sustainability

MDPI

Article

Improving Crowdfunding Decisions Using Explainable Artificial Intelligence

Andreas Gregoriades [1] and Christos Themistocleous [2,*]

[1] Department of Communication and Marketing, Cyprus University of Technology, Limassol 3036, Cyprus; andreas.gregoriades@cut.ac.cy

[2] Persuasive Tech Lab, Department of Communication and Marketing, Cyprus University of Technology, Limassol 3036, Cyprus

* Correspondence: c.themistocleous@cut.ac.cy

Abstract: This paper investigates points of vulnerability in the decisions made by backers and campaigners in crowdfund pledges in an attempt to facilitate a sustainable entrepreneurial ecosystem by increasing the rate of good projects being funded. In doing so, this research examines factors that contribute to the success or failure of crowdfunding campaign pledges using eXplainable AI methods (SHapley Additive exPlanations and Counterfactual Explanations). A dataset of completed Kickstarter campaigns was used to train two binary classifiers. The first model used textual features from the campaigns' descriptions, and the second used categorical, numerical, and textual features. Findings identify textual terms, such as "stretch goals", that convey both elements of risk and ambitiousness to be strongly correlated with success, contrary to transparent communications of risks that bring forward worries that would have otherwise remained dormant for backers. Short sentence length, in conjunction with high term complexity, is also associated with campaign success. We link the latter to signaling theory and the campaigners' projection of knowledgeability of the domain. Certain numerical data, such as the project's duration, frequency of post updates, and use of images, confirm previous links to campaign success. We enhance implications through the use of Counterfactual Explanations and generate actionable recommendations on how failed projects could become successful while proposing new policies, in the form of nudges, that shield backers from points of vulnerability.

Keywords: crowdfunding; startups; nudge theory; risk communication; machine learning; counterfactual explanations; SHAP

check for updates

Academic Editor: Hyunchul Ahn

Received: 22 November 2024
Revised: 24 January 2025
Accepted: 26 January 2025
Published: 7 February 2025

Citation: Gregoriades, A.; Themistocleous, C. Improving Crowdfunding Decisions Using Explainable Artificial Intelligence. *Sustainability* **2025**, *17*, 1361. https://doi.org/10.3390/su17041361

1. Introduction

Crowdfunding has provided a modern approach for financing entrepreneurial projects by requesting funds a priori from prospective customers. Contrary to traditional financing methods, like venture capital and bank loans, crowdfunding companies are not obligated to provide information pertaining to the background, forecasting, or financial viability of the idea upon which assessments are based and decisions of financing by the public are made. Instead, prospective backers make judgment calls based on information that is descriptive in nature, like the idea of the product, its description, its visual depiction, and the number of others in support of the pledge.

The global crowdfunding market has an estimated value of USD 20 billion [1], yet the societal and financial sustainability of the ecosystem is still under development. For example, fraudulent crowdfunding campaigns that acquired funds but did not deliver a final product while cutting communication ties with backers were the subject of recent fraud-funding investigations [2]. The rate of project failures (where campaigners received the money but did not deliver a product) between 2009 and 2012 was a reported 3.7% [3], registering an increase closer to the 10% mark [4] over recent years. Extrapolating such percentages within the overall crowdfunding market, billions are lost to failure to deliver final products either due to fraudulent solicitation or simply due to cost and planning mismanagement by the campaigner. These, in conjunction with the fact that no guarantee is provided by platforms for a refund in cases of unfulfilled projects, accentuate the need to examine predictors of success and failure in crowdfunding pledges. Such investigations generate implications for enhancing the sustainability of this financial ecosystem by improving decisions by both backers and campaigners for the reduction of opportunity costs (supporting a project that fails to receive the pledged amount and thus does not proceed with idea implementation) and minimization of channeled funds to projects that are prone to fail in receiving the pledged amount.

In this paper, we examine the success and failure of crowdfunding pledges using eXplainable AI techniques (XAI), an emerging area of crowdfunding research. Specifically, we investigate how the interactions between textual, numerical, and categorical factors of campaigns influence the campaigns positively (having managed to successfully receive the pledged amount) or negatively (failing to do so) using XAI. Through this, we build upon previous work that examines crowdfunding from a financial sustainability point of view [5–7]. Such examinations delineate questions on what markers backers should use to reduce the chances of supporting projects that are riskier and more susceptible to failing. Additionally, we examine questions that aim to help campaigners structure their campaign descriptions more accurately and informatively, using the necessary complexity of language that cognitively resonates with prospective backers and provides procedural information. As such, we offer actionable recommendations to campaigners to enhance the success rate of their campaigns while proposing new policies that shield consumers from points of vulnerability.

Previous research aimed to delineate success and failure predictors of crowdfunding campaigns using available parameters like the overall amount pledged, the number of supporters, and the length of the campaign. For example, longer fundraising periods are correlated with campaign failure [8]. Having a formal website to complement the pledge is correlated with success [9], and the inclusion of images and videos in the fundraising campaign positively influences its outcome [10,11]. Interestingly, such measurements are based primarily on numerical and categorical data found in crowdfunding datasets, yet Natural Language Processing (NLP) can make use of textual data available in the project's description. Although we note the emergence of some XAI applications in crowdfunds (i.e., [12]), the use of XAI to explain the influence of the communication style in campaign descriptions on project success/failure is not addressed. Applications of AI techniques in crowdfunding [13] showcase the analysis of textual data but do not address linguistic styles and language complexity of the project's description nor provide recommendations on changes that can improve their chances of getting funded. As such, this research aims to fill this gap through the application of language processing and XAI and proposes a methodology that can infer linguistic predictors of crowdfunding campaigns' success (or failure) and recommend ways to change non-successful projects to successful ones using textual features from campaign descriptions.

The research goal that drives this work is to investigate how XAI can be used to inform crowdsourcing campaigners so as to improve their chances of success. To that end, this paper makes the following contributions. First, this study introduces the use of XAI in crowdsourcing campaigns and, secondly, applies both global (SHapley Additive exPlanations—SHAP [14]) and local (Counterfactual [15]) explanations to investigate the most important properties of campaigns that contribute to success. In this regard, two classification models (XGBoost) were developed, one solely using textual data and the other using all three data types. Specifically, we investigate the application of SHAP to infer the global patterns of the models and Counterfactual Explanations to examine the smallest change to the feature values of campaigns that change the prediction of the model (see [16]). Kickstarter datasets were used to train two machine learning models, which were subsequently explained by the aforementioned XAI techniques. The textual features were pre-processed and vectorized, and the numerical features were normalized.

Our results show different factors that contribute to crowdfunding success, such as certain textual features and linguistic styles, as well as certain crowdfunding categories that are correlated with success. A key finding includes specific phrases in campaigns that communicate both elements of risk and innovation to attract more financial support than more transparent statements on the involved risks. We argue that risk communication transparency in campaign descriptions can bring forth worries that would have otherwise remained dormant by backers ultimately preferring to support projects that are less risk transparent. We argue for the adoption of mechanisms that promote risk transparency and specifically nudges that prompt backer deliberation prior to investing. Lastly, our application of Counterfactual Explanations in the crowdfunding arena aims to allow campaigners to test their descriptions, which, if unsuccessful, offer actionable recommendations on how to make them successful, through a clearer communication style.

The rest of this paper is structured as follows. A literature review is provided on success/failure predictors of crowdfunding campaigns along with previously employed language processing techniques in this context. This is followed by an elaboration of the empirical approach, including pre-processing and data-cleansing practices to mitigate ML biases [17]. SHAP and Counterfactual Explanation results are summarized, demonstrating specific words in pledge descriptions that are associated with crowdfunding success or failure. We analyze the generated words and produce discussions relating to implications for both crowdfunding campaigners and backers, along with avenues for future research.

2. Literature Review

2.1. Crowdfunding Success Factors

Since its introduction, crowdfunding has emerged as a viable finance tool for startup companies. Contrary to bank loans that require formal profitability statements or venture capital that are direction-steering and competitive, crowdfunding allows a majority of investors (backers) to share the risks of investment and bring to light new entrepreneurial propositions. The range of projects is vast, with Kickstarter having 15 distinct categories, from arts and theatre to electronics and technology. The majority of platforms employ the all-for-nothing model where money is paid out to entrepreneurs only when the funding goal has been reached. As such, it is important for campaigners to convince a significant number of investors of the project's viability, the importance of the idea, and the ability to meet set deadlines. Figure 1 provides a visual summarization of the crowdfunding process.

