

A Proposed Framework for Venture Capital Investment Analysis Using Cox Proportional Hazards Model and SHAP-Based Explainability

Nikhil Rajendrakumar Naik¹, Dr. Aisha Fernandes²

^{1,2}Department of Information Technology, Goa College of Engineering, Ponda, Goa, India

Abstract - The venture capital industry operates under conditions of persistent uncertainty, where investment decisions must be made on the basis of incomplete information about which early-stage companies will ultimately succeed, be acquired, or fail. Quantitative tools for investment analysis have improved significantly with the growth of machine learning, yet two critical limitations remain in most existing approaches: the absence of temporal modeling that accounts for the time-dependent nature of startup risk, and the opacity of predictive models that prevents their adoption in professional investment workflows requiring transparent, auditable reasoning. This paper proposes a conceptual framework that addresses both limitations by integrating two established analytical techniques: the Cox Proportional Hazards model, which estimates time-dependent survival probability for portfolio companies, and SHAP (SHapley Additive exPlanations), which explains the predictions of classification models through principled, game-theoretic feature attribution. The paper reviews the theoretical basis of each technique, motivates their integration as complementary rather than overlapping tools, and discusses how the combined framework provides a more complete analytical summary of investment risk than either method offers individually. Limitations and directions for future work are also identified.

Key Words: Venture Capital, Cox Proportional Hazards Model, SHAP, Explainable Artificial Intelligence, Startup Investment, Survival Analysis, Risk Assessment, Investment Decision Support

1.INTRODUCTION

Venture capital investment is among the most uncertain of financial activities. The decision to allocate capital to an early-stage company requires judgment about factors that are inherently difficult to quantify: the quality and resilience of founding teams, the size and timing of market opportunities, and the competitive dynamics that will shape a company's trajectory over the coming years.

Historically, these judgments have been made through qualitative assessment by experienced investors, drawing on pattern recognition built over careers of deal evaluation.

The increasing availability of structured company data has opened new possibilities for quantitative investment analysis. Databases cataloguing funding histories, operational milestones, and company outcomes across tens of thousands of startups provide the raw material for machine learning models capable of identifying statistical patterns in startup success and failure. Research in this domain has demonstrated that predictive models trained on such data can generate acquisition likelihood scores and failure risk estimates that carry meaningful predictive signal [1].

Despite this progress, most existing quantitative approaches to startup investment share two structural limitations. The first is that they treat outcome prediction as a static problem, assigning a probability of success or acquisition at a single point in time without modeling how that risk changes as a company ages, raises new funding, or hits product milestones. Startup failure is not a static risk — it evolves continuously, and a framework that captures only a snapshot misses the dynamic nature of portfolio risk management.

The second limitation is interpretability. The machine learning models that achieve the strongest predictive performance on startup data — gradient boosting ensembles and random forests — are also the most opaque. Investment professionals cannot interrogate the internal logic of these models, cannot explain their outputs to investment committees or limited partners, and cannot validate whether the model's predictions are grounded in economically meaningful signals or spurious statistical associations in historical data.

This paper proposes that the Cox Proportional Hazards model and SHAP (SHapley Additive exPlanations) together constitute a principled response to both limitations. The Cox model provides the temporal dimension that static classifiers lack. SHAP provides the interpretability layer that opaque ensemble models lack. Together, they form a conceptual framework for

investment analysis that is richer, more transparent, and more actionable than either approach in isolation.

2. BACKGROUND

2.1 Quantitative Investment Analysis

The application of statistical and machine learning methods to startup investment has a substantial academic history. Early empirical work established relationships between funding characteristics and investment outcomes using regression frameworks. More recent research has demonstrated that ensemble classifiers achieve statistically meaningful discrimination between acquired and non-acquired startups, with area under the receiver operating characteristic curve values indicating genuine predictive signal beyond naive baseline rates [1].

