FROM PEG TO PANIC: LIQUIDITY AND MARKET INDICATORS OF STABLECOIN BREAKDOWNS

by

Yevhen Nezghovorov

A thesis submitted in partial fulfillment of
the requirements for the degree of

MA in Economic Analysis.

Kyiv School of Economics

2026

Thesis Supervisor: _____ Professor Olesia Verchenko

Approved by _____
Head of the KSE Defense Committee, Professor

_____
_____
_____

Date _____

Kyiv School of Economics

Abstract

FROM PEG TO PANIC: LIQUIDITY AND MARKET INDICATORS OF STABLECOIN BREAKDOWNS

by Yevhen Nezghovorov

Thesis Supervisor: Professor Olesia Verchenko

Stablecoins are now a core part of digital financial markets, but periods of stress can quickly expose how fragile their stability really is. Even stablecoins that are backed by fiat have repeatedly deviated from the one-dollar peg, and some algorithmic designs have collapsed completely. While their importance is rising, there remains no widely accepted framework for detecting and predicting depeg events in advance. This thesis provides an empirical approach that seeks to formulate an empirical early warning system for stablecoin depeg risk. A daily panel of seven major USD-pegged stablecoins is employed to analyze daily. Depeg events are discovered with a dynamic threshold that adapts according to each token's volatility and trading conditions rather than fixed price cut-offs. The study estimates the likelihood of a depeg in the next seven days, as well as compares econometric and machine-learning models. The results indicate that liquidity deterioration is the most consistent predictor of stablecoin instability, providing a solid foundation for practical monitoring frameworks.

To my father, who sadly never saw the final result of this work, to my mother, who gave so much to my education, and to my wife, whose love, inspiration, patience, and support made this possible.

TABLE OF CONTENTS

LIST OF FIGURES v
LIST OF TABLES vi
ACKNOWLEDGEMENTS vii
AI USE DECLARATION viii
LIST OF ABBREVIATIONS ix
Chapter 1. INTRODUCTION 1
Chapter 2. LITERATURE REVIEW 5
2.1 Theoretical Foundations: Liquidity Crises and Coordination 5
2.2 Stablecoin Design and Structural Vulnerabilities 7
2.3 Measuring Depeg Events 8
2.4 Early Warning Systems for Crisis Prediction 9
2.5 Empirical Determinants of Stablecoin Instability 10
2.6 Gaps in the Literature 11
Chapter 3. METHODOLOGY 13
3.1 Defining the Depeg Event 15
3.2 Variables Construction 20
3.3 Machine Learning Extension 22
3.4 Econometric Model Specification 23
Chapter 4. DATA 31
4.1 Data Sources 31
4.2 Sample Selection 33
4.3 Raw Data Analysis 40
Chapter 5. ESTIMATION RESULTS 50
5.1 Predictive Performance 50
5.2 Discriminatory Power 52
5.3 Calibration and Operational Thresholds 57
5.4 Drivers of Depeg Risk and Interpretability 59
5.5 Robustness 64
5.6 Hypothesis Evaluation 65
Chapter 6. CONCLUSIONS AND POLICY RECOMMENDATIONS 68
6.1 Limitations and Future Research 69
6.2 Policy Recommendations 69
WORKS CITED 71

TABLE OF CONTENTS – CONTINUED

Appendix A: Threshold Policy Sensitivity 79
Appendix B: Robustness Checks Summary 81

LIST OF FIGURES

Number *Page*

Figure 1. Dynamic Threshold dt vs. rolling 30-day volatility by token type..... 20
Figure 2. Historical stablecoins price and the model's reaction to depeg events. 38
Figure 3. Mean depeg episode duration by token (days)..... 39
Figure 4. Depeg direction: episodes below vs above $1 by token. 39
Figure 5. Daily USD price by stablecoin. Source: CoinGecko. 40
Figure 6. Daily circulating supply (USD) by stablecoin. Source: DefiLlama..... 42
Figure 7. BTC and ETH daily prices (separate scales). Source: CoinGecko..... 43
Figure 8. Curve TVL by pool and token. Source: The Graph..... 44
Figure 9. Rolling 90-day positive rate (y_depeg_7d)..... 48
Figure 10. Correlation matrix of standardized predictors..... 49
Figure 11. ROC curves for Logit, Probit, Random Forest, and XGBoost. 54
Figure 12. Precision-Recall curves for Logit, Probit, Random Forest, and XGBoost..... 55
Figure 13. Lift curve (Top 1% and Top 5%) for Logit, Probit, Random Forest, and XGBoost..... 56
Figure 14. Reliability diagrams for probability calibration across models..... 58
Figure 15. SHAP summary: mean absolute impact of each feature on XGBoost predictions. 63
Figure 16. Event study: mean predicted P(depeg) around depeg days..... 64

LIST OF TABLES

Number *Page*

Table 1. Models Included in the Evaluation. 14
Table 2. Dynamic Threshold Parameters by Stablecoin Type. 18
Table 3. Train–Test Split Summary. 36
Table 4. Depeg episodes per token. 37
Table 5. Used DEX pools by token pairs. 44
Table 6. Summary statistics for model covariates. 45
Table 7. Mean values of covariates: non-crisis vs crisis token-days. 46
Table 8. Training Performance Summary. 51
Table 9. Test Performance (using train-selected thresholds). 52
Table 10. Discriminatory power: baseline and comparison models (test set). 53
Table 11. Calibration Metrics. 57
Table 12. Logit coefficients table 60
Table 13. Feature Group Ablation Results. 62
Table 14. Hypothesis evaluation summary. 66
Table 15. Alternative Threshold Sensitivity Results 79
Table 16. Summary of Robustness Validation Results. 81

ACKNOWLEDGEMENTS

My deep appreciation to my thesis advisor, Professor Olesia Verchenko, for the direction, assistance, and support throughout this work.

I sincerely thank the Kyiv School of Economics for the academic environment, high standards, and financial support to complete this project. I also want to thank the KSE faculty for their commitment and professional attitude in my studies.

AI USE DECLARATION

Various AI tools and language models have been employed in the preparation of this thesis as supplementary support instruments. Their use was limited to language editing, grammar correction, improving sentence clarity, idea generation related to empirical model design and robustness checks, assistance in identifying potentially relevant literature for independent verification by the author, and support for selected programming, raw data processing, and visualization tasks.

Any suggestions generated by AI tools were carefully reviewed, independently verified, and substantially revised by the author before being incorporated into the thesis. Code-related assistance was used only when fully understood and reproducible by the author. AI tools were not used to generate research findings, conduct autonomous data analysis, produce literature summaries used as source material, or create citations.

The research questions, theoretical framework, hypotheses, methodology, empirical analysis, interpretation of results, and conclusions were developed independently by the author. The author takes full responsibility for the originality, validity, reproducibility, and academic integrity of this thesis.

LIST OF ABBREVIATIONS

AMM. Automated Market Maker.

Brier. Brier Score.

BTC. Bitcoin.

CDF. Cumulative Distribution Function.

CEX. Centralized Exchange.

CI. Confidence Interval.

DAI. Dai Stablecoin (crypto-collateralized).

DeFi. Decentralized Finance.

DEX. Decentralized Exchange.

ECE. Expected Calibration Error.

ETH. Ether (Ethereum native token).

ETL. Extract, Transform, Load.

EWS. Early Warning System.

F1. F1-Score.

FRAX. Frax Stablecoin (fractional-algorithmic).

LUSD. Liquity USD Stablecoin.

MAD. Median Absolute Deviation.

MCC. Matthews Correlation Coefficient.

MIM. Magic Internet Money Stablecoin.

OR. Odds Ratio.

PR-AUC. Precision-Recall Area Under the Curve.

PYUSD. PayPal USD Stablecoin.

ROC. Receiver Operating Characteristic.

ROC-AUC. Receiver Operating Characteristic Area Under the Curve.

SHAP. SHapley Additive exPlanations.

SVB. Silicon Valley Bank.

TVL. Total Value Locked.

USDC. USD Coin (fiat-backed).

USDe. Ethena USDe (synthetic dollar).

USDT. Tether (fiat-backed).

UST. TerraUSD (algorithmic stablecoin).

USTC. TerraClassicUSD (formerly UST, algorithmic).

XGBoost. eXtreme Gradient Boosting.

Chapter 1

INTRODUCTION

The stablecoin market has seen unprecedented scale recently, demonstrating the growing significance of this asset class. Over 1% of U.S. dollars today exist as tokenised stablecoins on public blockchains, and stablecoin issuers are now the 17th largest holders of U.S. Treasuries globally, collectively managing over $150 billion in government debt (Andreessen Horowitz 2025). The significance of stablecoins as a critical monetary instrument is illustrated by this large macro capitalization, their annual transaction volume of $46 trillion (Andreessen Horowitz 2025), and their total market capitalization of US dollar-based stablecoins exceeded $300 billion (DeFiLlama 2026). As a result, they are widely used to exchange fiat currencies and volatile cryptocurrencies such as Bitcoin (BTC) and Ethereum (ETH).

But the explosive growth of this sector is also met with a profound fragility. Stablecoins are designed with a stable value, but that stability is a "developing and fragile state" that is almost entirely dependent on continued confidence in the market and on continuous liquidity. This concentration has made them vulnerable to devaluation threat that occurs, along the lines of the situation existing in pegged exchange rate regimes, a feature common to many traditional currencies (Eichengreen, Nguyen, and Viswanath-Natraj 2025). The recent events have highlighted these shortcomings clearly. The failure of TerraUSD (UST) stablecoin in May 2022 took the market capitalization down nearly USD 40 billion (Briola et al. 2023) due to a loss of confidence and mass withdrawals brought about by the run (Krause 2025). Fiat-backed stablecoins, which have long assumed the safest option, are still exposed. To illustrate, the USD Coin (USDC) lost its peg in March 2023 because of concerns about reserves stored at collapsed Silicon Valley Bank (SVB). This episode demonstrates how a lack of transparency and liquidity imbalances can quickly damage confidence. These events highlight the need to

construct sound and timely risk detection utilities. This study applies the Early Warning Systems (EWS) framework, in which speculative attacks on fixed exchange rates (Eichengreen, Rose, and Wyplosz 1995; Kaminsky, Lizondo, and Reinhart 1998) can be anticipated for the contemporary stablecoin market. Currency crisis theory (First-Generation Models: crises which flow from unsustainable fundamentals, e.g., Krugman 1979; Second-Generation Models: crises that are self-fulfilling products of expectations of investors, e.g., Morris and Shin 1998; Obstfeld 1996) provides the theoretical basis for a view of stablecoin runs as analogous speculative events (MacDonald and Zhao 2022). Similar to Bucevska (2015) and Frankel and Rose (1996), the aim of this paper is to apply an econometric approach and an estimation of the likelihood of a crisis occurring.

In this thesis, we will address the following central research question: Can we build and test an Early Warning System to predict when a stablecoin will lose its peg?

We want to introduce and test the system that would offer real-time signals for the stability of stablecoins. Based on the indicator framework introduced by Lee, Chiu, and Hsieh (2025), we examine the most important warning signs (liquidity imbalances and market stress and others) and provide a study on the predictive accuracy of the model in the case of recent episodes of depegging. One significant aspect of this research is the use of a dynamic threshold adjusted for trading volume as a criterion for depegging. This approach is derived from Lee, Chiu, and Hsieh (2025), which expands from static or subjective threshold approaches utilized in previous work, to facilitate a more accurate classification of crisis events for stablecoins demonstrating varying liquidity profiles.

Based on established theories about financial crises and recent empirical evidence, the following hypotheses are tested in this study.

In the first example, the Liquidity Hypothesis (H1) suggests that a decrease in on-chain liquidity substantially increases the chance of a stablecoin losing its peg. In

this context, liquidity corresponds to any lowering of Total Value Locked (TVL), lower trading volume, or severe imbalances in DEX pools.

Second, the Supply-Shock Hypothesis (H2) states that rapid 7-day changes in circulating supply (net issuance or redemption) are associated with a higher probability of a depeg episode in the following week. The expected sign for correspondent predictor is positive: large positive surges (sudden issuance) historically precede stress episodes in the sample; the hypothesis is therefore one-sided.

The Market-Return Hypothesis (H3) provides further evidence whereby a high level of systemic stress throughout the entire cryptocurrency market sharply raises the risk of a stablecoin crash. As per usual empirical practices, such stress is represented by high price volatility and negative returns of mainstream crypto assets namely, Bitcoin and Ethereum (Eichengreen, Nguyen, and Viswanath-Natraj 2025; Lee, Chiu, and Hsieh 2025).

Our study offers valuable lessons to stakeholders in the digital asset ecosystem. These results may be beneficial to stablecoin issuers and decentralized finance developers, for improving monitoring of risk and in order to maximize liquidity management practices. Crypto exchanges and financial risk managers might analyze the developed signaling patterns through these early warning indicators to detect signs of impending instability before it grows into a full depeg episode. These findings also promote policy response, empowering regulators and policymakers to employ visible liquidity and supply variables to track systemic risk in real time, designing fair balances within digital markets. This framework can also be further developed with the help of academic researchers as we strive to further extend crisis prediction models in cryptocurrency markets.

The thesis is organized roughly as follows: In Chapter 2 we review relevant literature on financial crises, stablecoin design and early warning models. Chapter 3 develops the methodology (depeg events and econometric and machine-

learning specifications). Chapter 4 reports data sources, sample construction and raw descriptive statistics. The estimation results, and robustness checks, are described in Chapter 5. Chapter 6 discusses policy implications.

The preliminary results indicate that liquidity conditions are the most consistent variable to predict stablecoin depeg risk with supply and broader market signals appearing to be weaker and context-dependent. Model comparative evidence further supports a tension between predictive ranking performance and operational signaling stability and encourages the use of adaptive monitoring alongside conventional threshold rules.