The lack of financial and forecasting reports in crowdfunding campaigns arguably exacerbates information asymmetry between campaigners and backers, with the literature interested in the importance of signals for achieving campaign success. Using signaling theory, Koch and Siering [18] identify how certain signals projected by campaigners affect investment decisions by backers. For example, the ability to communicate aspects of trustfulness and quality were found to be more important than signals relating to the campaigner's experience and popularity. The use of such signals by campaigners capitalizes on what decision-makers (backers) are interested in when in deliberation.

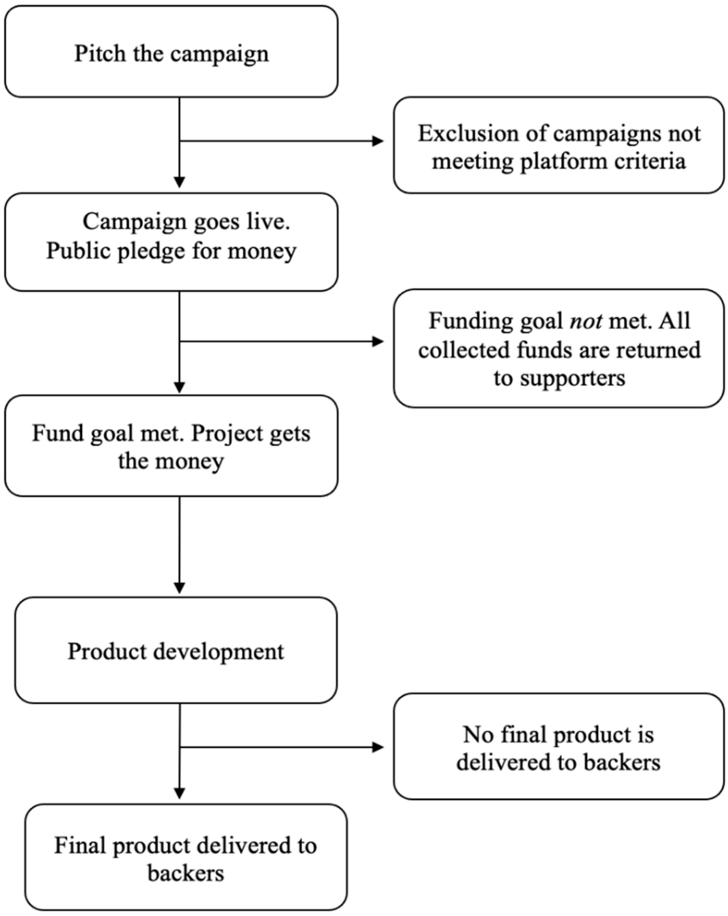

Figure 1. The crowdfunding process.

In relation to pledging conditions, both the funding goal and funding period are systematically documented as important markers of project success. Specifically, high funding goals are associated with success [3,8], with Kickstarter's fulfillment report mentioning that campaigns targeting below the 1000USD mark have significantly higher project failure percentages than ones requesting funds above that mark [4]. Funding periods are also crucial. Xiao et al. [19] identify longer-than-average periods to be correlated with failure, a finding confirmed by follow-up investigations [11].

Important success factors are also linked to the way the proposed idea is presented. In relation to audiovisual means, the use of images in the pledge is correlated with success [9]. Kaminski and Hopp [20] and Liu et al. [8] also identify the use of videos depicting the product-in-action as another success predictor. Relating to textual features in the projects' descriptions, subtitles, and titles, Zhang et al. [21] demonstrate that the description of projects that use exploratory language (words like discover, explore, experiment) received less financial support from backers compared to ones with exploitative language (execute, exploit, implement) in their description. The use of exploitative language can be aligned with textual features that evoke enthusiasm in backers [20] and thus assist with meeting funding goals. Emotional appeal [18] and the use of descriptive (but not exhaustive) language [9], both serve as additional success predictors. Table 1 offers a summarization of key papers and success/failure predictors of crowdfunding campaigns.

Table 1. Key literature examining success/failure predictors of crowdfunding campaigns.

Paper	Categories Associated with Campaign Success/Failure
Liu et al., (2023) [8]	Use of Images and Videos; Duration of Fundraising; Pledged Amount
Koch & Siering (2019) [18]	Emotional Appeal; Duration of Fundraising; Pledged Amount; Previous Project Experience
Zhang et al., (2023) [21]	Use of Exploitive Language
Xiao et al., (2014) [19]	Pledged Amount; Consistent Communication with Backers; Fewer Reward Tiers
Yeh et al., (2019) [9]	Use of Descriptive Language; Use of Images; Frequent Updates; Having a Website
Kaminski & Hopp (2020) [20]	Use of Enthusiastic Language; Use of Product-In-Action Videos
Barbi & Bigelli (2017) [11]	Length of the Description; Pledged Amount

2.2. Explainable AI in Crowdfunding

From the above, it is evident that to succeed in crowdsourcing campaigns, it is not enough to have an innovative idea. The description of the project and the way it is communicated is critical to its success. The application of machine learning (ML) and artificial intelligence (AI) in businesses across sectors is aimed at improving their competitiveness and performance through a range of applicable solutions. Learning to recognize patterns, AI can make predictions on what products can be more appealing to certain customer groups [22], nudge consumers for behavioral change [23], identify which crowdfunding investments are riskier than others [24], and serve as a decision support tool for effective and efficient business management [25]. Specifically in crowdfunding, AI has been applied for the identification of success and failure markers [20,26]. Yet, despite its potential, the lack of transparency of certain ML models has grown from a concern to a problem, and the inability to understand the logic followed by an ML model posits significant accountability and fairness questions.

The internal workings of complex ML models, which are invisible to the user, contribute to the black box problem and the lack of human understanding of the AI's processes when delivering results. XAI offers a pragmatic solution in shedding light inside these black boxes by explaining the reasoning behind the behavior and output of the system to the user [26]. XAI methods can be classified either as local or global, inherent or post hoc methods. Local/global refers to whether they aim to explain the whole logic of a model and follow the entire reasoning leading to all the different possible outcomes or to explain the reasons for a specific decision or single prediction. Artelt and Gregoriades [16] (p. 1) argue that XAI assists in interpreting and understanding complex models, contributing to their bias-free validation and increasing decision-makers' confidence and trust in their predictions. As mentioned, a prominent class of explanation methods are post hoc approaches, which refer to approaches applied to a trained model that makes predictions. Examples of post hoc methods include LIME [27], SHAP [14], and Counterfactual Explanations [15]. Such methods extract local explanations from complex models.

In this study, we use SHAP values and Counterfactual Explanations, both of which are common techniques for generating explanations of how certain features impact the model's prediction. SHAP was originally developed to estimate the importance of an individual player in a collaborative game. It aimed to distribute the total gain or payoff among players in a game, depending on the relative importance of their contributions to the final outcome. Shapley values are thus rationalized as a reasonable allocation of feature importance given a particular model output. Counterfactual Explanations take explanation one step further by identifying the minimum changes required to a case so as to change the output of a model to the desired state (for example, project success). Through the use of classification and XAI techniques, our research aims to provide insights regarding the causes of success and failure in crowdfunding campaigns and address the research questions expanded below.

2.3. Development of Research Questions

An integral part of crowdfunding is the ability to communicate risky yet ambitious ideas to prospective backers. Risk communication or risk disclosure in crowdfunds is key for backers, as close to 80% of Kickstarter campaigns fail to receive the pledged amount [4]. Nevertheless, there are mixed findings in the literature on how financiers and consumers react to risk disclosures. An increase in keywords conveying risk inherent to stock returns by a corporate firm, as an act to enhance transparency, increases investors' risk perceptions compared to a company that avoided its use and follows a positive framing approach [28]. On other occasions, like advertisements on remedies, the risk disclosures had the opposite effects, decreasing the severity of risk for consumers under certain conditions [29]. Kim et al.'s [30] work focused on risk communication in crowdfunding, concluding that direct risk disclosures by campaigners can lead to inferior funding outcomes. We build upon their work by assessing textual features that convey risk and their association with campaign success. Welkenhuysen et al. [31] posit that the framing effect was more prevalent in verbal and textual communication rather than in numerical iterations that are more neutral in their presentation. Based on these, we would expect positively framed textual features that highlight ambitiousness in conjunction with riskiness, to be met with more funds specifically because of prospective backers' interest in innovative products. We focus on answering the question of whether textual features associated with both innovation and risk are linked to campaign success. Formally, we phrase this as the following:

RQ1: Are textual features in campaign descriptions that communicate risk in conjunction with ambitiousness associated with success or failure?