However, the dominant paradigm in this literature is binary outcome classification: models are trained to predict whether a company will be acquired or will succeed, without modeling the time dimension of that outcome. This limits the practical utility of such models for portfolio management, where the timing of outcomes — not just their probability — is central to return calculation, reserve allocation, and exit strategy planning.

2.2 Survival Analysis in Business Contexts

Survival analysis is a well-established statistical methodology for modeling the time until an event of interest occurs, accounting for the presence of censored observations — cases where the event has not yet occurred at the time of analysis. Originally developed for clinical and engineering applications, survival analysis has been applied productively in business research to model customer churn, employee turnover, and company failure [2].

The application of survival analysis to startup failure modeling is a natural one. Many companies in a portfolio are still operating at any given analysis date — they are right-censored with respect to the failure event. Standard classification approaches discard this information by treating censored observations as successes, introducing systematic bias. Survival analysis handles censoring correctly, making use of the full observed history of each company regardless of whether the failure event has occurred.

2.3 Explainability in Machine Learning

The interpretability of machine learning models has become an increasingly important research and policy concern. The tension between predictive performance and model transparency — often described as the accuracy-interpretability trade-off — is particularly acute in high-stakes domains such as finance, medicine, and law, where opaque model outputs cannot be responsibly acted upon without understanding the reasoning behind them [3].

Several methodological approaches to model explanation have been developed, ranging from inherently interpretable models such as linear regression and shallow decision trees, to post-hoc explanation methods that approximate the behavior of complex models locally or globally. Among post-hoc methods, SHAP has emerged as theoretically preferred because it satisfies formal axiomatic properties — local accuracy, consistency, missingness, and symmetry — that heuristic alternatives do not guarantee [4].

3. THE COX PROPORTIONAL HAZARDS MODEL

3.1 Overview

The Cox Proportional Hazards model, proposed by David Cox in 1972, is a semi-parametric regression method for survival data [5]. It models the hazard — the instantaneous rate of occurrence of an event — as a function of a set of covariates, without requiring a specific parametric form for how the baseline hazard varies over time. This semi-parametric flexibility is one of the principal reasons for the model's widespread adoption: it allows the data to determine the shape of the underlying risk trajectory rather than imposing an assumption that may be incorrect.

The model specifies that the hazard for a given subject at time t is the product of a baseline hazard function, common to all subjects, and an exponential function of the subject's covariate values. The exponential form ensures that the hazard is always positive and allows the model to be estimated through partial likelihood methods that do not require explicit estimation of the baseline hazard. Covariate effects are summarized as hazard ratios, which express the multiplicative change in hazard associated with a one-unit increase in each covariate holding others constant.

3.2 Key Concepts

Three concepts are central to understanding the Cox model in the investment context. The survival function expresses the probability that a company will still be operating at a given time, and is a monotonically decreasing function from one at the time of founding toward zero as time increases. The hazard function expresses the instantaneous rate of failure at each time point given survival up to that point, and captures the dynamic risk profile of a company over its lifecycle. The concordance index, also called the C-statistic, measures the discriminative ability of the fitted model by calculating the proportion of pairs of companies for which the model correctly orders survival times, with a value of 0.5 representing random chance and 1.0 representing perfect discrimination.

The proportional hazards assumption — that the ratio of hazards between any two companies is constant over time — is the key modeling assumption of the Cox framework. This assumption implies that the covariates shift the hazard function up or down by a constant multiplicative factor at all time points, rather than modifying its shape. The assumption can be tested empirically using Schoenfeld residuals, and can be relaxed through stratification or time-varying coefficient extensions when violations are detected.

3.3 Relevance to Venture Capital

For venture capital applications, the Cox model provides answers to investment questions that static classification cannot address. It estimates the probability that a portfolio company will still be operating at a specific future time horizon — one year, three years, five years — given its current characteristics. It identifies which company attributes are most strongly associated with longer survival, expressed as hazard ratios that quantify the protective or detrimental effect of each characteristic. And it supports scenario-based projection, allowing analysts to estimate how a change in a company's funding or operational profile would affect its expected survival trajectory.