Chapter 2

LITERATURE REVIEW

For example the literature relating to the risks caused by stablecoins draws on classical models of financial crises such as bank runs and currency attacks (Diamond and Dybvig 1983; Morris and Shin 1998), and more recent digital asset risks, such as algorithmic "death spirals" or sudden panics caused by unclear reserves (Anadu et al. 2024; Liu, Makarov, and Schoar 2023). This combination allows us to view stablecoin crashes (depegs) as updated forms of classical liquidity crises (Eichengreen, Nguyen, and Viswanath-Natraj 2025). As with traditional pegged currencies, stablecoins often evaporate behind weak fundamentals or a sudden surge of panic, so traditional Early Warning System (EWS) models are highly applicable to the crypto market.

Yet to predict digital crises, we need to tweak these models. Where most EWS research has been based on slow macroeconomic data (Kaminsky, Lizondo, and Reinhart 1998), stablecoin markets output large quantities of fast, real-time blockchain data. Hence, more recent studies integrate common tools with specific crypto metrics including decentralized exchange (DEX) liquidity and market stress (Cintra and Holloway 2023; Lee, Chiu, and Hsieh 2025). Yet many things remain not sync among researchers. There is no accepted framework for defining a depeg, which warning signs are the best, or how to assess models when crises happen in exceptionally few cases. This chapter reviews basic financial crisis theories, recent empirical studies on stablecoins, and standard EWS methods, explaining how this thesis aims to contribute to filling the research gaps and explains our research strategy.

2.1 Theoretical Foundations: Liquidity Crises and Coordination

The primary objective of this research is to construct an Early Warning System (EWS) for stablecoin instability. Assigning a stablecoin depeg as a modern

speculative attack allows for the application of existing models of economic modeling to the crypto market (Eichengreen, Nguyen, and Viswanath-Natraj 2025).

This approach is strongly supported upon the currency crises literature. Theories of the economy "first-generation" models posit that crises are due to the fact of weak economic fundamentals. Krugman (1979), for example, demonstrates how excessive money creation gradually drains reserves until a speculative attack emerges as inevitable. Later, "second-generation" models confirmed that market expectations are as well a huge component. Even in a genuinely healthy economy, the coordinated panic of investors when suddenly certain a peg will fail may lead even to a devaluation (Obstfeld 1996). To build onto this, Morris and Shin (1998) illustrate that when investors get unclear signals, even small news updates trigger huge market movement and self-fulfilling crises.

Researchers designed Early Warning Systems to help them forecast such complex events. The early study by Kaminsky, Lizondo, and Reinhart (1998) put forward the "signals" approach. They showed that certain economic variables behaved abnormally in the months leading up to a crash. Following these leading indicators helps in estimating the likelihood of a future crisis.

The logic of EWS is still extremely relevant for stablecoins today. Because stablecoins are designed to keep a fixed value against the US dollar, they can have the same problem of credibility. If coin holders anticipate liquidity trouble, their collective panic can easily break the peg. As in the case of traditional currency crises, stablecoin instability results from a combination of weak fundamentals and investor expectations. As a result, the risk of a crash builds up gradually until finally it manifests itself in the coin's price (Eichengreen, Nguyen, and Viswanath-Natraj 2025; Liu, Makarov, and Schoar 2023).

2.2 Stablecoin Design and Structural Vulnerabilities

Stablecoins are based on the idea that they should keep the price at a stable point, typically against the US dollar, but they do this by very different means. These variant designs require different levels of risk (Bullmann, Klemm, and Pinna 2019). Fiat-backed stablecoins are based on conventional financial reserves. Their fundamental weakness is in reserve credibility; if investors have a doubt about whether these assets are good enough or liquid, then sudden panic can cause short-term price declines. Crypto-collateralized stablecoins provide backing from volatile crypto assets, utilizing over-collateralization for their protection. While doing so removes central control, they are extremely vulnerable to overall crypto market crashes. Algorithmic stablecoins utilize smart contracts and linked tokens rather than real reserves. All they rely on is market confidence. Without trust, these can be trapped in a "death spiral" and completely fail (Liu, Makarov, and Schoar 2023). Depeg risks are not identical for each and every token as a result of these variant designs. Accordingly, an Early Warning System (EWS) should account for token-related characteristics as well as market risk. In order to predict these breakdowns accurately, some current empirical researches have used a number of warning signals in combination. First, massive market stress (e.g., sudden plummeting or extreme volatility of Bitcoin) often induces market panic in all types of stablecoins (Lee, Chiu, and Hsieh 2025). Second, researchers emphasize the importance of on-chain liquidity dynamics. For instance, severe imbalances in decentralized exchange (DEX) pools can signal an impending depeg before it becomes fully visible on centralized markets (Cintra and Holloway 2023). Finally, massive reserve outflows or sudden spikes in trading volume often act as direct evidence of a "flight-to-safety" bank run (Anadu et al. 2024; Eichengreen, Nguyen, and Viswanath-Natraj 2025). By monitoring these combined indicators, an EWS can accurately estimate the probability of a future crisis.

2.3 Measuring Depeg Events

In this thesis, we expect one to arise, what we may like to call a stablecoin devaluation or "depeg" event. The risk of loss of fixed value of a stablecoin is essentially similar to the run risk of commercial banks (Diamond and Dybvig 1983) and the devaluation risk for national currencies with pegged exchange rate systems (Eichengreen, Nguyen, and Viswanath-Natraj 2025). To empirically quantify this risk, an early warning system requires specifying exactly what creates a crisis. The current literature, however, lacks a consensus framework to evaluate these events, resulting in dramatically differing crisis counts in different methodologies.

Most empirical studies do not use a single universal threshold to define a depeg. Alternatively, the academic literature combines threshold-based theoretical models with empirical approaches that incorporate depegging rules adjusted for trading volume (Lee, Chiu, and Hsieh 2025). Static rules are simple to implement, but can lead to comparison issues as stablecoins differ in baseline volatility, collateral structure and market liquidity. To sum up, a 1% decrease in price could be very painful for a fiat-backed coin, but the same change could be a normal noise for a more volatile token especially during low volume trading times (Lee, Chiu, and Hsieh 2025).

Considering these limitations, one recent study suggests that we can develop dynamic thresholds to capture prevailing market characteristics such as rolling volatility and volume (Lee, Chiu, and Hsieh 2025). It is this volume-adjusted approach which forces an event to experience significantly larger price deviations in low trading volume periods so that it might be defined as a real crisis. This in effect filters out false signals resulting from usual market noise and recognizes that the bottom-line interpretation of the price decrease is dependent and heavily influenced by liquidity conditions in the market.

Another key methodological decision is whether to consider every single depegged day as a distinct crisis observation or only when the depeg begins. Counting each

day under the threshold as a new event artificially inflates the number of positive observations, redirecting attention toward the duration of the crisis. On the other hand, establishing the event as the first day a threshold is breached emphasizes the analysis of how early a state of instability starts. This is much closer to the idea of an early warning system, which is to predict rising risks long before prolonged financial stress emerges (Kaminsky, Lizondo, and Reinhart 1998). Pursuant to this rationale, this thesis's empirical framework employs a dynamic, volume-adjusted threshold and models the onset of depeg episodes rather than their duration.

2.4 Early Warning Systems for Crisis Prediction

EWS (Early Warning Systems) were created to predict the likelihood of financial crises before they occur. The model of this domain is based on Kaminsky, Lizondo, and Reinhart's (1998) "signals" approach. They demonstrated that some economic factors behave oddly in the lead up to a crisis and serve as signals of imminent danger when they break some threshold. Two fundamental principles underlie this model: crises are driven by measurable economic deterioration, and predictions should be predictive. Later studies have advanced to multivariate logit and probit models (Comelli 2014; Frankel and Rose 1996) that gained popularity for predicting crises. These binary-response models describe the probability of a crisis from multiple values simultaneously. The most important benefit is transparency, since coefficient signs and marginal effects clearly describe the relationship between a single predictor and overall crisis risk. However, there is a problem in this space called the rare-event problem. Crises account for a very small part of the data, so standard accuracy metrics can be hugely misleading.

Hence, new EWS assessment involves estimation and ranking power not simple classification. The Receiver Operating Characteristic (ROC) curve and its Area Under the Curve (ROC-AUC) are often utilized to evaluate how well a model knows a threshold-free gap between crisis and non-crisis periods (Yildirim and Sanyal 2022). When dealing with highly imbalanced datasets, precision-recall

metrics (PR-AUC) provide significantly larger insight regarding the model performance on the minority class (Barthélémy, Gautier, and Rondeau 2024). Because the use of operations depends on a rule for decision-making, finding the final threshold is still an issue of policy compromise between the cost of false alarms and failure to detect a crisis.

The past few years witnessed a move towards high-dimensional modeling, which in turn made it possible to incorporate machine learning techniques (in particular the Random Forests and Gradient Boosting techniques) into existing EWS models. Flexible algorithms are implemented to represent complex non-linear interactions and their prediction accuracy can sometimes be better than that of classical parametric models (Lee, Chiu, and Hsieh 2025). Nevertheless, since machine learning models are significantly less interpretable than traditional statistical models, they are thus a trade-off between achieving predictive accuracy and interpretability (Barthélémy, Gautier, and Rondeau 2024). It is these methodological tools that permit us to define stablecoin depegs as rare binary events, within limits determined by a predictive, structured framework.

2.5 Empirical Determinants of Stablecoin Instability

On an empirical level, some empirical works have highlighted several of the major drivers of stablecoin instability systematically (Ante et al. 2021; Lee, Chiu, and Hsieh 2025). Historically, these contributors have been subdivided into four main or interconnected groups: liquidity situation, market stress, supply relationships, and token structure. For one, liquidity ensures stability in price. Total Value Locked (TVL), trading volume and liquidity concentration metrics (e.g. for DEX pools) measure market depth. The liquidity of a single pool, when liquidity is low or is very concentrated, means that small trades can suddenly turn all-time lows in prices and that stablecoins become extremely fragile in panic episodes for many investors (Cintra and Holloway 2023).

A second reason to consider this alternative is that market stress indicators constitute an index of the actual stresses. The market signals on crypto as a whole, particularly the returns and volatility of Bitcoin and Ethereum, represent a solid proxy for systemic risk. Further, there is enough evidence that periods characterized by weak returns or excessive volatility are closely associated with a high depegging risk (Lee, Chiu, and Hsieh 2025). Third, alterations in circulating supply are aligned to investor activity. A sudden supply decline usually leads to strong redemption pressure (Eichengreen, Nguyen, and Viswanath-Natraj 2025), where investors want to liquidate their wealth before a potential crash hits. Lastly, the design of a stablecoin has a clear and significant effect on its overall risk profile. Fiat-backed, crypto-collateralized, and algorithmic stablecoins have extremely different exposures in terms of reserve quality, collateral volatility, and investor coordination failures (Gadzinski et al. 2023). So, while structural design captures baseline differences in risk across stablecoins, a predictive model should primarily rely on time-varying liquidity and market variables. Together, these groups of variables provide a solid, theoretically grounded foundation for estimating the probability of a depeg event within an early warning framework.

2.6 Gaps in the Literature

Despite the tremendous progress towards stablecoins and crisis prediction research, there have been several large gaps that still exist in the existing research. One of these issues is that a depeg event has no common, dynamic definition. Many of the analyses also continue to use fixed price thresholds which do not take into account the essential liquidity status, volatility variations, and variation in token constructions (Duan and Urquhart 2023). As such, crisis detection is not consistently applied across its assets and environments. In order to describe and to measure financial stress in a meaningful way, it is becoming more important to have variable, volume-corrected definitions of event (Lee, Chiu, and Hsieh 2025). Moreover, while researchers generally agree that liquidity is a key driver of

stablecoin stability, DEX metrics (e.g., Total Value Locked (TVL), trading volume, pool concentration), are seldom combined into formal early warning systems. Rather than using these critical indicators as integrated points in structured, multivariate probability models, existing models almost exclusively address them as a descriptive or a single-case study only (Cintra and Holloway 2023). Finally, the existing literature rarely applies strict, time-based out-of-sample tests. However, since the training of a given model is not in a particular period nor separate from a test period, the predictive power can be greatly exaggerated, particularly in markets susceptible to abrupt structural breaks and complex post-crisis dynamics (Barthélémy, Gautier, and Rondeau 2024). Filling these gaps will require an integrative approach. This thesis explicitly addresses these limitations, as an empirical basis, it aims to clarify event scope, including DEX liquidity constructs, and also to adopt transparent threshold policies and hard temporal validation across competing models.

Chapter 3

METHODOLOGY

Our empirical strategy is guided by the leading-indicator approach used in early warning systems of currency crises (Kaminsky, Lizondo, and Reinhart 1998; Berg and Pattillo 1999). We view the depegging of a stablecoin as an all-or-nothing event. It is our intention to measure the likelihood of a token experiencing crisis within 7 days given anything that can be observed.

Conventional macroeconomic early warning systems typically forecast crises with a 12 to 24-month lead time. Nonetheless, cryptocurrency markets have much faster adjustments, so such long horizons are not realistic. Recent studies indicate that high-frequency data is vital in identifying distress in digital assets. For instance, Cintra and Holloway (2023) state that warning signals for stablecoin depegs often only emerge mere hours before the depeg. In the same manner, Lee, Chiu, and Hsieh (2025) develop daily risk prediction models and argue that low-frequency models are slow in responding to risk. In this study, we employ daily data to capture rapid changes in risk dynamics in cryptocurrency markets. This choice is also driven by data availability, as daily frequency ensures consistent coverage across all selected predictors while allowing the analysis to span a sufficiently long time period, which is essential for capturing rare depeg events.

Our methodology is sequential in nature. To start with, we define in a clear and rules-based manner what a depeg is, which further determines our dependent variable. The definition is not based on fixed arbitrary limits. Rather, it adjusts based on the recent volatility of every token.

We will now define our explanatory variables. The variables will be selected in ways to capture liquidity stress, trading activity, supply conditions, co-movement with others, and token design features.