Similar to the proposed research, previous examinations in crowdfunding analyzed the language used in crowdfunds. Findings focus on the emotional appeal of the description [18] as well as how enthusiastic the language is based on certain index scores [20]. Other examinations focused on the degree to which the language conveyed exploration or exploitation [21,32], yet, to our knowledge, none have examined the complexity of the language used in campaign descriptions. Yeh et al.'s [9] examination showed an interest in the length of the description, yet this is different from understanding overall complexity based on the difficulty of certain terms, for example, and whether they are positioned in long sentences. We note that the latter will require more cognitive effort for prospective backers to follow but might also be perceived as a signal of domain knowledge by the campaigner. By exploring this relationship between the combination of sentence length and language complexity in crowdfunding descriptions, we want to enhance the resolution of previous research that focused primarily on descriptive language analysis. As such, RQ2 considers the following:

RQ2: Are textual features relating to sentence length and language complexity in crowdfunding campaign descriptions associated with project success or failure?

As certain textual features are associated with success and others with failure, questions on what features can change in order to turn an unsuccessful project into a successful one remained unanswered with traditional means. Platforms also lack tailored feedback on how the product description could have been rephrased to boost support for backers, especially for projects that were close to raising the whole pledged amount. For example, Kickstarter may provide arbitrary explanations as to why a project was unsuccessful, such as a "poor description of the idea", yet such explanations do not help current and prospective campaigners to improve their chances of being successful, nor provide guidelines for improved and transparent communication of their ideas with backers. In certain cases, the factor that needs changing might not even be changeable, for example, the age of the project proposer. Yet, there are changeable dimensions, such as the complexity of the language in the description. From XAI, Counterfactual Explanation frameworks, like Diverse Counterfactual Explanations (DiCE; see [15]), provide what-if scenarios and thus can be applied to provide information on what needs to change to achieve this desired effect. The present study aims to test the application of DiCE in a crowdfunding setting to pinpoint actionable recommendations for an unsuccessful project and changes that can be made to convert it to a successful one. As such, this approach aims to examine the third research question, which reads as follows:

RQ3: Can an unsuccessful crowdfunding project be converted into a successful one?

RQ1–RQ3 were examined using a textual-data-only model (Model 1). We complemented the textual features model with a second model (Model 2) that also considers categorical and numerical data. The latter serves as the confirmatory part of this research and an opportunity to test and confirm previous findings using a different methodology, that of XAI. We test previous findings on three fronts. First, we examine whether the use of images in campaigns, as pointed out by Liu et al.'s [8] metanalysis, is associated with success. Secondly, we examine whether the duration of a campaign, as identified by Koch and Siering [18], especially longer campaigns, is associated with failure. Thirdly, we analyze whether consistent communication with backers through updates in the form of posts, as proposed by Xiao et al. [19], is associated with success.

RQ4: Are images, campaign length, and communication consistency between campaigners and backers predictors of project success?

3. Empirical Approach

The methodology employed to answer our research questions combines machine learning with explainable AI. The former is used to train two binary classifiers for predicting project success or failure, using different feature sets, such as features that emerge from the textual description of the projects alone. The second set of features combines textual, numerical, and categorical features, such as number of comments by people, duration of project, and project category. The second part of our method aims to explain the trained models using XAI techniques (specifically SHAP and Counterfactual Explanations). Figure 2 summarizes the methodology, which includes data management and model training.

3.1. Data

The data used in this work consist of textual descriptions of projects scraped from the Kickstarter website, constituting dataset 1, and project-related categorical, numerical, and textual data from the Web Robots Kickstarter datasets (dataset 2). The textual data refer only to projects that had a project description with more than 100 words and were not canceled. The integration of the 2 datasets was performed using project ID mapping. Initially, the raw text in dataset 1 was separated by a semicolon, and thus, it had to be split and expanded into new columns. Rows with no language information were dropped. Dataset 2 was originally in txt format and was thus read line by line and parsed to extract the pid and the description of the projects. After pre-processing the 2 datasets, they were integrated using pid, which is a common field in both of them. Rows that had non-integer pid were dropped. Projects that were not in English were also dropped.

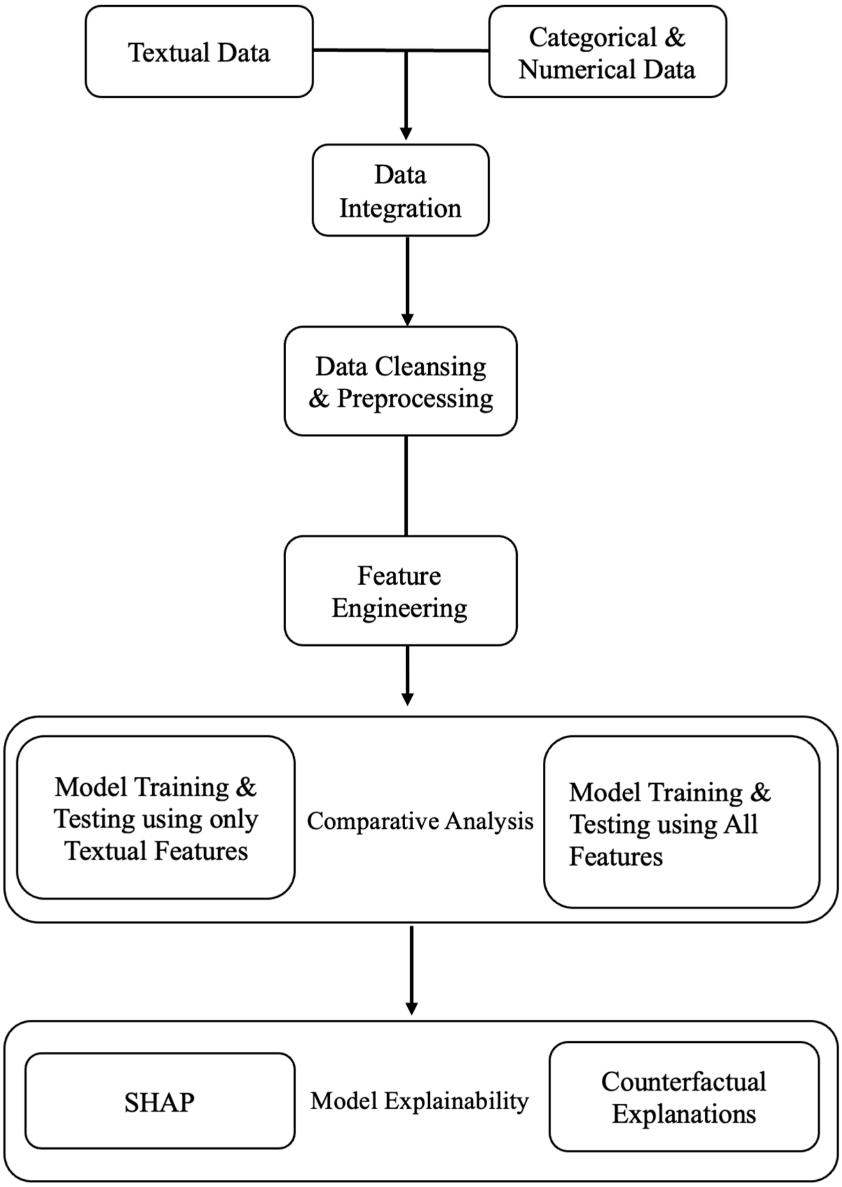

Figure 2. Methodological steps. Data management, model training, and explainability.

The compiled dataset consisted of 9258 campaigns finished in the specified timeframe by the projects' owners (projects did not diverge from plan) and allowed the application and evaluation of explainable AI techniques using textual features. Within this dataset, 70% of projects were successful, or they managed to reach their funding goal, and 30% of projects were unsuccessful in meeting their goals.

