlike traditional Explainable Artificial Intelligence methodologies, which primarily aim to clarify why a certain prediction has been generated, HERA-PRO addresses a more operational question:

what should change in order to obtain a more favorable rating outcome?

By shifting the focus from explanation to prescription, the proposed framework moves beyond model transparency and directly supports managerial decision-making.

From a methodological perspective, HERA-PRO introduces a genuine counterfactual approach. Rather than analysing the effect of predefined variations in selected indicators, as typically done in sensitivity analyses or scenario simulations, the framework starts from a desired rating outcome and searches for the minimum set of financial adjustments required to achieve it. In this sense, the counterfactual question addressed by HERA-PRO can be stated as follows:

what should be different in the financial profile of a firm for the model to assign a higher rating class?

This perspective represents a substantial departure from traditional explainability methods, moving from the interpretation of predictions to the identification of feasible paths toward alternative outcomes.

A central element of the proposed framework is the parameter α, which measures the minimum proportional adjustment simultaneously applied to the firm's financial indicators that is necessary to obtain a rating upgrade. Consequently, α can be interpreted as a measure of the minimum multidimensional financial effort required to move from the current rating class to the target next one-year-ahead rating category. The parameter therefore provides managers with a practical indication of how distant the firm is from the target class and quantifies, in a synthetic way, the magnitude of the financial changes required to improve future creditworthiness.

The empirical analysis demonstrates that local explainability information can be effectively combined with counterfactual optimization procedures to identify financially plausible modifications capable of altering rating classifications. The resulting recommendations are subsequently translated into managerial actions through a dedicated action dictionary, and ranked according to a Weighted Match Score, thereby creating a direct link between financial indicators and strategic corporate decisions. In this respect, HERA-PRO represents a bridge between predictive analytics and practical decision support, providing firms with an interpretable mechanism for evaluating potential pathways toward rating improvement.

An important distinction must nevertheless be emphasized. While the managerial action dictionary is general and derives directly from the accounting structure of the financial indicators, the counterfactual solutions generated by HERA-PRO remain inherently firm-specific. Consequently, firms belonging to the same rating category may receive substantially different recommendations and, in some cases, even recommendations pointing in opposite directions. This finding highlights the importance of avoiding excessive generalizations from aggregated results. Although the robustness analysis identified recurring patterns associated with production enhancement, debt reduction, and improvements in current assets, such evidence should not be interpreted as universally optimal strategies for all firms within a given sector. For this reason, the firm-level results reported in Appendix C are arguably as informative as the aggregated evidence presented in the main text, since they reveal the heterogeneity of the counterfactual solutions identified by the framework.

Another relevant contribution concerns the managerial interpretation of rating forecasting models. Corporate rating prediction is ultimately valuable only if it helps decision-makers understand which aspects of the firm's financial structure should be modified to improve future outcomes. HERA-PRO addresses this requirement by translating rating forecasts into operational recommendations, thereby transforming a forecasting system into a managerial planning instrument. In this respect, the framework does not merely predict whether a firm is likely to receive a better rating in the following year; it also identifies the minimum adjustments potentially required to achieve such an improvement.

The final methodological extension introduced in this study demonstrates that HERA-PRO is not limited to identifying the minimum effort required for a one-class rating upgrade. By extending the optimization process to the entire HERA architecture, firms can assess the financial adjustments necessary to achieve substantially higher rating targets and compare alternative improvement trajectories. At the same time, the introduction of a Weighted Match Score enables the Action Dictionary to rank managerial actions according to both directional consistency and financial relevance. This combination of global optimization and weighted managerial recommendations represents a significant step toward a truly prescriptive artificial intelligence framework for corporate rating management.

Despite these contributions, several limitations should be acknowledged. First, the prescriptive optimization procedure operates on a set of financial indicators derived from accounting statements rather than on the underlying accounting items themselves. Consequently, although the framework identifies the indicator adjustments associated with a target rating transition, it

does not directly determine the specific accounting operations required to achieve those adjustments. Second, the analysis considers each recommendation as technically feasible and does not explicitly incorporate implementation costs, managerial constraints, or strategic trade-offs that may affect the practical adoption of the proposed actions. Third, the present application adopts as benchmark the weakest firm belonging to the target rating class and located at the Youden threshold. Alternative implementations could instead consider stronger benchmark firms or multiple benchmark profiles, potentially generating different prescriptive trajectories.

Future research may extend HERA-PRO in several directions. A first promising development consists of moving from financial indicators to the underlying accounting variables and balance-sheet items that generate them. Although the current framework identifies the indicator adjustments associated with a target rating transition, managers do not directly control financial ratios but rather the accounting and operational variables from which those ratios are derived. Extending the optimization procedure to the balance-sheet level would therefore allow the framework to identify not only which indicators should change, but also which specific accounting, financial, and managerial interventions would be required to produce those changes.