These capabilities make the Cox model a natural complement to acquisition likelihood classification. Where a classifier answers whether a company is likely to be acquired, the Cox model answers how long the company is likely to survive if it is not — providing a second analytical dimension that is essential for portfolio management, exit timing, and reserve allocation decisions.

4. SHAP: SHAPLEY ADDITIVE EXPLANATIONS

4.1 Theoretical Foundation

SHAP draws its theoretical foundations from cooperative game theory, specifically from the Shapley value concept introduced by Lloyd Shapley in 1953 [6]. In the game-theoretic formulation, a set of players cooperate to generate a joint payoff, and the Shapley value provides the unique fair allocation of that payoff among players satisfying a set of axiomatic fairness properties. Translated to machine learning, the players are the input features of a model, the payoff is the model's prediction for a specific instance, and the Shapley value for each feature is its fair contribution to that prediction.

The Shapley value for a feature is computed as the weighted average of that feature's marginal contribution to the model output across all possible orderings of features. This formulation ensures that the contribution attributed to each feature accounts for the full range of contexts in which that feature might be encountered — alone, in combination with other features, or as the last feature added to a complete feature set. The result is an attribution that is both mathematically principled and intuitively meaningful: it represents how much a specific feature shifted the model's output, on average, relative to what the model would have predicted without that feature.

4.2 SHAP Axioms

SHAP satisfies four axioms that together constitute a uniqueness result: it is the only attribution method satisfying all four simultaneously. Local accuracy requires that the sum of all feature attributions equals the difference between the model's prediction for a specific instance and its average prediction across the training data, ensuring that attributions fully account for the prediction. Missingness requires that features absent from a specific input receive zero attribution, preventing phantom contributions from uninformative features. Consistency requires that if a model is modified such that a feature's contribution increases for all possible feature combinations, that feature's attribution should not decrease, ensuring stability across model updates. Symmetry requires that features contributing equally to all feature coalitions receive equal attribution, ensuring fair treatment of structurally equivalent features.

These properties distinguish SHAP from heuristic alternatives. Methods such as permutation importance

and local linear approximation satisfy some of these properties in some conditions, but none satisfies all four simultaneously. This theoretical advantage is particularly important in investment applications, where the credibility of an explanation depends on its mathematical guarantees rather than its intuitive plausibility alone.

4.3 Types of SHAP Explanations

SHAP supports several types of explanations that serve different purposes in an investment analytics workflow. Global feature importance summarizes each feature's average contribution across all predictions, providing a portfolio-level view of which company characteristics most influence acquisition likelihood. Directional effect analysis reveals whether high or low values of each feature increase or decrease predicted acquisition probability, enabling validation of model behavior against economic intuition. Individual prediction explanation attributes the acquisition score for a specific company to its constituent features, providing the company-specific reasoning required for investment committee review. Interaction analysis identifies pairs of features whose joint contribution to predictions exceeds the sum of their individual contributions, revealing compound signals that univariate analysis would miss. Together, these explanation types constitute a multi-level view of model behavior — from portfolio-wide patterns to company-specific justifications — that supports investment analysis at every stage of the deal lifecycle, from initial screening through due diligence to portfolio monitoring.

5. PROPOSED INTEGRATION OF COX MODEL AND SHAP

5.1 Motivation for Integration

The Cox model and SHAP are analytically complementary in a precise sense: they address different dimensions of the same investment question through different methodological approaches. The Cox model operates on time-to-event data and produces probabilistic estimates of company longevity — a continuous, temporal output. SHAP operates on the outputs of a classification model and produces feature-level attributions for discrete outcome predictions — a categorical, explanatory output. Neither technique alone provides a complete analytical picture of investment risk, but together they cover the two dimensions most

critical for professional investment decision-making: how long a company is expected to survive, and why it is predicted to succeed or fail.

This complementarity motivates an integrated framework in which the Cox model and a SHAP-explained classifier are applied in parallel to the same company data, with their outputs synthesized into a unified risk summary for each portfolio company or investment candidate. The survival probability estimates from the Cox model and the explained acquisition score from the classifier together provide a two-dimensional characterization of investment risk that supports a broader range of decision types than either output supports individually.