3.2. Pre-Processing and Feature Engineering

Two sets of binary classifiers (XGBoost, Logistic Regression, Random Forest) were trained, the first utilizing only textual features and the second using textual, numerical, and categorical features. In both cases, the data were split into training (80%) and testing (20%) sets. The train–test split of the data was performed using stratification to ensure that the proportion of samples from each class (success/failure) in the train and test sets was the same as in the original dataset. The textual descriptions of the projects were pre-processed to eliminate unnecessary text like URLs, non-ASCII characters, and numbers.

During the feature engineering step, the textual features underwent transformation using TF-IDF (Term Frequency—Inverse Document Frequency Technique) so as to convert these into numerical values (vectorization of text). Three variants of TF-IDF were used in this step with regard to text tokens and the use of unigrams, bigrams, and trigrams during TF-IDF vectorization. The vectorization technique that yielded a model with better performance was selected. Additionally, textual features underwent a comparison between the use and no use of stopword removal. The latter eliminates common English language words, such as “the”, “is”, and “and”, or custom stopwords that convey little meaning to the problem at hand (e.g., the word Kickstarter). Thus, only words that were more informative about success/failure were selected by the TF-IDF (for example, “definitely”, “would like”, etc.). Both techniques were evaluated, and the one with the best classification performance was selected.

Additional feature engineering steps included the generation of features based on the text’s linguistic properties, such as language complexity and extraction of part-of-speech features from the original text, since both have been reported to improve classification performance [33,34]. Therefore, linguistic analysis of the text was performed to identify features such as the number of adverbs, verbs, and nouns in the project description that could be informative of project success and failure. This was performed through the use of speech tagging using the spaCy library (an open-source library for advanced natural language processing). With regard to language complexity, language readability features were extracted from the text to further enhance the feature set used to train the classifiers. Readability is the ease with which a reader can understand a written text, and it can be extracted automatically using different techniques and Python libraries. Readability scores include the Dale–Chall, Flesch–Kincaid, and McLaughlin’s SMOG. In this work, all scores were extracted from the text and used as additional textual features.

A second set of classifiers was built using the above textual pre-processing and feature engineering step of the projects’ descriptions in combination with categorical and numerical features. For example, categorical features included the project’s category (e.g., technology, games) or whether the project had a product-describing image (or not). These were one-hot-encoded using CountVectorizer. Pre-processing for the numerical features included outlier elimination and missing value imputation (e.g., number of posts, project duration).

Feature selection steps were applied to reduce the number of features used while training the models. Several techniques exist to achieve this, namely filter, wrapper, embedded, and ensemble [35]. Here, a filter technique was used due to its performance and cost-effectiveness. In this work, a correlation-based technique was used for the trained model using numerical features and chi-square in combination with correlation-based when numerical (TF-IDF vectors) and categorical features (e.g., project type). The selected features had the highest correlation with the class and the lowest correlation between them. In the case of categorical features, the chi-square metric was used, and the features were selected based on the best chi-square results. In the case of mixed categorical and numerical features, the best features from each technique were combined, and the final feature set was found based on features with the highest accuracy of a naive Bayes classifier.

3.3. Model Training

The pre-processed data were used to train and test 2 sets of binary classifiers. XGBoost was used since it constitutes a popular ensemble machine learning technique that utilizes multiple gradient-boosted decision trees and has been previously employed in relevant settings [36]. Boosting is a technique that attempts to build a strong classifier from the number of weak classifiers. During XGBoost training, an initial model is built from the training data and complemented by a subsequent model that corrects the errors from the first one. This procedure was repeated for all data. XGBoost has been preferred over other machine learning techniques due to its incorporation of L1 (Lasso) and L2 (Ridge) regularization, that mitigate overfitting by penalizing complex models. L1 encourages sparsity by driving some coefficients to zero, while L2 penalizes large coefficients for smoother models [37]. Additionally, XGBoost efficiently handles missing values, supports parallel processing, and performs very well with medium and small datasets.

To verify the superiority of the XGBoost approach in predicting campaign success/failure, two alternative techniques were used for comparison, namely Random Forest and Logistic Regression. These were selected due to their popularity in binary classification tasks in the FinTech domain [13]. To evaluate the trained models, the area under the Receiver Operating Characteristic (ROC) curve AUC and F1 were used as performance metrics. These are used in binary classification tasks and in situations when the data are imbalanced [38]. The hyperparameters of the XGBoost model were optimized using grid search (max tree depth, learning rate, scale_pos_weight, reg_alpha, lambda; the last 2 refer to L1 and L2 regularization). To address the data imbalance, the scale_pos_weight hyperparameter was evaluated using different scores. Similarly, the Random Forest and Logistic Regression models had also undergone hyperparameter tuning. The comparative analysis results showed that the AUC and F1 scores of the XGBoost were superior to the baseline models, as seen in Table 2, and thus, the XGBoost was selected in subsequent steps.

Table 2. Comparative evaluation of the optimized XGBoost model against two other binary classification models.

	AUC (%)		F1 (%)	
Binary Models	**Textual Features Only**	**All Features**	**Textual Features Only**	**All Features**
XGBoost	78	92	78	91
Random Forest	74	86	70	87
Logistic Regression	69	78	71	81

3.3.1. Global Explanation Using SHAP Analysis

The trained XGBoost models were initially explained using tree-SHAP (Shapley Additive Explanations), a widely used post hoc explainability technique for calculating feature attributions [39]. It is model agnostic and, thus, can be used for both local and global feature attribution. The Shapley value of a feature is determined as the average of its contributions across all possible permutations of a feature set. In this work, the tree-based SHAP summary plot is used since we are using a tree-based classification model, and we are interested in identifying which features SHAP assigns importance to by calculating its average marginal contribution to the prediction across all potential combinations of features [40,41]. This ensures a fair and comprehensive assessment of each feature's influence, considering every possible interaction. Furthermore, the choice of the SHAP method was attributed to recent insights identifying SHAP values to align closely with human decision-making processes and intuitive judgments [42]. Other studies have highlighted SHAP's usefulness for feature selection, showing that it can lead to improved model performance [16,43].

3.3.2. Counterfactual Explanations Analysis

To be able to answer questions RQ3 and 'how can a previously unsuccessful project become successful?', we employed Counterfactual Explanations with the trained model. An example that illustrates how Counterfactual Explanations function is as follows. Consider a company that campaigns for crowdfunding yet fails to acquire the full amount. It provides decision-makers with feature-perturbed versions of the failed project but with recommended changes that could have led to its success. In essence, it offers what-if explanations for model output and can complement SHAP results. These what-if scenarios allow users to understand how small changes in input features affect the outcome, improving explainability and fostering confidence in the model's decisions (see [44]).

Counterfactual Explanations offer certain advantages over other XAI techniques like LIME, feature importance, or partial dependence plots. Counterfactuals are more actionable as they demonstrate exactly what changes in input features would result in a different outcome, providing clear guidance for users looking to influence the prediction [16]. Compared to LIME, which uses perturbed data to approximate local decision boundaries, Counterfactual Explanations identify minimal and realistic changes needed to alter an outcome, making them more precise and actionable in contexts like crowdfunding, where explanations are crucial for ensuring compliance and transparency in decision-making processes. Darias et al. [45] highlight that Counterfactual Explanations are especially intuitive for non-experts, echoing their applicability in financial decision-making processes.

4. Results

In answering the research questions, the first XGBoost model employed solely textual features, used SHAP and Counterfactual Explanations and aimed to enhance resolution on the impact of language complexity and used language for communicating risk. The second XGBoost model employed all types of data in an attempt to inductively identify patterns conveyed through categorical and numerical pieces of information. The objective was to delineate characteristics of crowdfunding campaigns associated with either success or failure through the existence of images, the length of the campaign, and the presence of consistent communication through updates via posts. SHAP analysis was used to assess these. We elaborate on the results of each model below.

4.1. Model 1—Textual Features

Of the two XGBoost models that were trained, the first aimed to scrutinize the textual features alone and achieved an AUC of 78% and an F1 score of 78%. The model predicted the probability of success or failure of a project with good accuracy using purely textual data. The SHAP summary plot of this model (Figure 3) shows how certain textual features that emerge from TF-IDF, linguistic analysis, and POS tagging are contributing positively or negatively to success/failure.