We use econometric and machine learning techniques to predict the chances of a depeg event. We adopt Logit and Probit binary-response models for testing hypotheses with well-defined economic interpretation (Comelli 2014). In addition, we use models, for example Random Forest and XGBoost, since they can model complex non-linear relationships without restricting them to a rigorous linear structure.

Table 1. Models Included in the Evaluation.

Model	Role	Interpretation
Logit	Baseline specification	Parametric binary-response model; log-odds.
Probit	Robustness benchmark	Alternative parametric specification
Random Forest	Non-linear extension	Captures complex feature interactions.
XGBoost	Flexible gradient boosting model	Benchmark for ranking performance.

Model performance is evaluated by two complementary dimensions: ranking ability and operational accuracy.

First, ranking performance is measured using the Area Under the Receiver Operating Characteristic Curve (ROC-AUC) and the Precision-Recall Area Under the Curve (PR-AUC). These metrics evaluate the model's ability to rank crisis risk independently of a chosen threshold. In highly imbalanced rare-event settings such as stablecoin depeg prediction, ROC-AUC may appear high even when the model poorly identifies actual crisis events (Barthélémy, Gautier, and

Rondeau 2024). Therefore, PR-AUC is particularly informative, as it focuses on the minority class and better reflects performance in detecting crises.

Second, operational performance evaluates the model based on a selected classification threshold, which determines alert frequency, precision, and recall. The threshold is chosen by maximizing the F1-score on the training set, balancing precision and recall (Xu, Li, and Wu 2024). If this threshold yields fewer than 0.05 alerts per day in the test sample, a minimum-alert constraint is imposed to maintain sufficient sensitivity. Finally, Youden's index is reported for comparison, but is not used as the primary decision rule.

The primary objective of this study is to predict and rank stablecoin depeg risks rather than identify causal relationships. To ensure a proper out-of-sample evaluation, models are tested on a strictly time-based split (Barthélémy, Gautier, and Rondeau 2024). According to threshold-based decision rules, predicted probabilities are turned into meaningful early warning signals (Bussière and Fratzscher 2006; Kaminsky, Lizondo, and Reinhart 1998). Ranking and operational metrics are used to measure model performance as above.

3.1 Defining the Depeg Event

Specifying precisely what is a "depeg" event has always been a problem in recent literature. Earlier empirical studies often relied on arbitrary, static thresholds to identify a crisis. For example, researchers and industry reports define a depeg as a fixed price deviation of 3% (Moody's 2023), or even 10% (S&P Global 2023) from the $1 peg.

But one single static rule on a stablecoin introduces a major methodological error. If we apply the fundamental price deviation ($d_t = |p_t - 1|$) (the absolute distance from parity) to stablecoins for all its real value, the model fails to capture the genuine crisis. Different types of stablecoins have intrinsic and inherent variations in price volatility. Roughly, a 1% plunge would represent the fiat-

backed coin going through a massive crisis, and in the case of an algorithmic coin would simply be the normal daily trading noise. Furthermore, price deviations alone are not enough, and low trading periods can create misleading signals and false alarms.

In order to mitigate this limitation, the current study deviates from a strict set of static rules, and introduces a dynamic, volume-adjusted threshold. This is also adapted to the volatility and market properties of each stablecoin, based on frameworks by Lee, Chiu, and Hsieh (2025), and Cintra and Holloway (2023). We derive an operational definition in which a crisis is identified in terms of a sequence of specific conditions.

This step constructs a dynamic threshold for detecting deviations from the peg. The baseline tolerance, denoted as α_t, adapts to recent price instability for each token over a rolling 30-day window:

The threshold is calculated with the following equation:

$$\alpha_t = \text{clip}\ (\alpha_{\text{base}} + 2 \cdot \sigma_{\text{roll}} + 0.5 \cdot \text{MAD}_{\text{roll}}, \alpha_{\text{base}}, \alpha_{\text{max}}) \quad (1)$$

where σ_{roll} is the rolling 30-day standard deviation of log returns, and MAD_{roll} is the rolling 30-day median absolute deviation in the price.

The clip(·) operator limits the threshold within predefined bounds, ensuring that $\alpha_t \in [\alpha_{\text{base}}, \alpha_{\text{max}}]$. This keep the threshold adaptive to volatility while avoiding economically unreasonable values.

The lower bound αbase represents a minimum tolerance level interpreted as an early-warning threshold below which deviations are considered normal. The upper bound αmax corresponds to stress conditions, beyond which deviations indicate significant peg instability. Together, these parameters separate usual fluctuations from meaningful dislocations.

The threshold design also adheres to the volatility-scaling methodology commonly employed in financial econometrics and risk modeling. Volatility is a time-varying scale parameter (Engle 1982; Bollerslev 1986), and operational thresholds are derived on top of this scaling of dispersion measures (Jorion 2007; Duffie and Pan 1997). The applied coefficient 2 for σ_{roll} serves as a bandwidth multiplier, tying admissible deviations to the current volatility regime in place. Because returns are heavy-tailed, scaling by 2 provides a statistically grounded and widely used rule to separate normal movements from anomalous movements.

MAD_{roll} introduces a more robust measure of dispersion, less sensitive to extreme observations (Huber 1981; Rousseeuw and Croux 1993). It appears to perform a secondary role; since volatility-driven adjustments still prevail, MAD provides a stabilizing correction without overstating the impact of outliers, i.e., 0.5. This weighting captures the trade-off between sensitivity and robustness.

Accordingly, this specification comprises a calibrated heuristic, rather than one constructed from one theoretical architecture. This is driven by the early-warning system literatures wherein dispersion measures are used to establish thresholds for crisis (Kaminsky, Lizondo, and Reinhart 1998; Comelli 2014; Eichengreen, Rose, and Wyplosz 1995), and recent work on volatility-adjusted thresholds in digital asset markets (Lee, Chiu, and Hsieh 2025). Both coefficients (2, 0.5) serve as economically interpretable parameters, and they are not considered generalisable; similar models normally check calibration and robustness to optimise detection accuracy and false signals (Berg and Pattillo 1999; Bussière and Fratzscher 2006).

The bounds α_{base} and α_{max} are fine-tuned as empirically reported and as structural characteristics vary among stablecoin designs. Fiat-backed stablecoins exhibit tighter price responses because active arbitrage operates, whereas crypto-backed and algorithmic designs have higher volatility and tail risk (Bullmann, Klemm, and Pinna 2019; Lyons and Viswanath-Natraj 2023; Liu, Makarov, and Schoar 2023).

Table 2. Dynamic Threshold Parameters by Stablecoin Type.

Type	α_{base}	α_{max}
Fiat-backed	0.4%	1.5%
Crypto-backed	0.7%	2.5%
Algorithmic	1.0%	4.0%

The volatility profile and instability tolerance of stablecoins, as a function of their structural design, dictates that differentiated threshold ranges are justified. The tightest price deviations for the fiat-backed stablecoins, on the other hand, are driven by the backing of liquid off-chain assets and the efficient nature of redemption and arbitrage that links prices near parity (Lyons and Viswanath-Natraj 2023). Meanwhile, crypto-collateralized stablecoins operate on the basis of volatile digital assets and are susceptible to liquidation cascades which disrupt wider ranges of value due to market stress and deviance due to the open liquidation of the digital asset (Klages-Mundt and Minca 2021). Algorithmic and hybrid stablecoins tend to have the highest volatility and tail risk but are inherently endogenous as well as dependent on market confidence, causing them prone to failure of reflexivity and coordination such as exhibited by the TerraUSD crash (Clements 2021; Liu, Makarov, and Schoar 2023).

Due to these structural differences, it would be inappropriate to apply a single static depeg threshold across all stablecoin types. The literature itself uses multiple thresholds of deviation, usually around 1–3% for operational definitions and considerably higher for systemic stress events (Cintra and Holloway 2023; S&P Global 2023). As a result, the parameter ranges in Table 2 are calibrated for both the empirical evidence and structural risk differences. While a tighter boundary is applied to fiat-backed stablecoins in order to allow early deviation detection, wider

bands will be observed in crypto-backed and algorithmic designs that take into consideration their greater inherent volatility and multiple equilibrium dynamics. That is, such values need to be interpreted as calibration values prescribed through empirical testing, not threshold numbers tossed around.

The second step corrects low volume to avoid false alarms in inactivity of the market. We add a volume ratio variable (vol_ratio_t, 30 days trading volume / median 180-day volume, capped at 10 to avoid extreme outlier effects):

$$vol_ratio_t = min\left(\frac{vol_{30d}}{median(vol_{180d})}; 10\right) \quad (2)$$

The specification employs common benchmarking procedures for trading activity and anomaly detection. *vol_ratio*$_t$: short term turnover scaled against the long term baseline. The use of 180-day median reduces sensitivity on extreme observations typical of crypto markets (Liu, Makarov, and Schoar 2023), and the cap at 10 serves as a winsorization rule limiting extreme spikes (Angrist and Pischke 2009).

The thresholds state economically meaningful regimes. *vol_ratio*$_t$ < $0.5\alpha_t$, means that trading activity is abnormally low (Amihud 2002; Kyle 1985), which means that the market is illiquid, which occurs when a lower number of small trades increases the price sensitivity of a stock (Amihud 2002). In order to compensate for this, the dynamic threshold α_t is increased by a factor of 1.3 to reduce false signals in thin markets, corresponding to Early Warning System design (Kaminsky, Lizondo, and Reinhart 1998).

A nominal day t is said to be a depeg day if $d_t > \alpha_t$ and, likewise, either the deviation is exceptionally large ($d_t > 1.5 \cdot \alpha_t$), according to a volatility-based event definition (Bussière and Fratzscher 2006) or trading activity is abnormally high (*vol_ratio*$_t$ > *3*), reflecting volume spikes triggered by the stress effect (Andersen 1996; Liu, Makarov, and Schoar 2023).

$$(d_t > \alpha_t) \wedge [(d_t > 1.5\ \alpha_t) \vee (vol_ratio_t > 3)] \quad (3)$$

A depeg episode starts on the first day these conditions are met following a period of non-depeg days. This allows the Early Warning System to focus on predicting the beginning of a crisis rather than counting every day a coin remains depegged.

Figure 1 illustrates the dynamic threshold by token type.

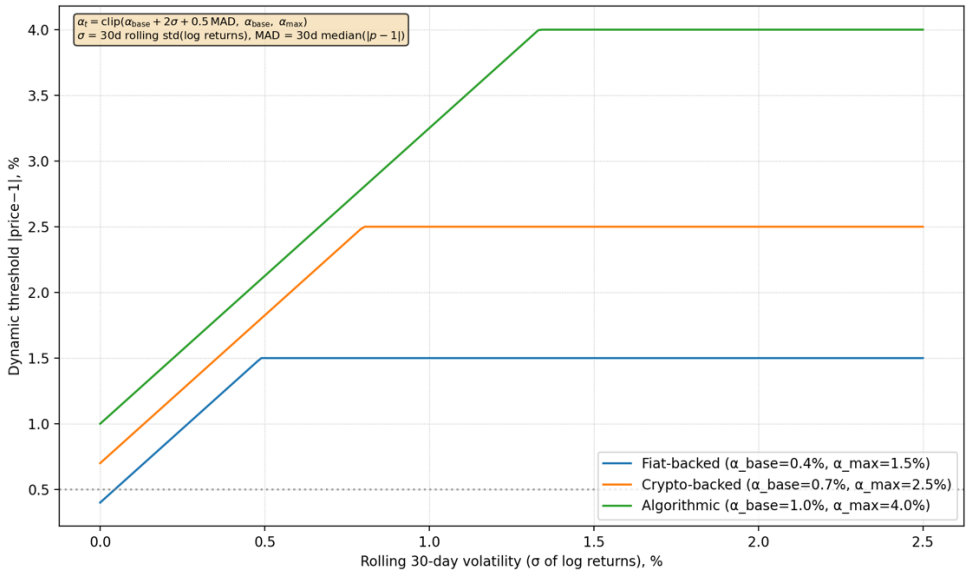


Figure 1. Dynamic Threshold dt *vs. rolling 30-day volatility by token type.*

3.2 Variables Construction

We define all model variables at the token-day level. The data sources employed to develop these variables are outlined in the Data part. The dependent variable and independent predictors here are defined:

y_depeg_7d – binary indicator 1 if a depeg episode begins for a given token for any of the subsequent 1 to 7 days; 0 otherwise.

Instruments vector of price and stress indicators include:

abs_dev – a continuous variable that represents the absolute deviation from the USD peg, calculated as $|price_usd - 1|$. It measures how far the stablecoin trades from its peg; larger values indicate higher price stress;

binance_volume_30d – a continuous variable measures the 30-day rolling trading volume for the stablecoin on Binance (in USD). This variable captures the centralized trading volume conditions and is used to test whether increases in CEX trading activity raise the probability of a depeg event.

The volume and liquidity covariates vector:

log_vol_30d – continuous variable and natural logarithm of 30-day rolling trading volume denoted in USD, defined as *ln(vol30d+1)* - a simple transformation which is used to reduce skewness, stabilize variance and retain zero observations in heavy-tailed financial data (Wooldridge 2010; Liu, Makarov, and Schoar 2023).

dex_tvl_usd – continuous variable of the total tokens in the DEX pools, which is the sum of the total amount of the token balance available in each of the associated pools as measure of the on-chain liquidity level;

dex_volume_30d – continuous variable for the rolling 30-day trading volume of the decentralized public exchanges;

dex_share_max – continuous variable representing the maximum ratio of token balance to pool TVL of all token pools. A decrease in this share indicates liquidity fragmentation or loss of capital to the primary pool;

cex_volume_30d – continuous variable for the rolling 30-day trading volume on the central market, zeroed if no CEX trading pair is represented.