5.2 Framework Description

The proposed framework operates through a structured analytical pipeline. In the initial stage, company data is prepared into a consistent feature representation capturing financial characteristics such as total funding raised and the number of funding rounds completed, operational characteristics such as the number of product milestones achieved and business relationships formed, and temporal characteristics such as company age and the time elapsed between consecutive funding events. This feature set serves as the input to both the survival model and the classification model, ensuring that the two analytical streams operate on a common information base.

In the modeling stage, the Cox Proportional Hazards model is fitted on a survival dataset derived from the company records, using company closure as the event of interest and treating actively operating and acquired companies as right-censored observations. Simultaneously, a gradient-boosted classifier is trained on the same feature set to predict acquisition probability as a binary outcome. The two models are estimated independently, reflecting their distinct methodological requirements — the Cox model uses partial likelihood estimation on time-to-event data, while the classifier uses cross-entropy loss minimization on binary outcome labels.

In the explanation stage, SHAP TreeExplainer is applied to the trained classifier to generate the full suite of feature attributions for each company in the analysis set. The resulting SHAP values quantify how much each feature characteristic shifted the acquisition probability prediction above or below the population average, providing the feature-level reasoning that transforms the

classifier from an opaque scoring engine into an auditable decision support tool.

In the synthesis stage, the outputs of the three analytical components — survival probability curve from the Cox model, acquisition likelihood score from the classifier, and SHAP attributions from the explainer — are combined into a unified investment risk summary for each company. This summary answers three questions simultaneously: how likely is the company to be acquired, why does the model predict that outcome based on the company's specific characteristics, and how long is the company expected to survive if an acquisition does not occur in the near term.

5.3 Benefits for Investment Practice

The integrated framework offers several practical benefits for investment practice that neither the Cox model nor SHAP provides in isolation. For portfolio risk management, the combination of survival trajectory and explained acquisition score enables systematic segmentation of portfolio companies into categories that support differentiated management strategies — companies with high acquisition scores and strong survival trajectories represent stable long-term holds, while companies with deteriorating survival trajectories require active intervention regardless of their acquisition attractiveness.

For deal screening, the framework enables structured prioritization of large investment pipelines. The Cox model provides an initial filter based on expected longevity, while the SHAP-explained classifier provides a second filter based on acquisition potential with company-specific justification. The combination reduces the cognitive burden on investment teams by directing analytical attention toward the specific characteristics of each candidate company most relevant to the investment decision.

For investment committee communication, the SHAP force plot for each screened company provides a structured, visual narrative of the quantitative reasoning behind an investment recommendation — which specific company characteristics are most aligned with the historical profile of successful acquisition targets, and which represent areas of concern. This narrative is more accessible to investment professionals without quantitative backgrounds than raw model scores, and more defensible than heuristic assessments without quantitative grounding.

6. ADVANTAGES OF THE PROPOSED APPROACH

6.1 Temporal Completeness

The most significant analytical advantage of incorporating the Cox model into an investment framework is temporal completeness. Static classification models assign an acquisition probability that is invariant to time, implicitly treating the risk of failure as constant over a company's lifecycle. This is empirically incorrect: startup failure rates are highest in the first two years, decline as companies achieve product-market fit and stable revenue, and may increase again as companies age beyond their initial growth phase without achieving scale. The Cox model captures this temporal variation through the hazard function, providing a dynamic risk estimate that evolves with company age and changes in covariate values.

Temporal completeness is not merely a statistical refinement — it has direct practical implications for investment decision-making. A company with high predicted acquisition probability but elevated early-stage failure risk requires a different portfolio management strategy than a company with the same acquisition score and a more stable survival trajectory. The integrated framework makes this distinction explicit and quantitative, enabling more precise risk management across diverse portfolio compositions.