When reading the SHAP plot, we note that red signals a high feature value while blue indicates a low feature value. Clustering of the red dots on the left side denotes a negative contribution and pushes the prediction away from the target—clustering of red on the right denotes a push of the prediction toward the target and thus has a positive contribution to success. In assessing language complexity and sentence structure of campaign descriptions, the SHAP plot reveals relevant findings. Linguistic scores like dale_chall and flesch-kincaid, which were found to be significant, assess language complexity and sentence complexity, respectively. Appendix A provides the formulas of the dale_chall and flesch-kincaid scores. Findings show that a high dale_chall score is associated with success, while a high flesch-kincaid score is associated with failure. In essence, when the vocabulary of the language increases and both the sentence length and the syllable count per word decrease, the probability of success increases.

Results here show that shorter sentences that utilize complex language are associated with success. This is arguably attributed to how perceptions of prospective backers are shaped. Language that is complex yet comprehensible and includes technical elements within short sentences is associated with the author's knowledge of the domain (i.e., [18]), an argument we expand on in the discussions section. Here, we expand on Yeh et al.'s [9] results, who identified that longer descriptions (yet not overwhelmingly long ones) are associated with success. Specifically, not only description length but also a balanced use of short sentences and complex/technical words can provide an effective communication of the idea in a comprehensible manner to the backer, thus enhancing the project's success.

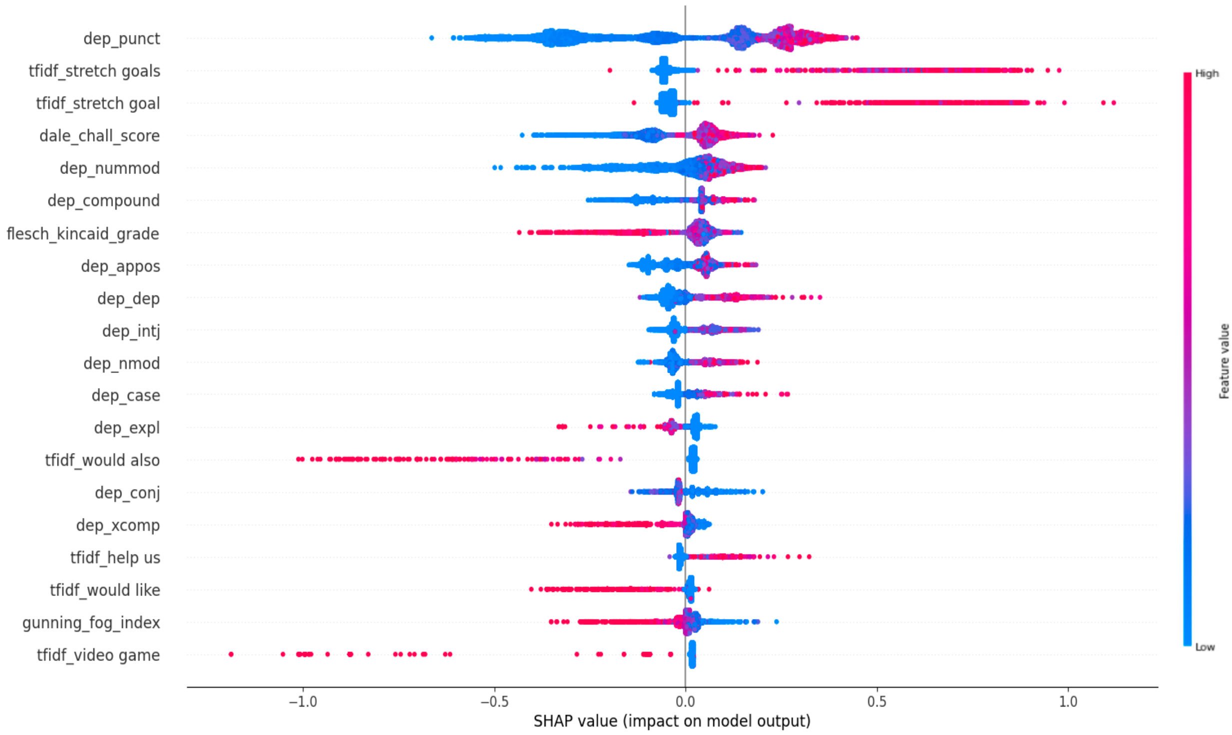

Figure 3. SHAP plot of the first XGBoost model trained using only textual features. Red indicates a high score for the feature, and blue indicates a low score.

In terms of risk communication, TF-IDF features that are most influential in terms of success include the phrase "*stretch goal(s)*". Stretch goals are associated with organizational performance and have a high-risk, high-reward character [46]. Findings here indicate that the expression of willingness to undertake risks in order to implement an ambitious project was met with the backer's support. This lies in the definition of the term *stretch goal(s)*. The term refers to a goal-setting business strategy that is characterized by ambitiousness, risk-taking, innovation, and the positioning of organizations outside their comfort zone [47,48]. It is not surprising that backers on crowdfunding platforms are interested in entrepreneurial projects that push boundaries and communicate ambitiousness and risk-taking in their descriptions. Here, we note that positively framed language that promotes innovation along with risk, in the form of stretch goals, is found to have a strong correlation with campaign success, contrary to language that communicates risks in a transparent manner, which had the opposite effects [30].

The SHAP plot also notes that phrases like "*would like*", "*help us*", and "*would also*" are significant in their association with success or failure. Specifically, we find that the projects that used language that was less direct and more exploratory in nature ("*would like*", "*would also*") were associated with failure, while language that was more direct to the pledge and with exploitative character ("*help us*") was associated with project success. Results demonstrate a dichotomy of exploratory and exploitative language and identify a tendency of backers to indulge pledges that use direct imperative language ("help us") instead of ones using a milder tone. The findings here expand the explore/exploit word pool of previous research [32] and align with those of Zhang et al. [21], who found exploitative language to be more effective in crowdfunding pledges than the use of exploratory language.

Our analysis brings forward additional findings. Another textual feature associated with success includes *dep_appos*, which refers to the statements that define or describe a noun, for example, "Chris, my brother, arrived". This noun-explanatory feature relates to enhancing clarity for the reader. This is very relevant for campaigners who must use technical terms to convey complicated ideas. The use of this noun-explanatory feature can assist with campaign success by accurately explaining and defining key nouns in their descriptions, such as technical terms that are likely to be unknown to backers. The *Dep_punct* feature, in the context of dependency parsing, is a type of syntactic dependency relationship that links punctuation marks to the words they are associated with in a sentence. The use of punctuation marks like "!" or "," seems to unveil excitement as well as clarity, a finding we link to Kaminski and Hopp's [20] results of enthusiastic language as a campaign success predictor.

4.2. Counterfactual Explanations of Model 1

To further enhance the explanation of the trained models and provide specific, actionable insights to project owners and backers, an additional XAI technique was utilized, namely Counterfactual Explanations. The XGBoost model was used with the DiCE Counterfactual Explanations library to infer actionable changes to unsuccessful project features that could convert them to successful projects. Most XAI methods, including SHAP, identify features associated with a target variable, in our case, success (or failure), but cannot provide information on how to change the outcome of the model.

We employed DiCE for the text-based features model for a specific project that was not successful. Our objective here was to examine RQ3 and demonstrate the approach's capabilities at the micro-level, prompting future research to employ and expand its use in this context. Table 3 summarizes the DiCE results for a specific campaign that failed. Recommendations indicate how the outcome of this failed project can change by altering two main textual features in its description, as seen in the output of model '1'. It specifically indicates that if the campaigner improved the language used in the description by using around 16 punctuation marks, then the project would have been successful (see dep_punct). More importantly, the analysis mentions a specific dale_chall_score (7.96) that needs to be satisfied to convert the unsuccessful project to a successful one. The latter is a metric of language complexity, meaning that the respective descriptions were either too difficult to understand by the majority of viewers or overly simplistic, which triggered concerns about the campaigner's knowledge of the subject matter. Achieving a Dale–Chall score of 7.96 can put the project in a better position to succeed. We note that the recommended changes that emerge from the Counterfactual Explanations constitute minimum actionable changes and are case-specific. The objective is to demonstrate the applicability of the method, and due to the case-specific nature of recommendations, results here should not be misinterpreted as generalizable. In the case that explanations are required for a group of cases, alternative approaches can be used (for a review, see [16]).

Table 3. DiCE results recommending minimum actionable changes to a specific failed project to make it successful.