A second research direction concerns the integration of causal inference and policy learning techniques. In its current version, HERA-PRO identifies counterfactual configurations associated with improved rating outcomes, but it does not explicitly model the causal mechanisms linking managerial actions to future rating transitions. Incorporating causal inference methods would make it possible to distinguish predictive associations from causal effects, while policy learning approaches could support the identification of optimal managerial strategies under alternative economic scenarios and operational constraints. Such an extension would represent an important step from prescriptive recommendations toward evidence-based policy design.

Furthermore, the framework is naturally scalable and could be expanded to consider alternative rating targets, multiple benchmark firms, higher rating jumps, and more complex optimization criteria incorporating implementation costs, feasibility constraints, risk preferences, and strategic priorities. Future versions of HERA-PRO could also compare alternative benchmark profiles, moving beyond the weakest firm belonging to the target rating class and considering stronger benchmark firms associated with more ambitious improvement pathways.

A final limitation concerns the computational complexity of the proposed framework. The combination of hierarchical classification, explainability techniques, and counterfactual optimization results in a computationally demanding architecture, especially when applied to large firm-level databases. Future research could therefore explore alternative learning algorithms within the HERA architecture, replacing the current Random Forest classifiers with more computationally efficient approaches such as XGBoost or other gradient boosting techniques. This would allow the fundamental logic of the HERA framework to be preserved while potentially improving scalability, estimation speed, and applicability to larger datasets.

Overall, HERA-PRO represents a further step in the evolution of the HERA research program, extending the framework from prediction (HERA) to explanation (HERA-XAI), and finally to prescription. By introducing counterfactual reasoning into hierarchical rating forecasting, the proposed approach opens new research avenues at the intersection of artificial intelligence, explainable machine learning, causal inference, policy learning, and strategic financial decision-making. In doing so, HERA-PRO contributes to transforming rating forecasting models from predictive tools into actionable decision-support systems capable of guiding firms toward improved future rating outcomes.

8. REFERENCES

Bertsimas, D., & Kallus, N. (2020). *From predictive to prescriptive analytics*. Management Science, 66(3), pp. 1025-1044.

Falavigna, G. (2026a). *HERA:Hierarchical Engine for Rating Assessment. A Sector-Specific Hierarchical Artificial Intelligence Framework for One-Year-Ahead Corporate Rating*

Forecasting (CNR-IRCrES Working Paper 4/2026). Istituto di Ricerca sulla Crescita Economica Sostenibile. http://dx.doi.org/10.23760/2421-7158.2026.004

Falavigna, G. (2026b). *HERA-XAI: Explaining One-Year-Ahead Corporate Rating Forecasts Through Hierarchical Artificial Intelligence* (CNR-IRCrES Working Paper 6/2026). Istituto di Ricerca sulla Crescita Economica Sostenibile. http://dx.doi.org/10.23760/2421-7158.2026.006

Guidotti, R. (2024). Counterfactual explanations and how to find them: literature review and benchmarking. *Data mining and knowledge discovery*, 38(5), pp. 2770-2824.

Huh, J. (2025). Prescriptive analytics for sustainable financial systems: A causal-machine learning framework for credit risk management and targeted marketing. *Systems*, 14(1), 16.

Karimi, A. H., Barthe, G., Schölkopf, B., & Valera, I. (2022). A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5), pp. 1-29.

Ustun, B., Spangher, A., & Liu, Y. (2019, January). Actionable recourse in linear classification. In *FACCT 19: Proceedings of the conference on fairness, accountability, and transparency* (pp. 10-19).

Verma, S., Boonsanong, V., Hoang, M., Hines, K., Dickerson, J., & Shah, C. (2024). Counterfactual explanations and algorithmic recourses for machine learning: A review. *ACM Computing Surveys*, 56(12), pp. 1-42.

9. APPENDIX A. FINANCIAL INDICATORS

Table A.1 Financial indicators used in the HERA Framework

Indicator	Financial Ratio	Economic Interpretation
DEBTA	Total Debt / Total Assets	Measures the overall level of indebtedness relative to total assets. Higher values indicate greater financial leverage and a stronger reliance on external financing.
DEBPROD	Total Debt / Production Value	Measures the amount of debt supported by the firm's productive activity. Higher values may indicate a heavier debt burden relative to operating capacity.
ROS	Operating Income / Production Value	Return on Sales. Measures operating profitability and the firm's ability to generate earnings from sales activities.
ROA	Operating Income / Total Assets	Return on Assets. Evaluates the profitability generated by the firm's asset base and the efficiency with which assets are employed.
DEBCOV	Operating Income / Total Debt	Debt Coverage Ratio. Measures the firm's ability to generate operating earnings relative to total debt. Higher values indicate stronger debt-servicing capacity.
FAPROD	Fixed Assets / Production Value	Fixed Assets to Production Ratio. Captures the intensity of fixed capital employed in productive activities. Higher values generally characterize more capital-intensive firms.
FCCOV	Financial Charges / Operating Income	Financial Charges Coverage Ratio. Measures the burden of financial expenses relative to operating earnings. Higher values may signal greater financial pressure.
FAEQ	Fixed Assets / Net Equity	Fixed Assets to Equity Ratio. Evaluates the extent to which fixed investments are financed by long-term equity sources. Lower values generally indicate a more balanced financial structure.