Vector of supply and market context variables includes:

supply_usd – continuous variable representing the total circulating supply of the token expressed in USD; while not necessarily stationary, it captures the level

effect and is complemented by supply_pct_change_7d to account for short-term dynamics.

supply_pct_change_7d – continuous variable for the percentage change in circulating supply over the prior 7 days, calculated as $(supply_usd_n - supply_usd_{n-7})/supply_usd_{n-7}$. It captures the effect of supply shocks (H2);

We specifically separate the warning signals from the event itself by specifying the target over a forward-looking 7-day ($y_{it}^{(7)}$) horizon while evaluating the predictors only at the current day$_t$ ($\mathbf{x}_{it}$). Also, this time-based separation diminishes simultaneity bias and guarantees that the model functions as a functional early warning system. It predicts on the basis of only the information available before the crisis, rather than on data collected during the period when market panic has already started.

3.3 Machine Learning Extension

dominance_eth – continuous variable for the share of the token's supply held on the Ethereum as the main DeFi network relative to its total supply (*supply_eth*/*supply_total*), as a ratio, it is more stable over time than level variables and captures chain concentration;

btc_return_7d and *eth_return_7d* – continuous variables for the 7-day logarithmic returns of Bitcoin and Ethereum prices, defined as ln(btc_price_t/ btc_price_{t-7}) and ln(eth_price_t/ eth_price_{t-7}), proxying the broad crypto market momentum;

token_type – categorical variable taking values for fiat-backed, crypto-backed, or algorithmic stabilization mechanisms. It is included in the model as dummy variables, with fiat-backed tokens serving as the reference category.

The estimation is performed after standardizing all continuous variables to zero mean and unit variance. The standardization parameters (mean and standard

deviation) are computed using only the training set to avoid data leakage and ensure proper out-of-sample evaluation (Barthélémy, Gautier, and Rondeau 2024). This transformation ensures that variables measured on different scales do not distort the estimation. It does not impose stationarity, which is not required in this modeling framework. This transformation ensures that variables measured on different scales are comparable and prevents features with large magnitudes from dominating the estimation, particularly in scale-sensitive models. The categorical token_type variable is predefined based on known token categories and included in the model as categorical indicators. Finally, to assess the model's predictive performance and simulate real-world forecasting, we split the dataset using time-based split criteria. The training set is assigned the first 80% of the observations in chronological order and the out-of-sample test set is assigned the last 20% (Hyndman and Athanasopoulos 2021). Early warning systems literature strongly implies that the temporal order of the financial time series should be maintained to mitigate look-ahead bias and thus standard random sampling splits would not be recommended (Sarlin 2014; Barthélémy, Gautier, and Rondeau 2024).

3.4 Econometric Model Specification

We take a discrete-dependent-variable based approach to explore the determinants of stablecoin instability. This decision was drawn from the many of publications on Early Warning Systems (EWS) of sovereign currency and banking crises. Although non-parametric signal extraction proposed by Kaminsky, Lizondo, and Reinhart (1998) predominated, multiple studies have shown that multivariate parametric models, such as Logit and Probit, are also superior in predictive ability, and in determining the simultaneous statistical significance of different warning signs (Frankel and Rose 1996; Berg and Pattillo 1999; Davis and Karim 2008).

Recently, the choice of this methodology is verified by empirical work relating to stablecoins. Lee, Chiu, and Hsieh (2025) applied logistic regression to predict depegging risks from on-chain and market data and is used as a benchmark for this type of asset. Moreover, Comelli (2014) emphasizes the usefulness of comparing both Logit and Probit specifications for maintaining robustness for characterizing the tail risks of financial stress.

Following the standard framework for limited dependent variables found in the currency crisis literature (Eichengreen, Rose, and Wyplosz 1995), we model the probability of a future depeg event as a function of the currently observed covariates. As defined in Section 3.1, the dependent variable y_{it+7} is a forward-looking binary indicator taking the value of 1 if a depeg episode starts for stablecoin i within the next 7 days, and 0 otherwise.

The conditional probability of a depeg occurring is modeled as:

$$P\left(y_{it}^{(7)} = 1 \mid \mathbf{x}_{it}\right) = G(\mathbf{x}_{it}'\boldsymbol{\beta}), \tag{4}$$

where $\mathbf{x}_{it}$ is a vector of standardized explanatory covariates (including liquidity, trading volume, and price stress indicators) observed at day t. To account for unobserved heterogeneity related to coin design, $\mathbf{x}_{it}$ it also incorporates categorical token-type dummies (fiat-backed, crypto-backed, or algorithmic) rather than individual coin fixed effects, which can cause estimation biases in non-linear panels with rare events.

The function $G(\cdot)$ represents a cumulative distribution function (CDF) that constrains the predicted probabilities between 0 and 1. Exactly, we estimate two primary specifications:

Logit model that assumes a logistic distribution of the error term, with the link function:

$$G(z) = \Lambda(z) = e^z/(1 + e^z), \quad (5)$$

Probit model that assumes a standard normal distribution of the error term, with:

$$G(z) = \Phi(z), \quad (6)$$

where $\Phi(z)$ is the cumulative distribution function (CDF) of the standard normal distribution.

Parameters β are estimated by maximum likelihood. The coefficients are reported as odds ratios (OR) for Logit, calculated as $OR_j = exp(\beta_j \cdot x_j)$, so that a one-standard-deviation increase in x_j multiplies the odds of a depeg by OR_j. Their statistical significance is evaluated based on robust standard errors, and we report two-sided p-values. The linear predictor $x_{it}'\beta$ becomes log-odds (Logit) or latent index (Probit) of depeg risk at desired covariate values, allowing marginal effects to be calculated. Endogeneity and simultaneity bias are critical aspects of estimating crisis probabilities. According to Berg and Pattillo (1999) and Bussière and Fratzscher (2006), relying on strictly contemporaneous variables to explain a crisis creates an endogeneity problem: by definition, the crisis distorts market metrics greatly (e.g., a depeg leads directly to the huge increase in trading volume or the precipitous decrease in Total Value Locked, so that the cause and effect become impossible to disentangle).

Despite having economic interpretability, parametric models like Logit and Probit can find it challenging to model complex, non-linear interactions (Florez-Lopez 2007). To overcome these shortcomings and to account for feature interactions without specifying them in advance, we extend our research framework with the use of two tree-based ensemble machine learning algorithms–Random Forest (Breiman 2001) and eXtreme Gradient Boosting (XGBoost) (Chen and Guestrin 2016). Both methods utilize the identical covariate vector and the 7-day-ahead binary target as the econometric

specification. Additionally, to ensure direct comparability across all four models, the features are standardized exclusively on the training set, and the same temporal train-test split is applied.

3.4.1 Random Forest

It is based on a model of decision trees, which is implemented by averaging the decisions of multiple trees to increase their accuracy and to decrease the variance. The model was built from B decision trees. Different bootstrap samples corresponding to the training data have been fitted on each tree. Moreover, at each node split, the algorithm considers only a random subset of m predictors (usually $m \approx \sqrt{p}$ for classification tasks) rather than the complete set of features. This twofold injection of randomness serves to decorrelate the individual trees, in addition to contributing to avoid overfitting, a recurrent issue in financial forecasting (Shrivastav and Kumar 2022). In case of binary classification, we consider each tree as returning a class proportion as the default, or as a simple 0/1 vote in terminal leaf. The overall ensemble then estimates a final probability by averaging of the output of each of the trees:

$$\hat{p}_{it}^{\mathrm{RF}} = \frac{1}{B}\sum_{b=1}^{B} \hat{f}_b\,(\mathbf{x}_{it}), \tag{7}$$

where $\hat{\hat{f}}_b$ represents the b-th decision tree. Tree splits are typically chosen to max the decrease in Gini impurity or cross-entropy.

3.4.2 XGBoost (Gradient Boosting)

XGBoost is a state-of-the-art gradient boosting algorithm well-documented for its efficiency and predictive ability for financial distress prediction (Carmona, Climent, and Momparler 2019). Unlike Random Forest, which builds trees independently, XGBoost fits an additive model of regression trees sequentially.

Each new tree is specifically designed to correct the prediction errors (residuals) made by the previous trees.

The predicted log-odds (or score) for observation i at round t is given by:

$$F(\mathbf{x}_{it}) = \sum_{k=1}^{K} f_k(\mathbf{x}_{it}), \quad f_k \in \mathcal{F}, \tag{8}$$

where each f_k is a regression tree and (F) is the functional space of all possible trees. Trees are added iteratively to minimize a regularized objective function. For binary classification, the primary loss is the negative log-likelihood (binary cross-entropy). Crucially, the objective at each boosting step includes an L2 penalty on the leaf weights to control model complexity:

$$\mathcal{L} = \sum_{i} \ell(y_i, \hat{y}_i) + \sum_{k} \Omega(f_k), \quad \Omega(f) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2. \tag{9}$$

Here, (*T*) represents the number of leaves in the tree, w_j are the leaf weights, and γ and λ are regularization parameters. Such built-in regularization serves to avoid the widely-cited risk of overfitting (Chen and Guestrin 2016). The final predicted probability of a depeg event is obtained by applying the sigmoid function to the aggregated score:

$$\hat{p}_{it}^{\mathrm{XGB}} = \sigma\big(F(\mathbf{x}_{it})\big) = \frac{1}{1 + e^{-F(\mathbf{x}_{it})}}. \tag{10}$$

In practice, both explicit regularization and early stopping techniques are applied to limit overfitting and improve the model's generalization capabilities on the temporal test set.

3.5 Model Evaluation and Robustness

Using a unified testing framework, we assess the predictive performance of Logit, Probit, Random Forest, and XGBoost models. In Chapter 5, we present all results and visualizations.

Performance of models is evaluated on the strictly time-ordered test set comprising the most recent 20% of observations, as mentioned in Section 3.2. All models follow the same temporal split and feature scaling to make sure their comparisons are fair. Because stablecoin depegs are infrequent events, there is a significant imbalance in the dataset. Under these conditions, the overall accuracy can be deceptive; a model could become highly accurate just by not predicting any crisis. Therefore we fall back on evaluation metrics which are created for rare financial events (Xu, Li, and Wu 2024; Barthélémy, Gautier, and Rondeau 2024).

We evaluate discrimination through ROC-AUC, indicating the trade-off of true vs. false positive rates (Drehmann and Juselius 2014) at all thresholds. However, because ROC-AUC is potentially too optimistic in cases of imbalanced data, we also present PR-AUC, which measures performance for the minority class (depegs). To preserve temporal dependence, we compute 95% confidence intervals from the block bootstrap for both metrics.

In practice, the early warning system should have a probability threshold to trigger crisis warning. We describe the F1-train and Youden cut-offs as a reference benchmark and suggest monitoring policies on a tail-based basis (Top 1%, Top 5%) or as the target-alerts-per-day rule under consideration, with the rule being more in line with the rare-event early-warning process to separate ranking and detecting (Kaminsky, Lizondo, and Reinhart 1998; Bussière and Fratzscher 2006). For each candidate cut-off we produce the confusion matrix and report Precision, Recall, F1-score, Matthews Correlation Coefficient (MCC) and balanced accuracy. Recall and F1-score are especially relevant for the prediction of financial distress,

as missing a true crisis (Type I error) is generally more costly than transmitting a false alarm (Type II error).

We also display Precision and Recall at the different fixed points (0.1, 0.2, 0.3, 0.5), to show the sensitivity / false alarm trade-off. Additionally, we analyze the lift curve to evaluate how well the model concentrates the depeg events among high risk days.

We assess parametric models vs predictive machine learning systems. Although Logit and Probit are interpretable, the use of tree-based models like Random Forest and XGBoost can show sophisticated non-linear relationships (Barboza, Kimura, and Altman 2017). Moreover, we cross-validate the SHAP values to confirm that the main predictors detected in parametric models still make sense for flexible machine learning specifications.

Since the most reliable early warning system should give realistic probabilities, we calibrate using both the Expected Calibration Error (ECE) and Brier score (Brier 1950). To establish systematic over- or under-prediction, reliability curves are utilized. To enhance probability estimates of the predicted XGBoost, we used Platt scaling (Platt 1999), where predicted probabilities are refined, but the ranking ability of the model is unchanged. Calibrated probabilities are an important requirement for risk management and regulatory applications.

Finally, we perform robustness checks. First, as a robustness check (Appendix B), we re-estimate the main specification with 90- and 180-day burn-ins per token to remove launch-phase volatility; the main specification uses no burn-in. Second, we evaluate the contribution of variable groups by removing them sequentially. Third, to verify that predicted risk increases before the crash, we examine event windows around depegs, confirming the early-warning interpretation.

Overall, this framework provides a transparent and reproducible evaluation of the early warning system under realistic risk management constraints.

Hypotheses H1–H3 are tested with Logit and Probit models within a binary-outcome framework. In line with conventional econometric methodology, each hypothesis is investigated based on the sign and statistical significance of the estimated coefficients (β), based on robust standard errors and respective p-values (Wooldridge 2010). For H1: the liquidity variables (*dex_tvl_usd, dex_volume_30d, log_vol_30d, dex_share_max,* and *ln_binance_volume_30d* as CEX-depth support) are expected to have negative coefficients, consistent with the hypothesis that lower liquidity leads to higher depeg probability. For H2, *supply_pct_change_7d* is expected to have a positive coefficient, implying rapid issuance spikes (net positive 7-day supply growth) precede periods of increased depeg risk. As for H3, *btc_return_7d* and *eth_return_7d* are expected to have negative coefficients entering the model due to adverse market conditions. When the estimated coefficient has the predicted sign and is statistically different from zero at conventional significance levels, a hypothesis is supported. As Logit and Probit represent non-linear models, we interpret effect sizes via odds ratios (Logit) and marginal effects to measure economic magnitude (Bussière and Fratzscher 2006). Hypothesis testing is conducted in-sample, whereas predictive validity is evaluated out-of-sample. Machine learning models are used only to corroborate predictive performance and are not employed for formal statistical inference, consistent with the early warning system literature (Berg and Pattillo 1999; Comelli 2014; Frankel and Rose 1996).