Failed Crowdfunding Project	Features		Output of Model
Number	dale_chall_score	dep_punct	Success
384	7.96	16	1

4.3. Model 2—All Features

The second XGBoost model aimed to scrutinize all features (numerical, categorical, and textual). The model achieved an AUC of 92% and F1 of 91%. The most influential features are the type of project, number of comments (posts), project duration, and use of images. Figure 4 shows the SHAP plot for the second XGBoost model.

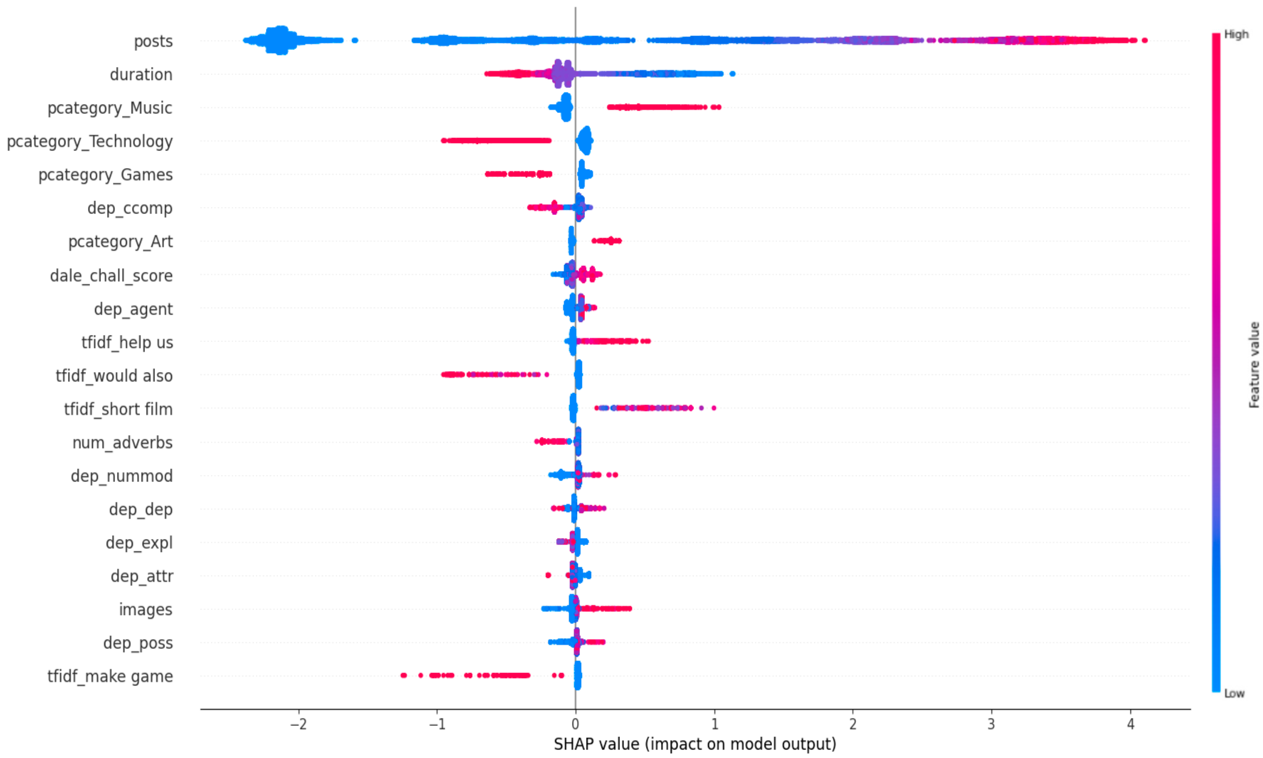

Figure 4. SHAP explanations of the second XGBoost model trained on textual, numeric, and categorical features.

We observe a series of findings that complement the previous XGBoost model, providing confirmatory links with the previous literature. In the second column, the *duration* of the campaign is found to be significant and associated with project failure. This confirms previous research identifying that longer than average periods are associated with campaign failure [19]. The utilization of images to convey the main idea and demonstrate the product in action was confirmed to be a predictor of success [8,20]. Finally, engagement of the campaigner through *posts* is associated with success.

Additional findings highlight certain campaign types that are doing better than others. Interestingly, projects in the *Art* and *Music* categories are associated with success, while projects in the categories of *Gaming* and *Technology* are more closely associated with failure. We note that these refer to trends present in the timeframe reflected in the current dataset and not necessarily general conclusions of campaign categories that are more likely to be successful nor reflective of future trends.

4.4. Validation Studies

To validate the results produced using SHAP and the two XGBoost models, additional XAI techniques were used, such as the XGBoost's feature importance function, surrogate XGBoost, tree-based explanation, and partial dependence plots. Results from these simpler techniques show that the features that have been identified as important by SHAP were also found important by these techniques. Indicatively, Appendix B depicts features using the inherent feature importance function of XGBoost, verifying most of the features in SHAP. Similar results are obtained from the other XAI techniques. However, SHAP ensures a fair and consistent allocation of feature contributions based on Shapley values derived from cooperative game theory and can show feature interactions and how each feature affects the output. In contrast to SHAP, the other techniques cannot account for feature interactions (except from partial dependence plots that show interaction between two features), and neither can show whether a feature influences the output positively or negatively.

5. Discussions

The rationale of this paper was to enhance decisions made within the crowdfunding ecosystem by both backers and campaigners in order to enhance its financial sustainability. Specifically, we contribute toward more accurate and efficient allocation of funds to projects that are less risky and prone to failure. In this paper, we demonstrate how eXplainable AI techniques can be used to analyze the textual features found in crowdfunding campaigns and assess how these translate to success or failure. We complemented the first XGBoost model with Counterfactual Explanations, aiming to present a recommendation-generating technique that can assist crowdfunding decision-makers and enhance managerial implications. Lastly, we followed this main analysis with a second XGBoost model that included all types of features (numerical, categorical, and textual) in order to expand on the findings of the first model.

Our results were based on a Kickstarter dataset that focused on entrepreneurial projects. We acknowledge that the inclusion of datasets from other platforms like IndieGoGo would enhance result robustness, yet we argue that risk communication and sentence complexity are features deeply embedded in human communication. As such, we would expect our results to be replicable with datasets from other entrepreneurial-focused crowdfunding platforms. Additionally, we acknowledge that this paper defined success as the ability to receive the pledged amount. As such, this paper's focus did not consider the deliverance of a final product but instead the reception of the pledged amount. We note that Kickstarter acknowledges that only a small percentage of 9% of projects fail to deliver a final product [4], yet future research can work with the latter variable for a better mapping of successful crowdfunds.

Implications and Conclusions

For RQ1, we found strong links between terms like "*stretch goal(s)*" and project success. We identify that the use of such terms that convey a conjunction of Ambitiousness AND Risk results in positive effects compared to language that communicates directly and transparently the involved risks [30]. This demonstrates an asymmetry and a potential weak point in the decision-making process of backers that potentially requires shielding through respective crowdfunding policies. We argue that transparency of involved risks in campaign descriptions can bring forth worries that would have otherwise remained dormant by backers ultimately preferring to support projects that have similar risks but are less risk-transparent. Findings point out the susceptibility of backers toward certain risk communication techniques. Similar findings that question decision rationality were reported in the context of disclosures where logical assurances of privacy reminded individuals of possible threats instead of reassuring them of the safety of their privacy, thus preferring to disclose information to sites that were less secure and did not provide such assurances that elevate worries [49]. Backers are susceptible to the aforementioned risk communication techniques, an argument we invite future research to investigate further. In handling this vulnerability, the authors propose the utilization of informative (system 2) precision nudges [23,50] that prompt deliberation by transparently reporting the exact risk associated with a project using historical data from similar projects. A success/failure percentage can be displayed in the platforms' interface when an individual is in deliberation or ready to invest. Adoption of such nudges aims to inform backers prior to a decision, and their utilization across platforms can generate a new baseline for risk communication in crowdfunds that is both more sustainable and ethical.