Source: Falavigna, 2026b.

10. APPENDIX B. DETAILED COUNTERFACTUAL RESULTS

Table B.1 provides the complete set of simulation results for the ten firms considered in the empirical analysis, including initial indicator values, counterfactual indicator values, AlphaStar, ProbStar, and the associated rating transition outcomes.

Table B1. Firm-Level Counterfactual Profiles and Prescriptive Recommendations.

FirmID	Indicator	Start Value	Counterfactual Value	Absolute Change	Relative Change	Direction	Alpha Star	Prob Star
1	DEBTA	0.995	0.985	-0.010	-1.037	Decrease	0.034	0.62
1	DEBPROD	1.036	1.019	-0.017	-1.661	Decrease	0.034	0.62
1	ROS	0.016	0.014	-0.002	-12.885	Decrease	0.034	0.62
1	ROA	0.015	0.013	-0.002	-16.289	Decrease	0.034	0.62
1	DEBCOV	0.015	0.012	-0.003	-21.942	Decrease	0.034	0.62
1	FAPROD	0.640	0.622	-0.018	-2.738	Decrease	0.034	0.62
1	FCCOV	0.688	0.664	-0.024	-3.462	Decrease	0.034	0.62
1	FAEQ	653.000	630.807	-22.193	-3.399	Decrease	0.034	0.62
2	DEBTA	0.993	0.873	-0.120	-12.101	Decrease	0.399	0.62
2	DEBPROD	1.981	1.402	-0.579	-29.222	Decrease	0.399	0.62
2	ROS	0.022	-0.004	-0.026	-118.392	Decrease	0.399	0.62
2	ROA	0.011	-0.016	-0.027	-244.238	Decrease	0.399	0.62
2	DEBCOV	0.011	-0.026	-0.037	-333.186	Decrease	0.399	0.62
2	FAPROD	0.019	0.061	0.042	215.759	Increase	0.399	0.62
2	FCCOV	1.375	0.821	-0.554	-40.263	Decrease	0.399	0.62
2	FAEQ	0.778	0.575	-0.202	-26.025	Decrease	0.399	0.62
3	DEBTA	1.906	0.993	-0.913	-47.905	Decrease	0.752	0.64
3	DEBPROD	12.708	3.550	-9.158	-72.063	Decrease	0.752	0.64
3	ROS	-0.813	-0.234	0.578	71.154	Increase	0.752	0.64
3	ROA	-0.122	-0.073	0.049	39.990	Increase	0.752	0.64
3	DEBCOV	-0.064	-0.078	-0.014	-21.807	Decrease	0.752	0.64
3	FAPROD	6.646	1.742	-4.904	-73.790	Decrease	0.752	0.64
3	FCCOV	-0.308	-0.086	0.222	72.145	Increase	0.752	0.64
3	FAEQ	0.444	0.313	-0.130	-29.358	Decrease	0.752	0.64
4	DEBTA	0.853	0.849	-0.004	-0.471	Decrease	0.025	0.64
4	DEBPROD	1.993	1.957	-0.037	-1.835	Decrease	0.025	0.64
4	ROS	-0.007	-0.008	-0.001	-13.675	Decrease	0.025	0.64
4	ROA	-0.003	-0.004	-0.001	-46.858	Decrease	0.025	0.64
4	DEBCOV	-0.003	-0.005	-0.002	-58.325	Decrease	0.025	0.64
4	FAPROD	0.872	0.853	-0.019	-2.143	Decrease	0.025	0.64
4	FCCOV	-2.000	-1.950	0.050	2.484	Increase	0.025	0.64
4	FAEQ	2.389	2.336	-0.053	-2.217	Decrease	0.025	0.64
5	DEBTA	0.844	0.843	-0.001	-0.072	Decrease	0.004	0.64
5	DEBPROD	0.721	0.721	-0.001	-0.106	Decrease	0.004	0.64