Chapter 4

DATA

The dataset constructed for the assessment and evaluation of the early warning system. The following section is organized sequentially, following the standard empirical structure established in recent applied EWS and stablecoin literature (Cintra and Holloway 2023; Lee, Chiu, and Hsieh 2025). We start by defining data sources the observation period and the unit of analysis. We then explain the sample selection and construction of the model-ready panel. Secondly, we present the raw sources and transformations used to engineer the codebook variables. With this in mind, we present the descriptive statistics which demonstrates selected stablecoins` historical behaviors, liquidity conditions, and distributional properties preceding the empirical estimation.

4.1 Data Sources

To determine the causes of depeg in stablecoins, we create an extensive daily panel dataset (January 1, 2018–February 11, 2026). Our unit of observation is a token-day–an observation corresponds to a stablecoin's features on a specific calendar date. The sample comprises several large USD-pegged stablecoins representing various stabilization mechanisms (fiat-backed stablecoins, crypto-backed stablecoins, and algorithmic stablecoins) from which we can easily identify and analyze. As these assets were issued at different points in time, the dataset is built up as an unbalanced panel. One token's observation window begins only on the first date on which all core variables are available. Such an unbalanced structure is often made visible in cross-asset early warning systems and provides a more efficient utilization of information available to an investor (Lee, Chiu, and Hsieh 2025).

All the data are automatically collected using an Extract, Transform, Load (ETL) pipeline and is stored in a centralized Google BigQuery environment where these

raw inputs are standardized and turned into features needed for modeling. We rely on four providers of data to cover different sides of market stress, liquidity, and supply. We have general data of the market by CoinGecko, which is one of the most popular analytics on cryptocurrency but it contains price, market capitalization and trading volume data based on centralized/decentralized exchanges. Daily USD-denominated price for price_usd, on CoinGecko is used as the reference price for our analysis. The absolute deviation from the peg (*abs_dev*) is computed on this series. We augment CoinGecko with total trading volume & market capitalization data along with USD value of circulating supply (*supply_usd*) in order to produce the rolling 30-day trading volume (*log_vol_30d*). btc_return_7d and eth_return_7d refer to short term market return controls derived from daily prices for Bitcoin and Ethereum extracted from the same source in 7-day returns. These traits reflect underlying crypto market trends. DefiLlama, its analytics ecosystem, allows us to observe supply-side responses (incoming and outgoing supply) to measure token circulation and protocol-level blockchains. These data allow us to estimate the percentage of circulating supply that changes over 7 days (*supply_pct_change_7d*), considering the issuance or redemption shocks in addition to the ownership percentage of each token's supply deployed on Ethereum (*dominance_eth*), this proxy that indicates the degree of concentration of the chain and the exposure. On-chain liquidity is obtained using The Graph; It is a decentralized blockchain indexing protocol that allows structure smart-contract data based on its subgraphs as part of one's own network. Via subgraphs for major decentralized exchanges such as Curve Finance and Uniswap we obtain measures of total value locked (*dex_tvl_usd*) as well as rolling 30-day trading volume (dex_volume_30d) and the maximum pool share (*dex_share_max*). These factors reflect both liquidity depth and fragmentation dangers worthy of attention (Lyons and Viswanath-Natraj 2023).

Finally, we include Binance data to show centralized exchange trading activity. Based on daily OHLCV data, we build a 30-day rolling sum of USD trading

volume across all stablecoin pairs on Binance (*binance_volume_30d*). We use its log transformation, *ln_binance_volume_30d* = ln(1 + *binance_volume_30d*), to reduce skewness while preserving all available observations. This variable captures overall CEX market depth and trading intensity and is included to assess whether changes in centralized liquidity conditions increase the probability of a stablecoin depeg.

4.2 Sample Selection

To ensure that all the elements in the stablecoin ecosystem can be examined together, our research universe consists of nine major USD-pegged tokens that cover different design mechanisms: fiat, crypto-collateralized, and algorithmic/synthetic. These tokens were included as they are the market capitalizations with the highest trading volumes and most diversity around stabilization mechanisms within decentralized finance (Kosse et al. 2023; Liao and Caramichael 2022).

The first kind includes fiat-backed stablecoins that maintain a one-to-one peg based on off-chain traditional financial assets, such as cash, bank deposits, and U.S. Treasuries. Issuers make their income from the yield derived from these traditional reserves (d'Avernas, Maurin, and Vandeweyer 2022). These include Tether (USDT), USD Coin (USDC), and PayPal USD (PYUSD). By market capitalization, USDT is the oldest and largest stablecoin, serving as the key quote currency of centralized crypto exchanges with its backing in cash, Treasuries, and corporate debt (Lyons and Viswanath-Natraj 2023). USDC, managed by the Centre consortium, is known for its tight short-term U.S. Treasuries backing and has strong transparency, though it suffered a significant depeg in March 2023 on account of reserve exposure to the failed Silicon Valley Bank (Anadu et al. 2024). PYUSD is a relatively new fiat-backed token launched by PayPal aimed at mainstream payments and remittances, in order to connect traditional finance with blockchain networks.

The second category includes crypto-collateralized stablecoins (those which exist fully on-chain) in that these tokens over-collateralize in order to withstand their underlying crypto assets' price volatility (Klages-Mundt and Minca 2021). This group comprises DAI, Liquity USD (LUSD), and Magic Internet Money (MIM). DAI, which is managed by MakerDAO, is the leading decentralized stablecoin and users deposit and lock collateral into smart contract vaults used for tokens and make a mint from them and the platform is monetized through charges for stabilizing (d'Avernas, Maurin, and Vandeweyer 2022). In LUSD, the minimum collateral is 110 percent funded by Ethereum and it charges a single borrowing fee as opposed to interest rate changeable (Anadu et al. 2024). This space is how MIM is innovating in the category, i.e., users deposit interest-bearing tokens as collateral and are able to borrow stablecoins that have been invested, but will earn yield elsewhere in the ecosystem.

The third category consists of algorithmic, hybrid, and synthetic stablecoins, which try to achieve stability via the mechanism of arbitrage and fractional reserves/financial derivatives, as opposed to merely supporting assets (Clements 2021). This includes FRAX, TerraClassicUSD (USTC), and Ethena USDe. FRAX is a mixed fractional-algorithmic model, supported by USDC and partially stabilized by its native token, that generates payment by deploying collateral in any of the different protocols (Kazemian, Iyer, and Huan 2020). USTC was formerly known as UST as a pure algorithmic stablecoin completely backed by an endogenous token that was managed by mint-and-burn arbitrage and is included in our sample as the primary example of a structural failure and downward spiral which took place from May 2022 onwards (Liu, Makarov, and Schoar 2023; Uhlig 2022). USDe is a nascent synthetic currency that keeps its peg in perpetual futures markets with a delta-neutral hedging mechanism instead of over-collateralizing, making money from staking yields and derivative funding rates. Although the descriptive behaviors of all nine tokens were observed, PYUSD and USDe were not included in the final estimation sample. These two

of the tokens suffered only two and one instances of depeg respectively and so not enough variation was found for all binary classification models present in this research. As a result the final estimation sample consists of seven tokens with three or more depeg episodes. In each episode, we record the duration, measured in consecutive depeg days beginning with the start, alongside the direction, showing if the price was below or above the $1 peg on the first day. This classification method corresponds with S&P Global's approach (2023). The episode number, duration, and direction statistics for the stablecoins in our universe are also shown in Table 3.

The obtained sample will then be utilized in training and test datasets. The dataset comprises 669 crisis observations, with a crisis share of around 4.7%. Due to the extreme rarity of depeg events, we used a strictly temporal evaluation design. The training sample includes the earliest 80% of data, the test sample includes the most recent 20%. Such an approach helps to account for look-ahead bias (data leakage), and it is a true representation of the actual state of forecasting in real time (Barthélémy, Gautier, and Rondeau 2024).

Table 3. Train–Test Split Summary.

Item	Value
Unit of observation	Token-day
Tokens (final sample)	7
Total observations	14,250
Crisis observations	669
Crisis share	4.7%
Time range	2018-01-01 to 2026-02-11
Train period	2018-01-01 to 2024-11-25
Test period	2024-11-25 to 2026-02-11
Split rule	First 80% train, last 20% test

All transformations and standardization are computed using training data only and then applied to the test set. No tuning is performed on the test data. To account for serial dependence, we compute confidence intervals using block bootstrap with a seven-day window.

Table 4. Depeg episodes per token.

Token	Collateral	Episodes	Mean duration (days)	Max duration (days)	Below $1	Above $1
USDT	fiat	49	2.3	28	17	32
USDC	fiat	16	2.8	9	4	12
DAI	crypto	18	2.2	8	7	11
LUSD	crypto	36	4.0	55	2	34
MIM	crypto	17	1.1	3	12	5
FRAX	algo	3	1.7	2	2	1
USTC	algo	3	1.3*	2*	3	0
PYUSD	fiat	2	1.5	2	1	1
USDe	algo	1	1.0	1	0	1

*Note on USTC (Terra): Data are trimmed at 2022-05-10 (the first "point of no return" day).

We keep the run-up and the collapse day but exclude the post-collapse "zombie" period, so the last episode has a duration of 1 day. This ensures duration statistics remain comparable across tokens and follows the standard econometric practice of censoring tokens after a structural collapse (Liu, Makarov, and Schoar 2023).

The historical price behavior and the model's reaction to depeg events across the research universe are visualized in Figure 2. The graph plots the stablecoin prices in USD against their respective dynamic threshold bands (α_t).

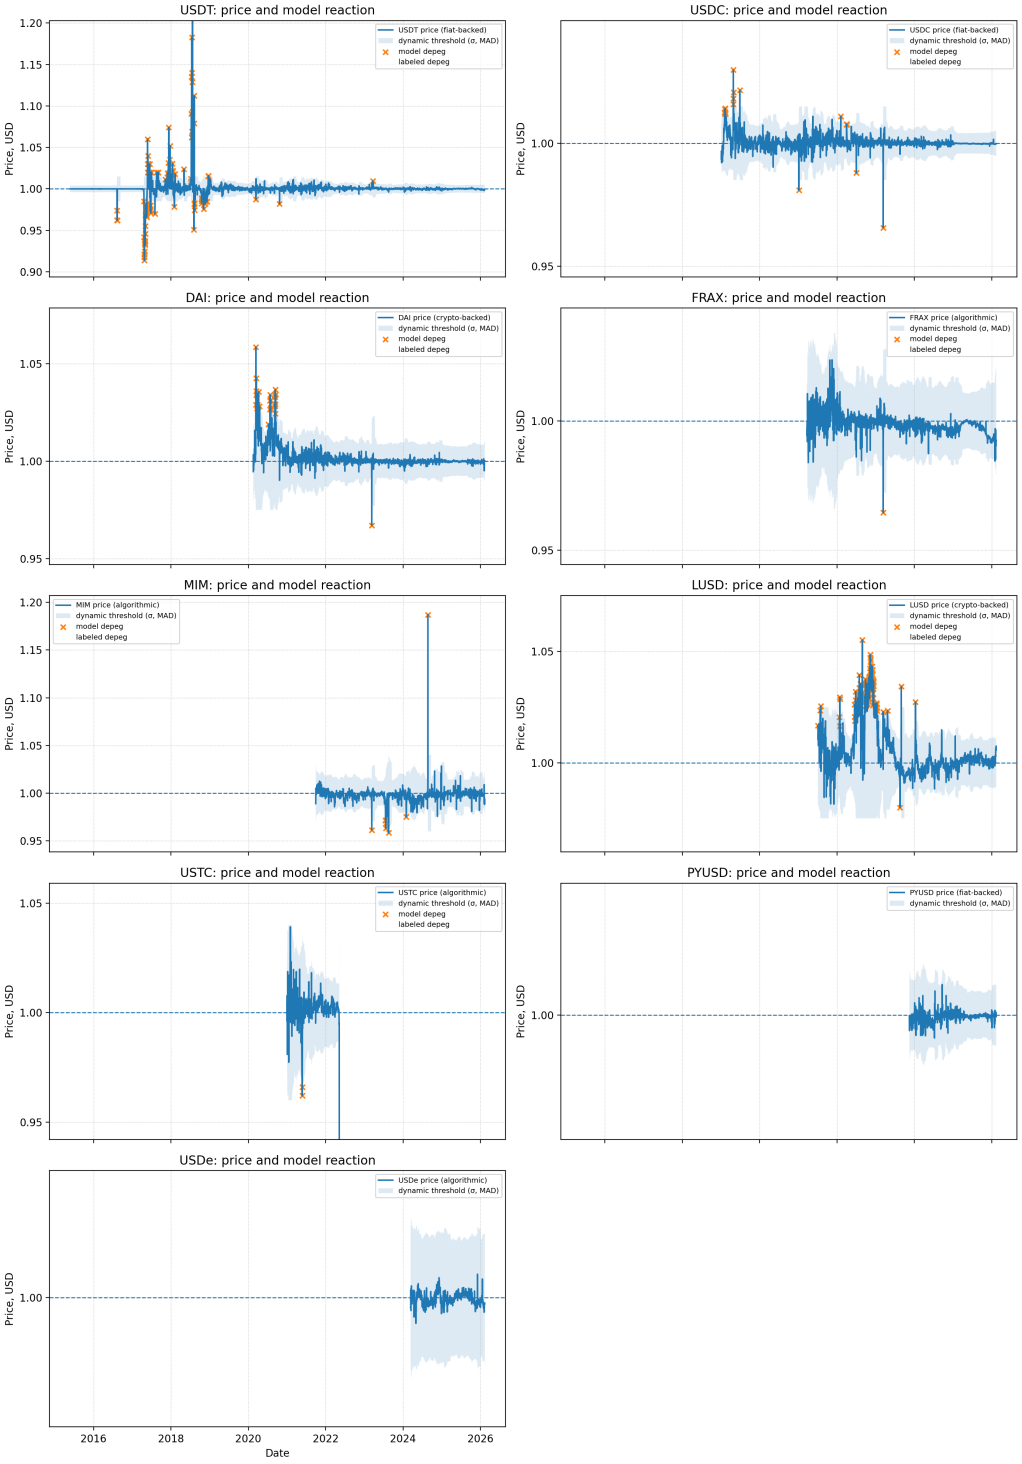


Figure 2. Historical stablecoins price and the model's reaction to depeg events.