Through RQ2, we fill in the gap in language complexity and sentence structure, expanding previous research (e.g., [9]). Specifically, we identify shorter sentences that employ more complex and technical words to be associated with success. Language that is complex yet comprehensible and includes technical elements within short sentences is perceived to be associated with the author's knowledge of the domain. We link the latter to signaling theory [18] and how campaigners communicate aspects of both their character and the investment proposition to influence decision-makers in deliberation. We argue that the projected knowledgeability of the campaigner through the sentence structure and language complexity in their campaign's description is a signal of the qualities of the campaigner. Implications from these findings include the adoption of XAI technology in crowdfunding platforms that would automatically guide campaigners on the complexity of their campaign's description. AI tools can make recommendations on how to tone down (or increase) language complexity after analyzing and identifying the text as either ultra-complicated or too simplistic, respectively, for the public. This can allow for a more balanced communication method founded on clarity, understandability, and transparency toward prospective backers. Furthermore, long sentences that challenge the public's comprehension can be identified, and the campaigners' attention can be requested to restructure them, thus enhancing communication quality.

The proposition of employing Counterfactual Explanation in the crowdfunding setting was explored via RQ3 and aimed to offer actionable recommendations to campaigners of failed projects. We expand previous applications (e.g., [16]) by demonstrating the applicability of Counterfactual Explanation in the context of crowdfunds. Counterfactual Explanation, DiCE results specifically, delineate questions on the minimum actionable changes that need to be performed in a project's description to enhance the chances of conversion from an unsuccessful to a successful one. For example, a certain number of punctuation marks and a specific Dale–Chall score for word complexity can assist a project's success. Albeit we acknowledge such changes to be incremental for the project's success (the actual product idea is naturally a more significant determinant of success), such language propositions on how to communicate an idea can assist projects that were close to meeting the pledged amount but fell short.

For RQ4, we provide a confirmatory analysis of previous research. We verify previous research with results indicating that the inclusion of images is linked with success [8,20]. We verify that the long duration of campaigns is associated with failure [8,16]. Finally, updates in the form of posts that demonstrate a consistent communication pattern by the campaigner toward its backers were associated with success, confirming the results of Yeh et al. [9] and Xiao et al. [19]. We complement these findings through the examination of crowdfunding categories. Results show that the *Art* and *Music* categories are associated with success, while projects in the categories *Gaming* and *Technology* are linked to failure. We note that these findings should not be taken at face value as they present trends during the covered timeframe of the dataset. If prospective campaigners have datasets reflective of the most recent patterns, then decisions could be based on such findings with more confidence, further channeling research on improving the economic sustainability and viability of the crowdfunding ecosystem.

Author Contributions: Conceptualization, C.T. and A.G.; writing—original draft preparation, C.T. and A.G.; writing—review and editing, C.T. and A.G.; methodology, C.T. and A.G.; supervision, A.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets generated and analyzed during the current study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Readability Formulas

The Dale–Chall formula used for language complexity score:

$$\text{Score} = 0.1579 \times \left(\frac{Total\ words}{Difficult\ words} \times 100\right) + 0.0496 \times \left(\frac{Total\ sentence}{Total\ words}\right)$$

The Flesch–Kincaid formula is used to measure sentence readability and complexity. The higher the grade level, the more complex the sentence is.

$$\text{Grade Level} = 0.39 \times \left(\frac{Total\ sentences}{Total\ words}\right) + 11.8 \times \left(\frac{Total\ words}{Total\ syllables}\right) - 15.59$$

Appendix B. XGBoost's Feature Importance Function

Table summarizing the inherent feature importance function of the XGBoost model.

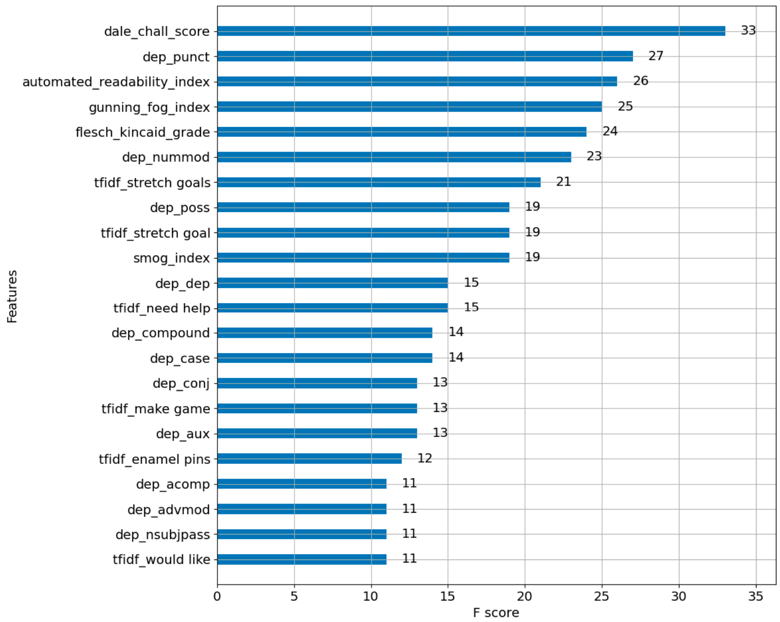

References

1. Polaris Market Research. *Crowdfunding Market Size to Grow USD 72.88 Billion by 2032, at a CAGR of 16.1% from 2024 to 2032*; Polaris Market Research & Consulting LLP.: Dover, DE, USA, 2024. Available online: https://www.globenewswire.com/news-release/2024/08/12/2928427/0/en/Crowdfunding-Market-size-to-grow-USD-72-88-Billion-by-2032-at-a-CAGR-of-16-1-from-2024-to-2032-Polaris-Market-Research.html#:~:text=Key (accessed on 18 November 2024).
2. Cumming, D.; Hornuf, L.; Karami, M.; Schweizer, D. Disentangling Crowdfunding from Fraudfunding. *J. Bus. Ethics* **2023**, *182*, 1103–1128. [CrossRef] [PubMed]
3. Mollick, E. The Dynamics of Crowdfunding: An Exploratory Study. *J. Bus. Ventur.* **2014**, *29*, 1–16. [CrossRef]
4. Kickstarter. *The Kickstarter Fulfilment Report*; Kickstarter: New York, NY, USA, 2024.
5. Torres, B.; Serrasqueiro, Z.; Oliveira, M. Crowdfunding in Portugal—Using the Raize Platform to Ensure the Sustainability of Companies and Projects. *Sustainability* **2024**, *16*, 4439. [CrossRef]