5	ROS	0.039	0.038	0.000	-0.852	Decrease	0.004	0.64
5	ROA	0.045	0.045	0.000	-0.904	Decrease	0.004	0.64
5	DEBCOV	0.054	0.053	-0.001	-1.015	Decrease	0.004	0.64
5	FAPROD	0.091	0.091	0.000	0.147	Increase	0.004	0.64
5	FCCOV	0.333	0.332	-0.001	-0.415	Decrease	0.004	0.64
5	FAEQ	0.603	0.602	-0.001	-0.221	Decrease	0.004	0.64
6	DEBTA	7.043	2.089	-4.954	-70.338	Decrease	0.78	0.64
6	DEBPROD	1.131	0.662	-0.468	-41.434	Decrease	0.78	0.64
6	ROS	0.002	-0.034	-0.036	-1746.736	Decrease	0.78	0.64
6	ROA	0.013	-0.042	-0.054	-427.658	Decrease	0.78	0.64
6	DEBCOV	0.002	-0.064	-0.066	-3637.588	Decrease	0.78	0.64
6	FAPROD	0.159	0.132	-0.027	-17.070	Decrease	0.78	0.64
6	FCCOV	4.444	0.968	-3.476	-78.219	Decrease	0.78	0.64
6	FAEQ	2.583	0.779	-1.804	-69.832	Decrease	0.78	0.64
7	DEBTA	0.428	0.429	0.002	0.371	Increase	0.006	0.64
7	DEBPROD	5.217	5.189	-0.028	-0.539	Decrease	0.006	0.64
7	ROS	0.121	0.120	-0.001	-0.817	Decrease	0.006	0.64
7	ROA	0.010	0.010	0.000	-4.057	Decrease	0.006	0.64
7	DEBCOV	0.023	0.023	-0.001	-2.737	Decrease	0.006	0.64
7	FAPROD	10.776	10.712	-0.064	-0.593	Decrease	0.006	0.64
7	FCCOV	0.296	0.294	-0.002	-0.625	Decrease	0.006	0.64
7	FAEQ	1.598	1.590	-0.008	-0.498	Decrease	0.006	0.64
8	DEBTA	1.285	0.832	-0.453	-35.268	Decrease	0.764	0.62
8	DEBPROD	0.893	0.616	-0.278	-31.066	Decrease	0.764	0.62
8	ROS	0.027	-0.027	-0.054	-198.120	Decrease	0.764	0.62
8	ROA	0.039	-0.034	-0.074	-186.840	Decrease	0.764	0.62
8	DEBCOV	0.031	-0.056	-0.086	-281.534	Decrease	0.764	0.62
8	FAPROD	0.113	0.122	0.009	7.982	Increase	0.764	0.62
8	FCCOV	0.111	0.017	-0.094	-84.995	Decrease	0.764	0.62
8	FAEQ	-0.569	0.072	0.642	112.700	Increase	0.764	0.62
9	DEBTA	0.591	0.594	0.003	0.515	Increase	0.03	0.62
9	DEBPROD	0.568	0.567	-0.001	-0.203	Decrease	0.03	0.62
9	ROS	0.039	0.037	-0.002	-6.325	Decrease	0.03	0.62
9	ROA	0.041	0.038	-0.003	-7.177	Decrease	0.03	0.62
9	DEBCOV	0.069	0.065	-0.005	-6.565	Decrease	0.03	0.62
9	FAPROD	0.479	0.468	-0.011	-2.220	Decrease	0.03	0.62
9	FCCOV	0.410	0.397	-0.013	-3.091	Decrease	0.03	0.62
9	FAEQ	1.296	1.265	-0.031	-2.374	Decrease	0.03	0.62
10	DEBTA	0.623	0.624	0.001	0.221	Increase	0.02	0.64
10	DEBPROD	1.411	1.394	-0.018	-1.249	Decrease	0.02	0.64
10	ROS	0.045	0.043	-0.002	-3.942	Decrease	0.02	0.64

10	ROA	0.020	0.018	-0.002	-7.742	Decrease	0.02	0.64
10	DEBCOV	0.032	0.030	-0.002	-7.170	Decrease	0.02	0.64
10	FAPROD	0.043	0.045	0.002	3.731	Increase	0.02	0.64
10	FCCOV	0.793	0.777	-0.016	-2.032	Decrease	0.02	0.64
10	FAEQ	0.051	0.055	0.004	8.625	Increase	0.02	0.64

11. Appendix C. Counterfactual Analyses and Managerial Recommendations (Section 5)

Table C.1 reports the original financial indicators, the counterfactual indicator values generated by HERA-PRO, the corresponding AlphaStar values, and the predicted probabilities associated with the target rating class for each firm included in the robustness analysis.

Table C.1. Firm-Level Counterfactual Indicator Configurations for Ten Randomly Selected CCC Firms