Finally, Figure 3 shows the average duration in days of depeg episodes for each token. Figure 4 displays the direction of these episodes, comparing instances of the price falling below versus rising above the $1 peg. The values of USTC in both curves are capped to maintain a stable visual scale.

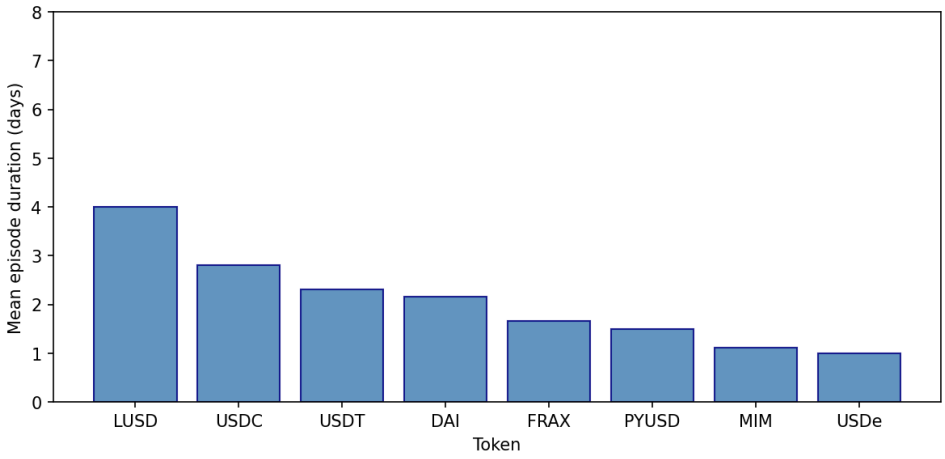


Figure 3. Mean depeg episode duration by token (days).

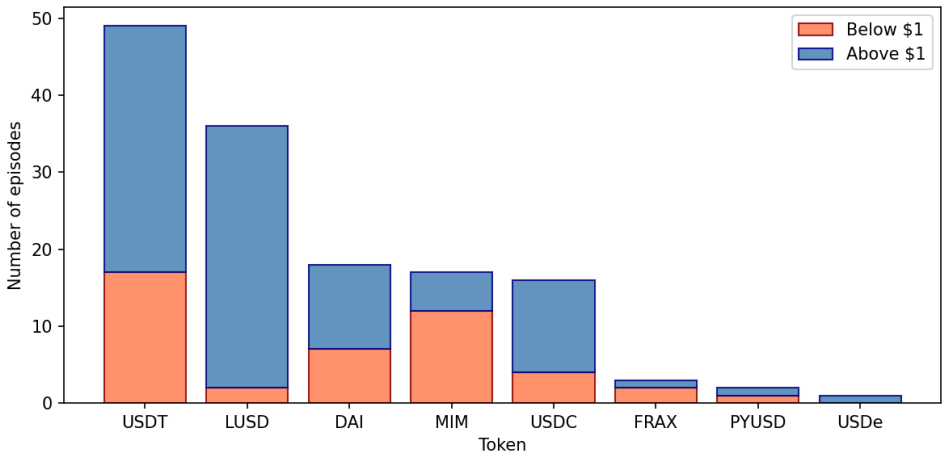


Figure 4. Depeg direction: episodes below vs above $1 by token.

Note for Figures 3-4: USTC values are capped due to the structural collapse of TerraUSD.

4.3 Raw Data Analysis

We begin with daily USD price series we got from CoinGecko as our principal prices. These series serve as the base for the core market variables of the empirical framework.

In Figure 5 we reflect the daily price of all stablecoins available in the research universe for the entire sample. Most stablecoins remain on the one dollar peg thus the deviations from the norm are extremely rare and small and short term. USDT volatility, for instance, is significantly higher during the first few years because the market environment becomes thin. The USDC, DAI, and FRAX are all highly peg-linked with their desired pegs and mild anomalies caused by acute stresses are present however.

The structural break in the panel is the impact of Terra (USTC) collapse in May 2022. To illustrate the size of the break, the chart is intentionally not severed prior to the 2022 collapse, the price diverging sufficiently from parity and never settling down sharply. To avoid bias, we will truncate the USTC series within the estimation sample from May 10, 2022, well before the terminal stage of the collapse (Liu, Makarov, and Schoar 2023).

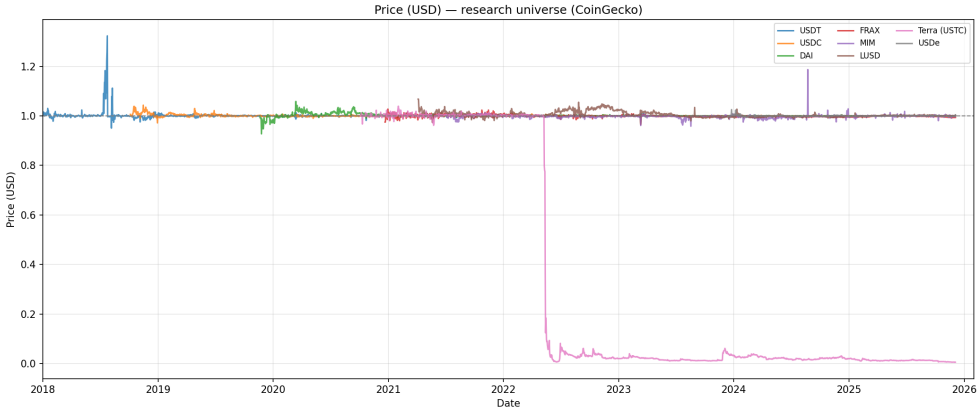


Figure 5. Daily USD price by stablecoin. Source: CoinGecko.

So now we look at supply dynamics which also serve to explain issuance, redemption and overall adoption per token. For daily circulating supply, we use USD data from DefiLlama, a widely used source that tracks token activity on multiple blockchains and DeFi protocols. We derive three variables from this data: total circulating supply (*supply_usd*), 7-day percentage change in supply (supply_pct_change_7d), and the share of supply held on the Ethereum network (*dominance_eth*), this captures the concentration of the network.

Figure 6 illustrates how the circulating supply in USD evolves over time for each stablecoin. The plot reflects the exponential growth in demand for major fiat-backed tokens after 2020 and shows changes in demand for stablecoins in the cryptocurrency system as a whole. Conversely, the figure also illustrates periods of sharp supply contraction. As recently stressed in the literature, a reduction such as this usually indicates mass market redemptions with its underlying market stress (Eichengreen, Nguyen, and Viswanath-Natraj 2025).

The most visible structural change is the resulting structural dislocation that followed the disastrous fall of Terra (USTC) in May 2022. Following this event, the market became highly skewed toward fiat- and over-collateralized crypto-backed tokens as investors fled the system to more secure investments (Anadu et al. 2024). As a result, algorithmic stablecoins have either entered the market later or were much smaller in size. Such supply patterns are of particular interest for our early warning system, as abrupt supply contractions of a circulating type often reflect waning confidence from investors ahead of depeg materializing (Liu, Makarov, and Schoar 2023).

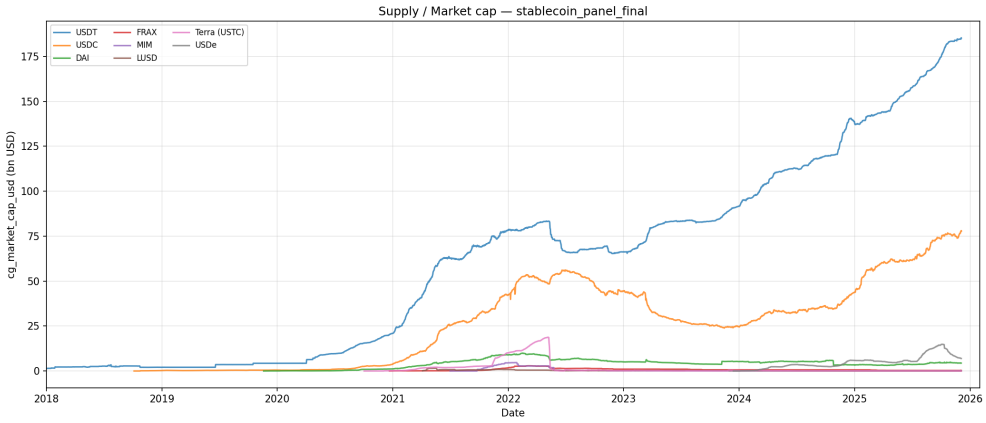


Figure 6. Daily circulating supply (USD) by stablecoin. Source: DefiLlama.

We use daily prices of Bitcoin and Ethereum from CoinGecko to compute 7-day logarithmic returns (*btc_return_7d* and *eth_return_7d*), which are proxies for overall crypto market momentum. These are added as the required control variables so that the early warning algorithm captures the stablecoin-specific stress and does not react solely to broad market fluctuations.

Daily price levels corresponding with Bitcoin and Ethereum are shown as illustrated in Figure 7 for comparison using two different vertical scales. They both are the basic assets in the cryptocurrency ecosystem and have had enormous historical volatility. As the recent empirical evidence reflects, large market cycles such as the bull and bear of 2021–2022 and their resets in 2024–2025 are often associated with periods of amplified stress in stablecoin markets. In particular, steep falls and extreme volatility in Bitcoin and Ethereum frequently result in flight-to-safety dynamics, margin calls, and mass redemptions of stablecoins (Eichengreen, Nguyen, and Viswanath-Natraj 2025; Lee, Chiu, and Hsieh 2025).

Hence including these benchmarked returns, the predictive model can discern between system-wide market panics and the idiosyncratic depeg risk of an individual stablecoin.

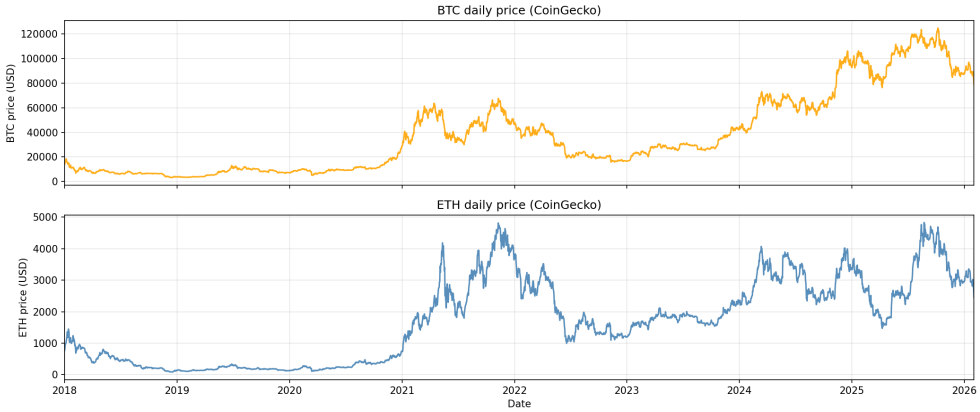


Figure 7. BTC and ETH daily prices (separate scales). Source: CoinGecko.

To capture the decentralized exchange (DEX) liquidity, we gather the daily balances of tokens (in USD) and volumes of trading with The Graph subgraphs for major protocols (for example Curve Finance and Uniswap on the Ethereum mainnet). We then aggregate the above raw data into three key liquidity covariates defined as the total value locked across related pools (dex_tvl_usd), 30-day rolling DEX trading volume (dex_volume_30d), and the maximum token share within a single pool (dex_share_max). In order to measure on-chain liquidity depth and pool concentration, critical structural imbalances on DEX pools often occur prior to or even coincide with depeg events (Cintra and Holloway 2023), thus these variables are critical to quantifying the depeg process. We remove LP tokens (for instance: 3CRV) during the transformation and map the historical UST observations to USTC for consistency to prevent double counting.

The Curve Total Value Locked (TVL) for pool and token is plotted in Figure 8 over the panel period. The visualization illustrates well the structural rupture in decentralized liquidity after Terra and the recent crypto bear market in 2022. Following these systemic shocks, the aggregate Curve TVL and pool-level liquidity collapsed. Accordingly, the figure shows the shift towards post-2022 market scenarios in which many stablecoins are operated under significantly

lower DEX depths (Liu, Makarov, and Schoar 2023). The data are summarised in the table below, showing the well-defined liquidity pools underpinning these metrics.

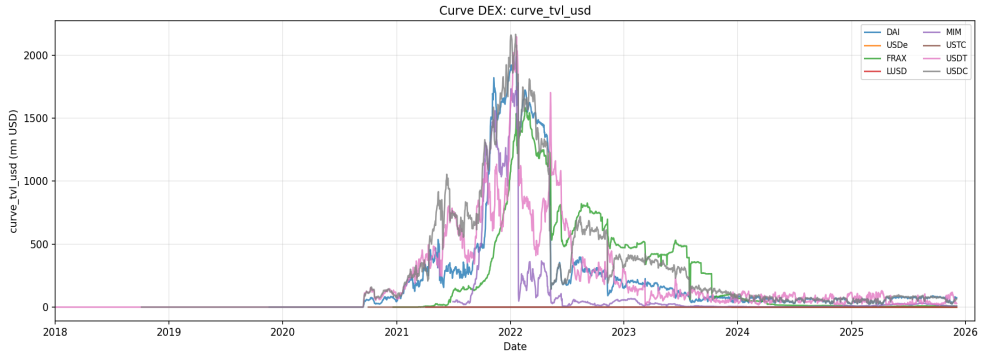


Figure 8. Curve TVL by pool and token. Source: The Graph.

Table 5. Used DEX pools by token pairs.

Pool	DEX	Tokens
3pool	Curve	DAI, USDC, USDT
FRAX-USDe	Curve	USDe, FRAX
PYUSD-USDC	Curve	PYUSD, USDC
MIM-3pool	Curve	MIM, 3CRV
LUSD-3pool	Curve	LUSD, 3CRV
FRAX-USDC	Curve	FRAX, USDC
UST-3pool	Curve	UST, DAI, USDC, USDT
USDe-USDT	Uniswap v4	USDe, USDT

4.4 Descriptive Statistics

We start with a description of the general organization of the empirical panel and the incidence of depeg events within the sample. Table 3 summarizes the scale of the dataset.