6. Zhao, K.; Wang, H.; Wang, W. Women Entrepreneurs Who Break through in Reward-Based Crowdfunding: The Influence of Entrepreneurial Orientation. *Sustainability* **2023**, *15*, 9276. [CrossRef]
7. Talukder, S.C.; Lakner, Z. Exploring the Landscape of Social Entrepreneurship and Crowdfunding: A Bibliometric Analysis. *Sustainability* **2023**, *15*, 9411. [CrossRef]
8. Liu, Z.; Ben, S.; Zhang, R. Factors Affecting Crowdfunding Success. *J. Comput. Inf. Syst.* **2023**, *63*, 241–256. [CrossRef]
9. Yeh, T.L.; Chen, T.Y.; Lee, C.C. Investigating the Funding Success Factors Affecting Reward-Based Crowdfunding Projects. *Innovation* **2019**, *21*, 466–486. [CrossRef]
10. Colombo, M.G.; Franzoni, C.; Rossi-Lamastra, C. Internal Social Capital and the Attraction of Early Contributions in Crowdfunding. *Entrep. Theory Pract.* **2015**, *39*, 75–100. [CrossRef]
11. Barbi, M.; Bigelli, M. Crowdfunding Practices in and Outside the US. *Res. Int. Bus. Financ.* **2017**, *42*, 208–223. [CrossRef]
12. Chaudhary, S. Enhancing Success Prediction In Crowdfunding Platform Using LIME-Explainable AI Method. In Proceedings of the 2024 IEEE Students Conference on Engineering and Systems (SCES), Prayagraj, India, 21–23 June 2024; pp. 1–6. [CrossRef]
13. Pekar, V.; Candi, M.; Beltagui, A.; Stylos, N.; Liu, W. Explainable Text-Based Features in Predictive Models of Crowdfunding Campaigns. *Ann. Oper. Res.* **2024**, 1–31. [CrossRef]
14. Owen, A.B.; Prieur, C. On Shapley Value for Measuring Importance of Dependent Inputs. *SIAM ASA J. Uncertain. Quantif.* **2017**, *5*, 986–1002. [CrossRef]
15. Wachter, S.; Mittelstadt, B.; Russell, C. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *Harv. J. Law Technol.* **2018**, *31*, 841–887. [CrossRef]
16. Artelt, A.; Gregoriades, A. Supporting Organizational Decisions on How to Improve Customer Repurchase Using Multi-Instance Counterfactual Explanations. *Decis. Support Syst.* **2024**, *182*, 114249. [CrossRef]
17. Kouros, T.; Theodosiou, Z.; Themistocleous, C. Machine Learning Bias: Geneaology, Expression and Prevention. In *AI in Social Research*; Christou, P., Ed.; CABI Publishing: Wallingford, UK, 2025.
18. Koch, J.A.; Siering, M. The Recipe of Successful Crowdfunding Campaigns: An Analysis of Crowdfunding Success Factors and Their Interrelations. *Electron. Mark.* **2019**, *29*, 661–679. [CrossRef]
19. Xiao, S.; Tan, X.; Dong, M.; Qi, J. How to Design Your Project in the Online Crowdfunding Market? Evidence from Kickstarter. In Proceedings of the Thirty Fifth International Conference on Information Systems (ICIS), Auckland, Australia, 14–17 December 2014.
20. Kaminski, J.C.; Hopp, C. Predicting Outcomes in Crowdfunding Campaigns with Textual, Visual, and Linguistic Signals. *Small Bus. Econ.* **2020**, *55*, 627–649. [CrossRef]
21. Zhang, Y.; DeCarlo, T.E.; Manikas, A.S.; Bhattacharya, A. To Exploit or Explore? The Impact of Crowdfunding Project Descriptions and Backers' Power States on Funding Decisions. *J. Acad. Mark. Sci.* **2023**, *51*, 444–462. [CrossRef]
22. Mehta, V.; Mehra, R.; Verma, S.S. A Survey on Customer Segmentation Using Machine Learning Algorithms to Find Prospective Clients. In Proceedings of the 2021 9th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO), Noida, India, 3–4 September 2021. [CrossRef]
23. Themistocleous, C.; Karapanos, E. Precision Nudging: The Future of Behavioural Change. In *Artificial Intelligence (AI) in Social Research*; Christou, P., Ed.; CABI: Wallingford, UK, 2024.
24. Mahbub, N.; Le, A.; Zhuang, J. Online Crowd-Funding Strategy: A Game-Theoretical Approach to a Kickstarter Case Study. *Ann. Oper. Res.* **2022**, *315*, 1019–1036. [CrossRef]
25. Agrawal, A.; Gans, J.; Goldfarb, A. How to Win with Machine Learning. *Harv. Bus. Rev.* **2020**, *98*, 126–133.
26. Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; Pedreschi, D. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* **2018**, *51*, 1–42. [CrossRef]
27. Ribeiro, M.T.; Singh, S.; Guestrin, C. "Why Should I Trust You?" Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 1135–1144. [CrossRef]
28. Kravet, T.; Muslu, V. Textual Risk Disclosures and Investors' Risk Perceptions. *Rev. Account. Stud.* **2013**, *18*, 1088–1122. [CrossRef]
29. Bolton, L.E.; Cohen, J.B.; Bloom, P.N. Does Marketing Products as Remedies Create "Get Out of Jail Free Cards"? *J. Consum. Res.* **2006**, *33*, 71–81. [CrossRef]
30. Kim, K.; Park, J.; Pan, Y.; Zhang, K.; Zhang, X. Risk Disclosure in Crowdfunding. *Inf. Syst. Res.* **2022**, *33*, 1023–1041. [CrossRef]
31. Welkenhuysen, M.; Evers-Kiebooms, G.; D'Ydewalle, G. The Language of Uncertainty in Genetic Risk Communication: Framing and Verbal versus Numerical Information. *Patient Educ. Couns.* **2001**, *43*, 179–187. [CrossRef] [PubMed]
32. Moss, T.W.; Payne, G.T.; Moore, C.B. Strategic Consistency of Exploration and Exploitation in Family Businesses. *Fam. Bus. Rev.* **2014**, *27*, 51–71. [CrossRef]
33. Chotirat, S.; Meesad, P. Part-of-Speech Tagging Enhancement to Natural Language Processing for Thai Wh-Question Classification with Deep Learning. *Heliyon* **2021**, *7*, e08216. [CrossRef]

34. Shams, R.; Mercer, R.E. Classifying Spam Emails Using Text and Readability Features. In Proceedings of the 2013 IEEE 13th International Conference on Data Mining, Dallas, TX, USA, 7–10 December 2013; pp. 657–666. [CrossRef]
35. Ige, O.P.; Gan, K.H. Ensemble Filter-Wrapper Text Feature Selection Methods for Text Classification. *Comput. Model. Eng. Sci.* **2024**, *141*, 1847–1865. [CrossRef]
36. Park, J.; Na, H.J.; Kim, H. Development of a Success Prediction Model for Crowdfunding Based on Machine Learning Reflecting ESG Information. *IEEE Access* **2024**, *12*, 197275–197289. [CrossRef]
37. Nielsen, D. Tree Boosting with Xgboost-Why Does Xgboost Win "Every" Machine Learning Competition? Master's Thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2016.
38. Jeni, L.A.; Cohn, J.F.; De La Torre, F. Facing Imbalanced Data—Recommendations for the Use of Performance Metrics. In Proceedings of the 2013 Humaine Association Conference on Affective Computing and Intelligent Interaction, Geneva, Switzerland, 2–5 September 2013; pp. 245–251. [CrossRef]
39. Vij, A.; Nanjundan, P. Comparing Strategies for Post-Hoc Explanations in Machine Learning Models. *Lect. Notes Data Eng. Commun. Technol.* **2022**, *68*, 585–592. [CrossRef]
40. Retzlaff, C.O.; Angerschmid, A.; Saranti, A.; Schneeberger, D.; Röttger, R.; Müller, H.; Holzinger, A. Post-Hoc vs Ante-Hoc Explanations: XAI Design Guidelines for Data Scientists. *Cogn. Syst. Res.* **2024**, *86*, 101243. [CrossRef]
41. Glanois, C.; Weng, P.; Zimmer, M.; Li, D.; Yang, T.; Hao, J.; Liu, W. *A Survey on Interpretable Reinforcement Learning*; Springer: Berlin/Heidelberg, Germany, 2024; Volume 113, ISBN 0123456789.
42. Weerts, H.J.P.; van Ipenburg, W.; Pechenizkiy, M. A Human-Grounded Evaluation of SHAP for Alert Processing. *arXiv* **2019**, arXiv:1907.03324.
43. Liu, Y.; Liu, Z.; Luo, X.; Zhao, H. Diagnosis of Parkinson's Disease Based on SHAP Value Feature Selection. *Biocybern. Biomed. Eng.* **2022**, *42*, 856–869. [CrossRef]
44. Del Ser, J.; Barredo-Arrieta, A.; Díaz-Rodríguez, N.; Herrera, F.; Saranti, A.; Holzinger, A. On Generating Trustworthy Counterfactual Explanations. *Inf. Sci.* **2024**, *655*, 119898. [CrossRef]
45. Darias, J.M.; Díaz-Agudo, B.; Recio-Garcia, J. A Systematic Review on Model-Agnostic XAI Libraries. In Proceedings of the 29th International Conference, ICCBR 2021, Salamanca, Spain, 13–16 September 2021; pp. 28–39.
46. Gary, M.S.; Yang, M.M.; Yetton, P.W.; Sterman, J.D. Stretch Goals and the Distribution of Organizational Performance. *Organ. Sci.* **2017**, *28*, 395–410. [CrossRef]
47. Sitkin, S.B.; See, K.E.; Miller, C.C.; Lawless, M.W.; Carton, A.M. The Paradox of Stretch Goals: Organizations in Pursuit of the Seemingly Impossible. *Acad. Manag. Rev.* **2011**, *36*, 544–566. [CrossRef]
48. Thompson, K.R.; Hochwarter, W.A.; Mathys, N.J. Stretch Targets: What Makes Them Effective? *Acad. Manag. Exec.* **1997**, *11*, 48–60. [CrossRef]
49. Acquisti, A.; Brandimarte, L.; Loewenstein, G. Privacy and Human Behavior in the Age of Information. *Science* **2015**, *347*, 509–514. [CrossRef]
50. Sunstein, C.R. People Prefer System 2 Nudges (Kind of). *Duke Law J.* **2016**, *66*, 121. [CrossRef]