FirmID	Indicator	Start Value	Counterfactual Value	Absolute Change	Relative Change	Direction	Alpha Star	Prob Star
1	DEBTA	0.995	0.985	-0.010	-1.037	Decrease	0.034	0.62
1	DEBPROD	1.036	1.019	-0.017	-1.661	Decrease	0.034	0.62
1	ROS	0.016	0.014	-0.002	-12.885	Decrease	0.034	0.62
1	ROA	0.015	0.013	-0.002	-16.289	Decrease	0.034	0.62
1	DEBCOV	0.015	0.012	-0.003	-21.942	Decrease	0.034	0.62
1	FAPROD	0.640	0.622	-0.018	-2.738	Decrease	0.034	0.62
1	FCCOV	0.688	0.664	-0.024	-3.462	Decrease	0.034	0.62
1	FAEQ	653.000	630.807	-22.193	-3.399	Decrease	0.034	0.62
2	DEBTA	0.993	0.873	-0.120	-12.101	Decrease	0.399	0.62
2	DEBPROD	1.981	1.402	-0.579	-29.222	Decrease	0.399	0.62
2	ROS	0.022	-0.004	-0.026	-118.392	Decrease	0.399	0.62
2	ROA	0.011	-0.016	-0.027	-244.238	Decrease	0.399	0.62
2	DEBCOV	0.011	-0.026	-0.037	-333.186	Decrease	0.399	0.62
2	FAPROD	0.019	0.061	0.042	215.759	Increase	0.399	0.62
2	FCCOV	1.375	0.821	-0.554	-40.263	Decrease	0.399	0.62
2	FAEQ	0.778	0.575	-0.202	-26.025	Decrease	0.399	0.62
3	DEBTA	1.906	0.993	-0.913	-47.905	Decrease	0.752	0.64
3	DEBPROD	12.708	3.550	-9.158	-72.063	Decrease	0.752	0.64
3	ROS	-0.813	-0.234	0.578	71.154	Increase	0.752	0.64
3	ROA	-0.122	-0.073	0.049	39.990	Increase	0.752	0.64
3	DEBCOV	-0.064	-0.078	-0.014	-21.807	Decrease	0.752	0.64
3	FAPROD	6.646	1.742	-4.904	-73.790	Decrease	0.752	0.64
3	FCCOV	-0.308	-0.086	0.222	72.145	Increase	0.752	0.64

3	FAEQ	0.444	0.313	-0.130	-29.358	Decrease	0.752	0.64
4	DEBTA	0.853	0.849	-0.004	-0.471	Decrease	0.025	0.64
4	DEBPROD	1.993	1.957	-0.037	-1.835	Decrease	0.025	0.64
4	ROS	-0.007	-0.008	-0.001	-13.675	Decrease	0.025	0.64
4	ROA	-0.003	-0.004	-0.001	-46.858	Decrease	0.025	0.64
4	DEBCOV	-0.003	-0.005	-0.002	-58.325	Decrease	0.025	0.64
4	FAPROD	0.872	0.853	-0.019	-2.143	Decrease	0.025	0.64
4	FCCOV	-2.000	-1.950	0.050	2.484	Increase	0.025	0.64
4	FAEQ	2.389	2.336	-0.053	-2.217	Decrease	0.025	0.64
5	DEBTA	0.844	0.843	-0.001	-0.072	Decrease	0.004	0.64
5	DEBPROD	0.721	0.721	-0.001	-0.106	Decrease	0.004	0.64
5	ROS	0.039	0.038	0.000	-0.852	Decrease	0.004	0.64
5	ROA	0.045	0.045	0.000	-0.904	Decrease	0.004	0.64
5	DEBCOV	0.054	0.053	-0.001	-1.015	Decrease	0.004	0.64
5	FAPROD	0.091	0.091	0.000	0.147	Increase	0.004	0.64
5	FCCOV	0.333	0.332	-0.001	-0.415	Decrease	0.004	0.64
5	FAEQ	0.603	0.602	-0.001	-0.221	Decrease	0.004	0.64
6	DEBTA	7.043	2.089	-4.954	-70.338	Decrease	0.78	0.64
6	DEBPROD	1.131	0.662	-0.468	-41.434	Decrease	0.78	0.64
6	ROS	0.002	-0.034	-0.036	-1746.736	Decrease	0.78	0.64
6	ROA	0.013	-0.042	-0.054	-427.658	Decrease	0.78	0.64
6	DEBCOV	0.002	-0.064	-0.066	-3637.588	Decrease	0.78	0.64
6	FAPROD	0.159	0.132	-0.027	-17.070	Decrease	0.78	0.64
6	FCCOV	4.444	0.968	-3.476	-78.219	Decrease	0.78	0.64
6	FAEQ	2.583	0.779	-1.804	-69.832	Decrease	0.78	0.64
7	DEBTA	0.428	0.429	0.002	0.371	Increase	0.006	0.64
7	DEBPROD	5.217	5.189	-0.028	-0.539	Decrease	0.006	0.64
7	ROS	0.121	0.120	-0.001	-0.817	Decrease	0.006	0.64
7	ROA	0.010	0.010	0.000	-4.057	Decrease	0.006	0.64
7	DEBCOV	0.023	0.023	-0.001	-2.737	Decrease	0.006	0.64
7	FAPROD	10.776	10.712	-0.064	-0.593	Decrease	0.006	0.64
7	FCCOV	0.296	0.294	-0.002	-0.625	Decrease	0.006	0.64
7	FAEQ	1.598	1.590	-0.008	-0.498	Decrease	0.006	0.64
8	DEBTA	1.285	0.832	-0.453	-35.268	Decrease	0.764	0.62
8	DEBPROD	0.893	0.616	-0.278	-31.066	Decrease	0.764	0.62
8	ROS	0.027	-0.027	-0.054	-198.120	Decrease	0.764	0.62
8	ROA	0.039	-0.034	-0.074	-186.840	Decrease	0.764	0.62
8	DEBCOV	0.031	-0.056	-0.086	-281.534	Decrease	0.764	0.62