6.2 Principled Interpretability

The integration of SHAP addresses the interpretability limitation of high-performing ensemble classifiers in a theoretically principled way. Unlike ad-hoc explanation methods that approximate model behavior through local linearization or feature permutation, SHAP provides exact attributions derived from the model's actual decision logic through the Shapley value computation. This exactness means that SHAP explanations accurately represent the model's reasoning rather than providing a simplified approximation that may diverge from the true decision boundary in complex feature regions.

Principled interpretability is particularly important in venture capital because investment decisions made on the basis of model outputs must be defensible to multiple audiences with different levels of quantitative sophistication. The additive structure of SHAP attributions — where each feature's contribution can be read directly from the force plot and the contributions

sum to the total prediction — supports clear communication of model reasoning to investment committee members, limited partners conducting due diligence, and regulatory reviewers assessing algorithmic decision systems.

6.3 Auditability and Governance

A practical requirement in institutional investment management is the maintenance of auditable records documenting the reasoning behind investment decisions. Traditional qualitative investment memos fulfill this requirement for human-driven decisions; the proposed framework extends this governance standard to data-driven decisions by generating a structured, quantitative record of the model reasoning behind each investment recommendation. The SHAP attribution vector for each company prediction, stored alongside the acquisition score and survival probability estimate, constitutes a machine-readable audit trail that supports retrospective analysis of model behavior, regulatory examination, and post-hoc evaluation of prediction quality.

This auditability property aligns the proposed framework with emerging governance standards for artificial intelligence in financial services, including the requirements of the European Union's AI Act for high-risk AI systems and the documentation standards of the Financial Stability Board's guidance on artificial intelligence and machine learning in financial services.

7. LIMITATIONS AND FUTURE DIRECTIONS

7.1 Modeling Assumptions

The Cox Proportional Hazards model rests on the proportional hazards assumption, which requires that the ratio of hazards between any two companies remains constant over time. In startup populations, this assumption may be violated for certain covariates — particularly funding-related features whose protective effect against failure may be stronger in early-stage companies than in mature ones. Violations of this assumption can bias hazard ratio estimates and reduce the accuracy of survival probability projections. Future implementations of the proposed framework should incorporate systematic testing of the proportional hazards assumption and consider time-varying coefficient extensions or stratified modeling approaches where violations are detected.

Standard SHAP computation assumes feature independence when marginalizing over excluded features, which can produce misattributed contributions when input features are correlated. In startup datasets, funding-related characteristics are structurally correlated — companies that have raised more total capital have typically completed more funding rounds — which may cause standard SHAP to distribute the collective funding signal inaccurately across individual features. Conditional SHAP estimation methods that account for feature dependence structures represent an important refinement for future work in this domain [7].

7.2 Data Quality Considerations

Startup databases are subject to several forms of systematic data quality limitation that affect any quantitative framework built upon them. Survivorship bias arises because companies that fail early and quietly are underrepresented relative to companies that raise substantial capital and receive significant media coverage. Reporting bias arises because the completeness of company records varies systematically with company size and visibility. Selection bias arises because the companies included in structured databases are not a random sample of the full startup population. These biases affect model training and may cause both the Cox model and the classifier to overestimate the importance of visibility-correlated features such as total funding and relationship count. Future work should incorporate bias correction techniques or triangulate findings across multiple data sources to reduce systematic distortion.

7.3 Static Feature Representation

Both the Cox model and the SHAP-explained classifier in the proposed framework treat company features as static values measured at a single point in time. In practice, the characteristics most relevant to startup survival and acquisition attractiveness evolve continuously as companies progress through their lifecycle. Future extensions of the framework should incorporate time-varying covariates in the Cox model to capture how changing funding conditions and operational metrics affect survival probability dynamically, and recurrent prediction architectures for the classification component to enable the acquisition likelihood score to be updated as new information about each company becomes available.

7.4 Causal Inference Extension

The current proposal is associative rather than causal: SHAP attributions and Cox hazard ratios reflect statistical associations learned from historical data, not causal relationships between company characteristics and outcomes. Future work integrating causal discovery methods — such as the PC algorithm or NOTEARS — could enable the framework to distinguish features that causally drive acquisition outcomes from those that are merely correlated with them, strengthening the actionability of framework recommendations for portfolio company management.