Table 6 details the summary statistics for the continuous covariates included in the model. On average, the absolute deviation from the peg (*abs_dev*) is minimal (mean = 0.003). This confirms that, under normal market conditions, stablecoins successfully maintain parity and are highly resilient to minor fluctuations (Ahmed, Aldasoro, and Duley 2024). At the same time, the wide dispersion observed in supply changes (*supply_pct_change_7d*) and broad crypto market returns (*btc_return_7d, eth_return_7d)* reflects the presence of severe liquidity stress and large systemic shocks within the sample (Lee, Chiu, and Hsieh 2025).

Table 6. Summary statistics for model covariates.

Variable	Mean	Std	p25	p50	p75
abs_dev	0.003	0.008	0.0003	0.001	0.003
log_vol_30d	22.20	4.03	19.12	22.10	25.76
*dex_tvl_usd**	170.87	336.81	2.71	50.94	129.36
*dex_volume_30d**	1,320.60	2,339.23	0.004	94.01	2,087.32
dex_share_max	0.353	0.276	0.083	0.355	0.542
supply_pct_change_7d	0.147	6.046	−0.006	0.000	0.019
btc_return_7d	0.010	0.086	−0.035	0.004	0.050
eth_return_7d	0.014	0.124	−0.053	0.003	0.070

Source: Model-ready panel (7 tokens, USTC trimmed at 2022-05-10).

Note: Liquidity variables (*dex_tvl_usd, dex_volume_30d*) are expressed in USD millions.

To further understand the behavior of our predictors leading up to a crisis, Table 7 compares the unconditional mean values of the model covariates between tranquil (non-crisis) and pre-crisis token-days. Pre-crisis observations correspond to periods where a depeg episode begins within the subsequent seven days (*y_depeg_7d*=1). The table reports the group means and their absolute differences, with statistical significance evaluated using Welch's unequal-variance t-tests.

Table 7. Mean values of covariates: non-crisis vs crisis token-days.

Variable	Non-crisis mean	Crisis mean	Difference
abs_dev	0.0028	0.0292	0.0263***
log_vol_30d	22.52	20.34	-2.17***
dex_tvl_usd (millions)	183.7	69.6	-114.1***
dex_volume_30d (millions)	1,419.1	626.8	-792.3***
dex_share_max	0.378	0.225	-0.154***
supply_pct_change_7d	0.12	5.85	5.73
btc_return_7d	0.008	-0.033	-0.041***
eth_return_7d	0.010	-0.029	-0.039**

*** p<0.01, ** p<0.05 (Welch t-test).

The empirical evidence shows distinct structural differences between tranquil and pre-crisis periods. The impact of *supply_pct_change_7d* must be treated with caution. The average difference, though large, results from a highly skewed distribution, having rare but great positive spikes. Most of the observations are centred around zero, and the statistical significance is not robust to trimming out

extreme values, highlighting that the result indicates episodic shocks, not a systematic difference between groups. It is not surprising that the absolute deviation from the peg (*abs_dev*) is substantially higher just before a depeg, as expected, reflecting increasing price stress and a deteriorating market condition.

Liquidity compression is also well documented in pre-crisis periods. Total decentralized liquidity (*dex_tvl_usd*) and rolling trading activity (*log_vol_30d, dex_volume_30d*) drops dramatically. This sharp drop in maximum pool share (*dex_share_max*) also suggests increasing liquidity fragmentation and flagging depth in primary automated market maker (AMM) pools, which are both tightly correlated with impending stablecoin failures (Cintra and Holloway 2023).

Lastly, larger fluctuations in crypto markets clearly manifest ahead of depeg episodes. Bitcoin and Ethereum returns (*btc_return_7d, eth_return_7d*) are weakly positive in the regular time but they consistently turn negative prior to crises. It is consistent with recent empirical findings that the broad market decline and increase in volatility in benchmark cryptocurrencies further push up the massive redemption pressures, flight-to-safety, and instability in stablecoin pegs (Eichengreen, Nguyen, and Viswanath-Natraj 2025; Lee, Chiu, and Hsieh 2025).

Figure 9 presents the rolling 90-day share of crisis observations (*y_depeg_7d*) over its length. The series shows considerable temporal variation which suggests that the depeg intensity varies in the sample and is not relatively constant. Extreme periods of crisis are interspersed with relatively quiet intervals, suggesting that instability tends to cluster around certain segments of the market. These patterns imply that depeg risk is regime-dependent and varies in time rather than being uniformly distributed (Gadzinski et al. 2023). The Early Warning System is approximated on a panel that spans the tranquil and turbulent periods, so this variability is important to consider when interpreting model performance and out-of-sample stability (Barthélémy, Gautier, and Rondeau 2024).

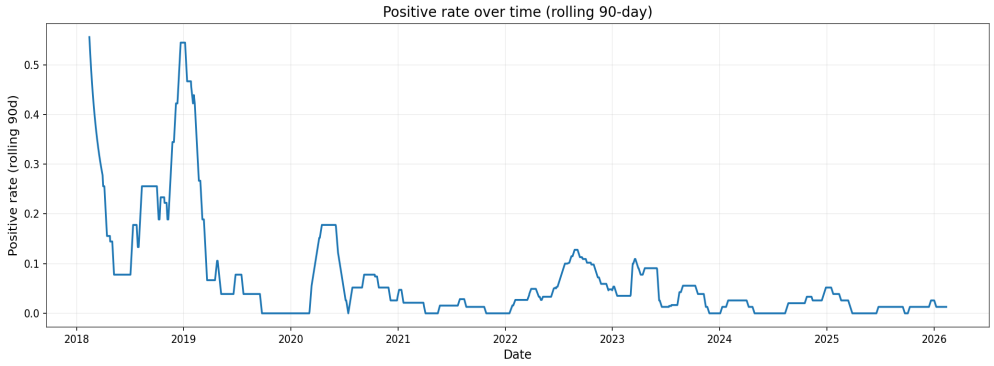


Figure 9. Rolling 90-day positive rate (y_depeg_7d).

The results are shown in Figure 10, suggesting that the pairwise correlations between predictive variables are, in general, moderate. This confirms the lack of substantial multicollinearity that, otherwise, would impair the efficiency and stability of the Logit estimates. As would theoretically be expected, there is a positive correlation between total decentralized liquidity (*dex_tvl_usd*) and decentralized trading volume (*dex_volume_30d*), which represents their common affiliation with broader on-chain liquidity conditions. However, the magnitude of this correlation remains well within acceptable econometric bounds. Thus, both variables are retained within the model, due to capturing distinct and complementary dimensions of market depth and trading activity (Cintra and Holloway 2023).

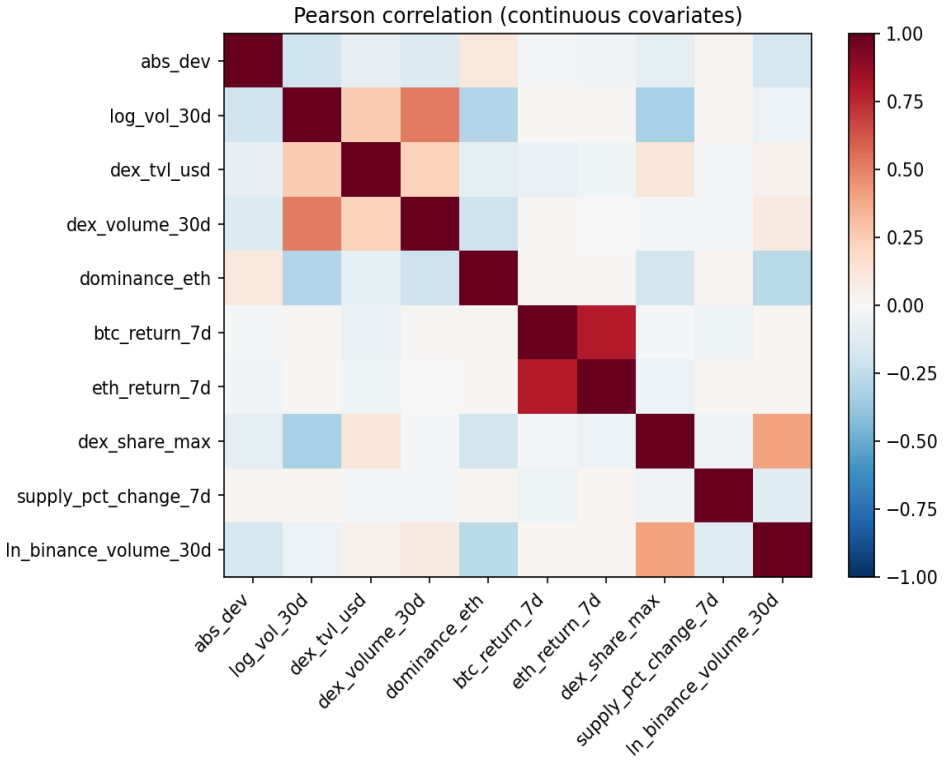


Figure 10. Correlation matrix of standardized predictors.

Based on the descriptive statistics, this implies that whereas most major stablecoins are generally stable, they experience distinct episodes of stress that vary significantly by design and liquidity conditions. The variation provides the empirical basis for testing our EWS.

Chapter 5

ESTIMATION RESULTS

This section introduces the experimental framework used to study the out-of-sample predictive performance of the model. As described in Chapter 3, our baseline parametric models consist of Logit and Probit regressions. Moreover, we evaluate Random Forest and XGBoost to assess whether more flexible, non-linear machine learning approaches generate gains in predictive accuracy (Barthélémy, Gautier, and Rondeau 2024; Xu, Li, and Wu 2024).

We report model performance in a logical order that distinguishes between in-sample fit and out-of-sample validity. Training-stage results are presented first as a benchmarking tool and threshold selection document, then the performance of the model is evaluated using the temporal test set to assess generalization under realistic conditions. This subsection tests if the models can distinguish between crisis and non-crisis observations in the temporal test sample. In this phase, performance is scored independent of any classification threshold (Yildirim and Sanyal 2022).

Robustness checks are shown in Appendix B.

5.1 Predictive Performance

This section assesses predictive performance in two steps: one, whether the model can be adequately represented in the training data; secondly, the model generalization out of the sample when tested in a temporally separated test sample. In line with previous early warning literature, model performance is evaluated through ranking metrics (ROC-AUC and PR-AUC), while classification results using thresholds should be accounted for separately from the actual discriminatory power. Output from training-stage models provides a measure of model capacity, but thresholds in the rare-event realm can cause model performance to be overestimated, so training-stage results should not be

treated as final evidence of model quality. Table 8 shows strong in-sample performance, especially for the tree models.

Table 8. Training Performance Summary.

Model	ROC-AUC (train)	PR-AUC (train)	F1-score (train)	Optimal threshold
Logit	0.811	0.254	0.423	0.139
Probit	0.817	0.247	0.418	0.139
Random Forest	0.969	0.624	0.559	0.821
XGBoost	0.992	0.822	0.781	0.809

XGBoost and Random Forest significantly outperform both Logit and Probit in training discrimination, but this probably reflects more overfitting rather than strictly better crisis-signal extraction. The assessment of the out-of-sample test serves as the main basis for model comparison. Table 9 shows the predictive performance on the temporal holdout sample.

Table 9. Test Performance (using train-selected thresholds).

Model	ROC-AUC (test)	PR-AUC (test)	F1-score (test)	Alerts per day (test)
Logit	0.852	0.050	0.088	1.000
Probit	0.857	0.054	0.088	0.991
Random Forest	0.899	0.072	0.000	0.000
XGBoost	0.927	0.092	0.000	0.000

All models worsen somewhat compared with their training results, and the distribution shifts between estimation and evaluation periods are confirmed. Yet differences across models are informative.

The test results show that tree-based models continue to perform better in ranking out of sample, while Logit and Probit remain more stable under threshold-based evaluation.

At the same time, thresholds selected by train reveal a major operational limitation. Specifically, thresholds calibrated in-sample for Random Forest and XGBoost generate zero alerts in the test sample, producing F1 equal to zero despite strong ranking performance.

5.2 Discriminatory Power

This subsection demonstrates the out-of-sample discriminatory effectiveness of the models based on ROC-AUC and PR-AUC.

Table 10. Discriminatory power: baseline and comparison models (test set).

Model	ROC-AUC	ROC 95% CI	PR-AUC	PR 95% CI
Logit (baseline)	0.852	[0.815, 0.889]	0.050	[0.033, 0.084]
Probit	0.857	[0.818, 0.890]	0.054	[0.033, 0.093]
Random Forest	0.899	[0.875, 0.922]	0.072	[0.043, 0.127]
XGBoost	0.927	[0.911, 0.943]	0.092	[0.057, 0.172]

Switching to more flexible algorithms, machine learning models provide incremental improvements in out-of-sample discrimination, particularly in the precision-recall dimension, which is especially relevant in rare-event settings. XGBoost shows the strongest ranking performance, while Random Forest also outperforms the parametric models on both ROC-AUC and PR-AUC.

But these gains should be viewed as incremental, not structural. Logit and Probit already provide substantial discriminatory power and remain important benchmark specifications. The suggestion is not that machine learning overturns the baseline results, but that it sharpens risk ranking, especially in the extreme right tail where crisis observations concentrate.