8	FAPROD	0.113	0.122	0.009	7.982	Increase	0.764	0.62
8	FCCOV	0.111	0.017	-0.094	-84.995	Decrease	0.764	0.62
8	FAEQ	-0.569	0.072	0.642	112.700	Increase	0.764	0.62
9	DEBTA	0.591	0.594	0.003	0.515	Increase	0.03	0.62
9	DEBPROD	0.568	0.567	-0.001	-0.203	Decrease	0.03	0.62
9	ROS	0.039	0.037	-0.002	-6.325	Decrease	0.03	0.62
9	ROA	0.041	0.038	-0.003	-7.177	Decrease	0.03	0.62
9	DEBCOV	0.069	0.065	-0.005	-6.565	Decrease	0.03	0.62
9	FAPROD	0.479	0.468	-0.011	-2.220	Decrease	0.03	0.62
9	FCCOV	0.410	0.397	-0.013	-3.091	Decrease	0.03	0.62
9	FAEQ	1.296	1.265	-0.031	-2.374	Decrease	0.03	0.62
10	DEBTA	0.623	0.624	0.001	0.221	Increase	0.02	0.64
10	DEBPROD	1.411	1.394	-0.018	-1.249	Decrease	0.02	0.64
10	ROS	0.045	0.043	-0.002	-3.942	Decrease	0.02	0.64
10	ROA	0.020	0.018	-0.002	-7.742	Decrease	0.02	0.64
10	DEBCOV	0.032	0.030	-0.002	-7.170	Decrease	0.02	0.64
10	FAPROD	0.043	0.045	0.002	3.731	Increase	0.02	0.64
10	FCCOV	0.793	0.777	-0.016	-2.032	Decrease	0.02	0.64
10	FAEQ	0.051	0.055	0.004	8.625	Increase	0.02	0.64

Table C.2 reports the managerial actions associated with the counterfactual indicator adjustments identified by HERA-PRO for each of the ten randomly selected CCC firms. Actions are obtained through the HERA-PRO action dictionary and represent the operational interpretation of the counterfactual recommendations.

Table C.2. Firm-Level Managerial Actions Derived from Counterfactual Solutions

FirmID	Action	Matches
1	Increase Production	3
1	Reduce Total Debt	2
1	Reduce Fixed Assets	2
1	Increase Current Assets	2
1	Reduce Operating Costs	1
1	Increase Equity	1
1	Reduce Financial Charges	1
2	Increase Production	2
2	Reduce Total Debt	2
2	Increase Current Assets	2
2	Reduce Operating Costs	1
2	Increase Equity	1

2	Reduce Financial Charges	1
2	Reduce Fixed Assets	1
3	Increase Production	4
3	Reduce Fixed Assets	3
3	Reduce Operating Costs	2
3	Reduce Total Debt	2
3	Increase Equity	1
3	Increase Current Assets	1
3	Reduce Financial Charges	0
4	Increase Production	2
4	Reduce Total Debt	2
4	Reduce Fixed Assets	2
4	Increase Current Assets	2
4	Increase Equity	1
4	Reduce Operating Costs	0
4	Reduce Financial Charges	0
5	Increase Production	2
5	Reduce Total Debt	2
5	Increase Current Assets	2
5	Reduce Operating Costs	1
5	Increase Equity	1
5	Reduce Financial Charges	1
5	Reduce Fixed Assets	1
6	Increase Production	3
6	Reduce Total Debt	2
6	Reduce Fixed Assets	2
6	Increase Current Assets	2
6	Reduce Operating Costs	1
6	Increase Equity	1
6	Reduce Financial Charges	1
7	Increase Production	3
7	Reduce Fixed Assets	2
7	Reduce Operating Costs	1
7	Reduce Total Debt	1
7	Increase Equity	1
7	Reduce Financial Charges	1

7	Increase Current Assets	1
8	Increase Production	2
8	Reduce Total Debt	2
8	Increase Current Assets	2
8	Reduce Operating Costs	1
8	Reduce Financial Charges	1
8	Increase Equity	0
8	Reduce Fixed Assets	0
9	Increase Production	3
9	Reduce Fixed Assets	2
9	Reduce Operating Costs	1
9	Reduce Total Debt	1
9	Increase Equity	1
9	Reduce Financial Charges	1
9	Increase Current Assets	1
10	Increase Production	2
10	Reduce Operating Costs	1
10	Reduce Total Debt	1
10	Reduce Financial Charges	1
10	Increase Current Assets	1
10	Increase Equity	0
10	Reduce Fixed Assets	0

ISSN (online) 2421-7158

CNR-IRCrES Working Papers

2026

N.6/2026 HERA-XAI. Explaining One-Year-Ahead Corporate Rating Forecasts Through Hierarchical Artificial Intelligence. Greta Falavigna

N.5/2026 Can AI enhance inclusion? A Participatory Approach to Defining and Evaluating New Technological Innovations. Maddalena Bartolini, Elena Ragazzi, Elia Moscoso Thompson, Caterina Rizzo, Alice De Luca, Monica Gori, Katia Lupinetti, Mauro Migliavacca, Lisa Sella, Michela Spagnuolo.