8. CONCLUSIONS

This paper has proposed a conceptual framework for venture capital investment analysis that integrates the Cox Proportional Hazards model and SHAP-based explainability as two complementary analytical tools. The Cox model addresses the temporal dimension of startup risk, providing survival probability estimates that capture the dynamic, time-dependent nature of company failure in ways that static classification models cannot. SHAP addresses the interpretability dimension of investment AI, providing principled, game-theoretically grounded attributions that make the predictions of opaque ensemble classifiers transparent, auditable, and actionable for investment decision makers.

The two techniques are argued to be complementary rather than redundant: the Cox model answers how long a company is likely to survive, while SHAP answers why a company is predicted to succeed or fail. Together, they produce a more complete investment risk summary — survival trajectory and explained acquisition score — that is more informative for investment practice than either technique provides in isolation.

The proposed framework addresses two critical gaps in existing quantitative venture capital analysis: the absence of temporal modeling in most predictive pipelines, and the opacity of high-performing machine learning models that limits their adoption in professional investment workflows. By combining statistical rigor, predictive performance, and principled interpretability in a single integrated pipeline, the framework offers a practical pathway toward data-driven investment analysis that meets the transparency and accountability standards of modern investment governance.

Future work should focus on empirical validation of the framework on real-world startup datasets, investigation

of conditional SHAP estimation for correlated financial features, integration of time-varying covariates into the Cox model component, and development of a user-facing decision support interface that makes the framework accessible to investment professionals without statistical expertise.

REFERENCES

- [1] Sharchilev, Boris, Roizner, Michael, Rumyantsev, Andrey, Ozornin, Denis, Serdyukov, Pavel, and de Rijke, Maarten. (2018). Web-based Startup Success Prediction. In Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM), pages 2283–2291. ACM. <https://doi.org/10.1145/3269206.3272011>
- [2] Lee, Elisa T. and Wang, John W. (2003). Statistical Methods for Survival Data Analysis. 3rd ed. Wiley-Interscience, Hoboken, NJ.
- [3] Rudin, Cynthia. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, vol. 1, no. 5, pages 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [4] Lundberg, Scott M. and Lee, Su-In. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, vol. 30. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- [5] Cox, David R. (1972). Regression Models and Life-Tables. Journal of the Royal Statistical Society: Series B, vol. 34, no. 2, pages 187–202. <https://doi.org/10.1111/j.2517-6161.1972.tb00899.x>
- [6] Shapley, Lloyd S. (1953). A Value for N-Person Games. Contributions to the Theory of Games, vol. 2, pages 307–317. Princeton University Press. <https://doi.org/10.1515/9781400881970-002>
- [7] Aas, Kjersti, Jullum, Martin, and Løland, Anders. (2021). Explaining Individual Predictions When Features Are Dependent: More Accurate Approximations to Shapley Values. Artificial Intelligence, vol. 298, page 103502. <https://doi.org/10.1016/j.artint.2021.103502>

[8] Gompers, Paul and Lerner, Josh. (2004). The Venture Capital Cycle. MIT Press, Cambridge, MA.

[9] Lundberg, Scott M., Erion, Gabriel, Chen, Hugh, DeGrave, Alex, Prutkin, Jordan M., Nair, Bala, Katz, Ronit, Himmelfarb, Jonathan, Bansal, Nisha, and Lee, Su-In. (2020). From Local Explanations to Global Understanding with Explainable AI for Trees. Nature Machine Intelligence, vol. 2, no. 1, pages 56–67.
<https://doi.org/10.1038/s42256-019-0138-9>

[10] Therneau, Terry M. and Grambsch, Patricia M. (2000). Modeling Survival Data: Extending the Cox Model. Springer-Verlag, New York.
<https://doi.org/10.1007/978-1-4757-3294-8>