N.4/2026 HERA: Hierarchical Engine for Rating Assessment. A sector-specific hierarchical artificial intelligence framework for one-year-ahead corporate rating forecasting. Greta Falavigna.

N.3/2026 Il detective sociale: laboratorio di divulgazione delle scienze sociali per la scuola primaria. Antonella Emina, Anna Perin, Lisa Sella, Isabella Maria Zoppi.

N. 2/2026 Il calcio come infrastruttura sociale? Riflessioni ed esempi dal contesto italiano. Chiara D'Amato, Luca Tricarico.

N. 1/2026 Regolamentazione ambientale europea e impatti economici sulle imprese. Greta Falavigna, Giampaolo Vitali.

2025

N. 9/2025 Finanziamento delle comunità energetiche: attori, strumenti e aspettative dell'ecosistema ESG nel contesto italiano. Maksym Koltunov. Luca Tricarico, Adriano Bisello, Chiara D'Amato.

N. 8/2025 Importazioni, esportazioni e strategie green: il caso di studio della chimica in Italia. Greta Falavigna, Giampaolo Vitali.

N. 7/2025 Mapping international collaboration in Nanosciences and Nanotechnologies: a bibliometric and econometric study of South African research. Ugo Finardi, Jeanne C.M. Vallette d'Osia.

N. 6/2025 Dispersione salariale e performance nel calcio in Italia: un'analisi degli incentivi tra sforzo e risultato nella serie A (2010-2022). Mario Nosvelli, Matteo Colzato.

N. 5/2025 Designing a multiple case study approach for studying universities' community engagement. Andrea Orazio Spinello, Valentina Carazzolo, Ugo Finardi, Stefania Frongia, Valentina Ghibellini, Emanuela Reale, Andrea Vargiu.

N. 4/2025 Tools and Methods for Assessing Active Ageing and Inclusion: The Rima Experimental Project. Lisa Sella, Nicola Pollo, Valentina Lamonica.

N. 3/2025 Sinfonie urbane: il ruolo della ritmanalisi nella rigenerazione degli spazi. Il caso di Piazza San Venceslao a Praga. Vittoria Virgili, Marco Marucci.

N. 2/2025 From Single to Multi-Project Participation: SME Absorptive Capacity in Horizon 2020. Loreta Isaraj.

N. 1/2025 Invecchiamento e welfare responsabile: la "condizione anziana" nelle periferie di Roma Capitale. Cristiana Di Pietro.

2024

N. 5/2024 National frameworks for universities' community engagement: perspectives from Italy, France and the UK. Valentina Carazzolo, Ugo Finardi, Emanuela Reale, Andrea Orazio Spinello.

N. 4/2024 DAPIL – Digital archive of pilgrimage literature. Un archivio digitale per la letteratura medioevale di pellegrinaggio. Giuseppe Cuscunà.

N. 3/2024 Thesaurus e glossario ragionato della lingua della IA e della robotica. Thesaurus and reasoned glossary of AI and robotic language. Grazia Biorci.

N. 2/2024 Comprendere le trasformazioni di San Lorenzo, Roma. Un tentativo di ordine concettuale tra gentrificazione e politica urbana. Antonia Astore, Luca Tricarico.

N. 1/2024 How to measure the social acceptability of alternative environmental management solutions in wetlands and other ecosystems. Lisa Sella, Francesca Silvia Rota, Nicola Pollo.

2023

N. 5/2023 Laboratori creativi sperimentali a distanza per la valutazione. Isabella Maria Zoppi, Antonella Emina.

N. 4/2023 Eco-innovations and labor in the European automotive industry: an econometric study. Anna Novaresio.

N. 3/2023 Climate change adaptation planning: tools and methods for effective and sustainable decisions. Filippo Fraschini, Marco Pregnolato.

N. 2/2023 The effect of propensity to savings on rate of profit. Mario De Marchi.

N. 1/2023 UNESCO o non UNESCO? Quando il brand incentiva il turismo: il caso dei paesaggi vitivinicoli delle Langhe e del Monferrato. Orsola Borsani, Greta Falavigna.

2022

N. 5/2022 Logit statico, Logit dinamico e modelli hazard. Franco Varetto.

N. 4/2022 Evolution of Deep Learning from Turing machine to Deep Learning next generation. Greta Falavigna.

N. 3/2022 A Simulation Model of Technology Innovation of a Territory. Angelo Bonomi.

N. 2/2022 Technology and Environmental Policies. Angelo Bonomi.

N. 1/2022 Le donne marittime: fra stereotipi di genere, discriminazioni e scarse opportunità occupazionali. Barbara Bonciani, Silvia Peveri.

2021

N. 9/2021 Management of open access research infrastructures in large EU projects: the "CultureLabs" case. Andrea Orazio Spinello, Danilo Giglitto, Eleanor Lockley.

N. 8/2021 Francia-Italia: parole in campo. Intorno alla narrazione del Campionato del Mondo di calcio donne. Antonella Emina.

N. 7/2021 Covid-19 e rischio di insolvenza: il punto di vista del mercato azionario. Franco Varetto.

N. 6/2021 Institutional efficiency and budget constraints: a Directional Distance Function approach to lead a key policy reform. Greta Falavigna, Roberto Ippoliti.

N. 5/2021 Different waves and different policy interventions in 2020 Covid-19 in Italy: did they bring different results?. Mario Nosvelli.

N. 4/2001 On Search of a General Model of Technology Innovation. Angelo Bonomi.

N. 3/2021 Design and implementation of a web survey on the effects of evaluation on academic research. Andrea Orazio Spinello, Emanuela Reale, Antonio Zinilli.

N. 2/2021 An online survey on the effects of agile working in Italian Public Research Organisations. Serena Fabrizio, Valentina Lamonica, Andrea Orazio Spinello.

N. 1/2021 Technology Transfer Activities in Universities and Public Research Organizations: A Literature Overview. Ugo Finardi, Rolfo Secondo, Isabella Bianco.

2020

N. 12/2020 Unexpected loss multiperiodale e pricing del rischio di credito. Franco Varetto.

N. 11/2020 La ricerca in Nanotecnologie e Nanoscienze in Italia: spesa del settore pubblico e aree tematiche prevalenti. Ugo Finardi, Andrea Orazio Spinello.

N. 10/2020 Persistent fast growth and profitability. Lucio Morettini, Bianca Potì, Roberto Gabriele.

N. 9/2020 Binomio *Burnout* e *Mindfulness* nelle organizzazioni. Alcuni studi e scenari di applicazione. Oriana Ippoliti, Riccardo Briotti, Bianca Crocamo, Antonio Minopoli.

N. 8/2020 Innovation and communication of companies on Twitter before and during COVID-19 crisis. José N. Franco-Riquelme, Antonio Zinilli, Joaquín B. Ordieres-Meré and Emanuela Reale.

N. 7/2020 The proposal of a new hybrid methodology for the impact assessment of energy efficiency interventions. An exploratory study. Monica Cariola, Greta Falavigna.

N. 6/2020 The technology innovative system of the Silicon Valley. Angelo Bonomi.

N. 5/2020 Storia dell'industria delle macchine utensili in Piemonte dalle origini alla seconda guerra mondiale. Secondo Rolfo.

N. 4/2020 Blockchain e Internet of Things per la logistica Un caso di collaborazione tra ricerca e impresa. Edoardo Lorenzetti, Lucio Morettini, Franco Mazzenga, Alessandro Vizzarri, Romeo Giuliano, Paolo Peruzzi, Cristiano Di Giovanni.

N. 3/2020 L'impatto economico e fiscale di un evento culturale: misure e scala territoriale. Giovanna Segre, Andrea Morelli.

N. 2/2020 Mapping the tangible and intangible elements of the historical buildings and spaces. Edoardo Lorenzetti, Nicola Maiellaro.

N. 1/2020 Il lavoro agile negli enti pubblici di ricerca. Emanuela Reale, Serena Fabrizio, Andrea Orazio Spinello.

2019

N. 6/2019 Women's candidatures in local elections: does the context matter? Empirical evidence from Italian municipalities. Igor Benati, Greta Falavigna, Lisa Sella.

N. 5/2019 Research activities in Nanotechnologies and Nanosciences: an analysis of Piedmont's nanotech research system. Ugo Finardi.

Numeri precedenti/Previous issues

How can artificial intelligence do more than predict a company's credit rating? Can it also suggest the concrete actions needed to improve it?
This Working Paper presents HERA-PRO, an innovative decision-support framework that extends the capabilities of rating assessment from prediction to prescription. Building on recent advances in explainable artificial intelligence (XAI), HERA-PRO not only estimates the rating a firm is likely to receive but also identifies the financial adjustments most likely to lead to a stronger rating outcome.
Combining hierarchical machine learning models, sector-specific analysis, and counterfactual optimization techniques, the framework explores alternative scenarios and determines the minimum changes required to cross critical rating thresholds. Rather than simply asking "What rating will be assigned?", HERA-PRO addresses a far more strategic question: "What should change to obtain a better rating?"
By integrating prediction, explanation, and actionable guidance within a single methodology, this work introduces a novel perspective on corporate credit risk assessment. It demonstrates how artificial intelligence can become a practical tool for decision-making, helping firms, financial institutions, and analysts translate complex data into targeted strategies for improvement.
Positioned at the intersection of machine learning, credit rating analysis, and prescriptive analytics, HERA-PRO offers both a methodological contribution and a forward-looking vision of how AI can support more informed and effective financial decisions.

CNR - Consiglio Nazionale delle Ricerche
IRCrES - Istituto di Ricerca sulla Crescita Economica Sostenibile