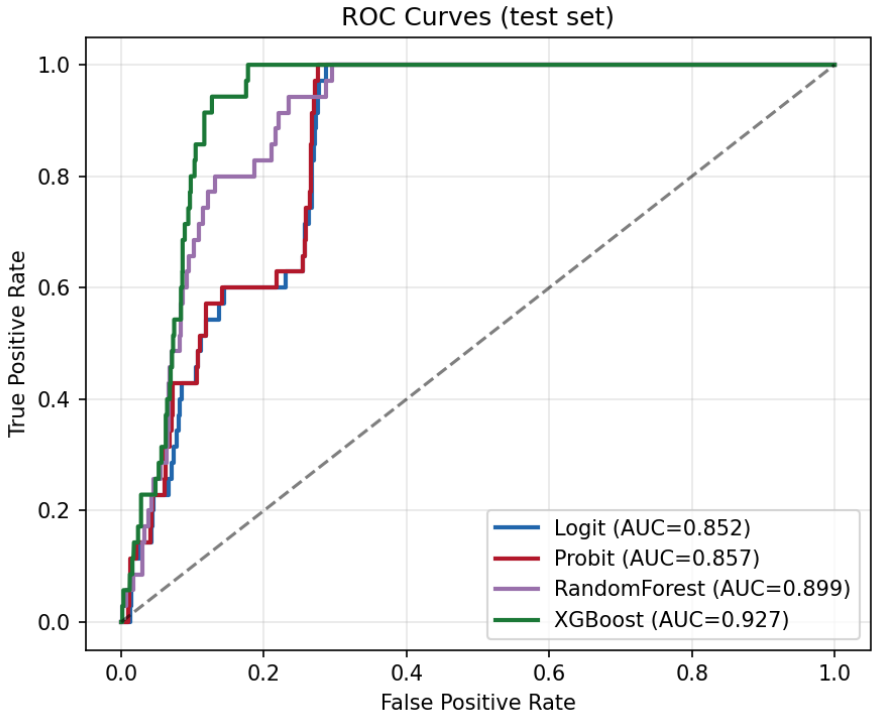


Figure 11. ROC curves for Logit, Probit, Random Forest, and XGBoost.

As per baseline predictions, the Logit and Probit data have identical curves (Comelli 2014), which confirms that the prediction outcomes are relatively robust using the underlying econometric link function that is selected (Comelli 2014). The ROC curves validate better ranking performance of tree-based models than Logit and Probit as well.

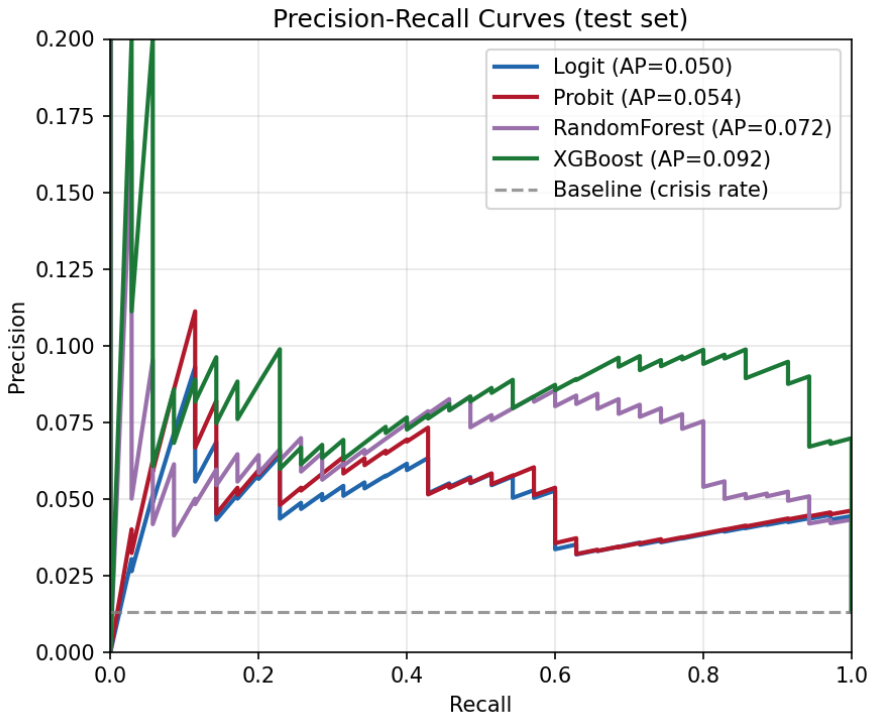


Figure 12. Precision-Recall curves for Logit, Probit, Random Forest, and XGBoost.

Such an unconditional positive rate is the horizontal "no-skill" benchmark. XGBoost performs the strongest PR-AUC among candidates. A high rating corresponds to more realistic crisis observations clustered close to the top percentiles of the projected risk distribution which is a required criterion for reducing false alarms in datasets with high level of imbalances (Xu, Li, and Wu 2024). On the other hand, the Logit and Probit models exhibit a moderate and stable PR performance that exactly agrees with the rigidity of their, yet interpretable parametric structure (Comelli 2014). These results reveal that flexible machine learning models can indeed greatly improve both tail ranking precision and classification accuracy, but the baseline econometric models already cover at least a fraction of what contributes the actual variations within crises.

The lift curve, depicted in Figure 13, indicates the concentration of real crises in the highest predicted-risk percentiles. However, in the out-of-sample test set, the top 5% of the predicted risks captures approximately 22.9% of crises under the Logit model and 28.6% under XGBoost. At the restrictive top 1% threshold, the Logit model captures zero crisis observations, whereas XGBoost successfully identifies 8.6%. This comparison points out that the tail separation is much more robust for XGBoost. These lift findings strengthen both ROC and PR evidence, indicating that model differences are particularly evident in the extreme right tail of the predicted risk. Specifically, XGBoost leads to a better concentration of crisis observations in higher risk percentiles, which is operationally relevant for rare-event monitoring.

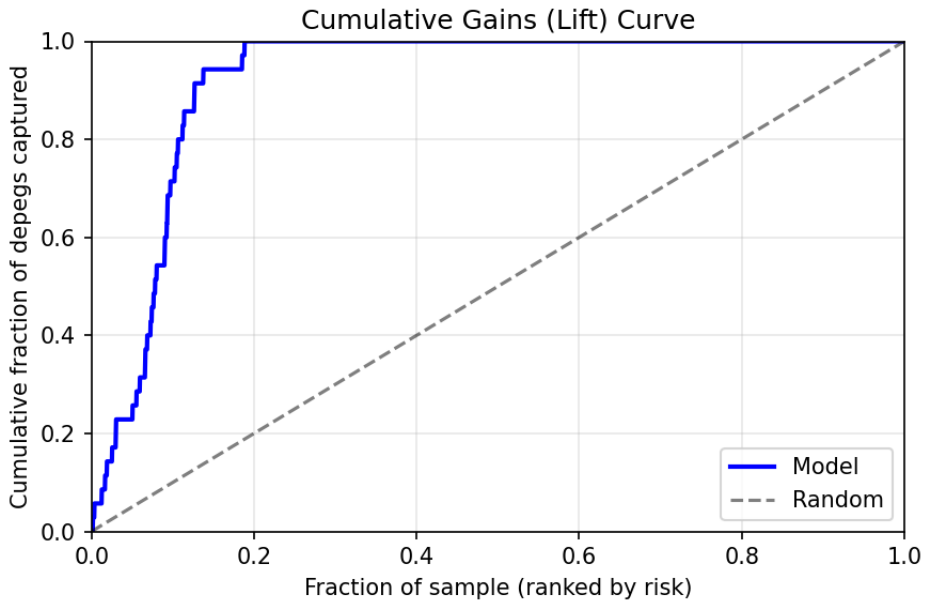


Figure 13. Lift curve (Top 1% and Top 5%) for Logit, Probit, Random Forest, and XGBoost.

5.3 Calibration and Operational Thresholds

Performance on ranking does not by itself ensure a strong early warning system. Supervisory detection in practice requires calibrated probabilities and a decision rule to turn expected risk into actionable alerts. Hence, this section investigates model calibration, alongside a comparison to alternative threshold policies for operational signaling.

Table 11. Calibration Metrics.

Model	Brier	ECE	MCC (at F1-optimal threshold)
Logit	0.019	0.048	0.178
Probit	0.018	0.044	0.183
Random Forest	0.056	0.160	0.188
XGBoost	0.013	0.010	0.239

Table 11 shows that the quality of calibration is not necessarily linked with ranking. XGBoost has a good degree of discrimination and comparatively good calibration. Logit and Probit, on the other hand, remain benchmarks that are stable and interpretable. Random Forest has lower calibration; however, it performs strong ranking, which demonstrates the gap between discrimination and probability reliability.

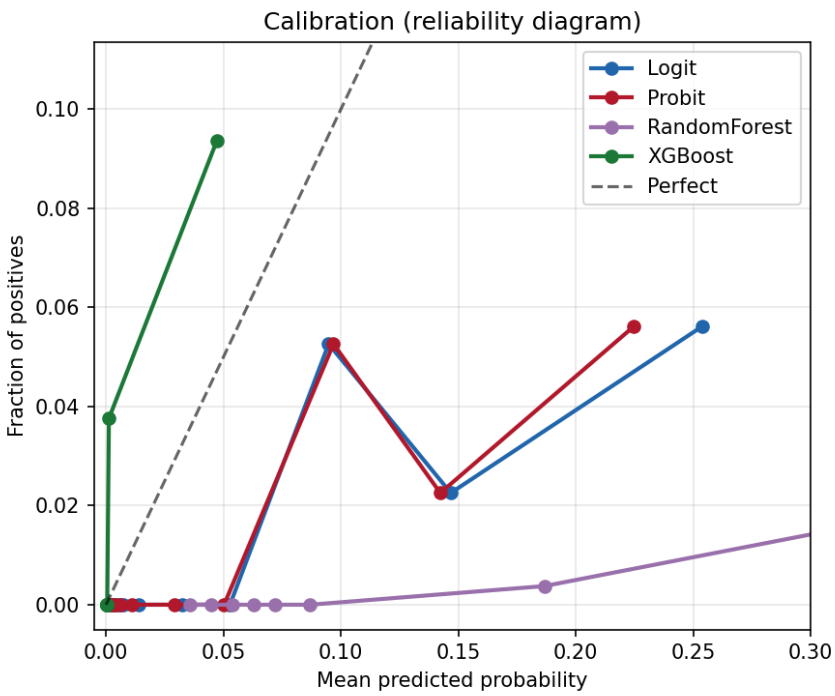


Figure 14. Reliability diagrams for probability calibration across models

As Figure 14 shows, calibration varies by model, and calibrated probabilities should guide the use of operational thresholds, particularly for the machine-learning specifications.

5.3.1 Operational Threshold Policies

In rare-event settings, thresholds mechanically tuned using F1 or Youden may generate unstable operational signaling under distribution shift. This is especially relevant for tree-based models, where train-selected thresholds can imply zero alerts out of sample.

Rather than treating threshold selection as purely statistical optimization, this study treats it as an operational policy choice. We therefore compare reference thresholds (F1-train, Youden) with tail-based monitoring rules (Top 1%, Top 5%) and target alert-frequency policies.

Detailed threshold-policy results are reported in Appendix Table A1, which shows two main patterns. First, parametric models exhibit more stable threshold behavior across policies. Second, XGBoost concentrates risk more strongly in the extreme upper tail, which makes tail-based policies operationally preferable despite instability under conventional cut-offs.

These findings justify distinguishing risk ranking from alert-generation policy. As an operational tool, tail-based policies or limited alert-frequency rules would be preferable to mechanically optimized thresholds.

In general, calibration and threshold evidence indicate that good model discrimination does not automatically translate into reliable operational signaling. Practical deployment consists thus in a dual approach that melds calibrated probabilities with purposefully selected alert policies.

5.4 Drivers of Depeg Risk and Interpretability

This section examines which indicators drive depeg risk and whether the model signals are economically interpretable in addition to being predictive. Whereas Sections 5.1–5.3 concentrated on predictive capability, discrimination, and threshold behavior, the objective here is to gauge which of these variables has the most significant explanatory signal and to what extent the machine-learning results converge with the parametric evidence.

Table 12. Logit coefficients table.

Variable	Coef	OR	p-value	Interpretation
const	−3.687	0.025	<0.001 ***	Intercept
abs_dev	+0.137	1.147	<0.001 ***	Higher deviation → higher P(depeg) (control)
log_vol_30d (H1)	−0.882	0.414	<0.001 ***	Higher 30-day volume → lower P(depeg)
dex_tvl_usd (H1)	−0.037	0.964	0.623	More DEX TVL → lower risk (not significant; collinear)
dex_volume_30d (H1)	+0.005	1.005	0.956	Not significant; sign opposite to theory (collinear)
dex_share_max (H1)	−0.738	0.478	<0.001 ***	Drop in primary-pool share → higher risk
ln_binance_volume_30d (H1, aux)	−0.278	0.757	<0.001 ***	Higher CEX aggregate 30d volume → lower risk
supply_pct_change_7d (H2)	+0.050	1.051	0.027 **	Higher 7-day supply growth → higher P(depeg) (tail-sensitive)
btc_return_7d (H3)	−0.339	0.712	<0.001 ***	Higher BTC 7-day return → lower P(depeg)
eth_return_7d (H3)	+0.196	1.217	0.014 **	Higher ETH 7-day return → higher P(depeg) (wrong sign)
dominance_eth	−0.306	0.737	<0.001 ***	Higher Ethereum supply share → lower risk
tt_crypto	+0.221	1.247	0.017 **	Crypto-backed vs fiat (control)
tt_algo	−0.752	0.471	<0.001 ***	Algorithmic vs fiat after controls (OR ≈ 0.47)

The baseline Logit coefficient estimates for the 7-day-ahead depeg specification are reported in Table 12. Among the five liquidity-related regressors, three have a negative sign at $p<0.01$ as H1 has predicted: the log 30-day token volume (*log_vol_30d*), the primary-pool share (*dex_share_max*), and the log Binance aggregate 30-day volume (*ln_binance_volume_30d*). DEX Total Value Locked (*dex_tvl_usd*) is directionally correct but not statistically significant, likely because it is collinear with *log_vol_30d* and *dex_share_max*, while the rolling DEX volume (*dex_volume_30d*) is insignificant with a theoretically inconsistent positive sign. Block-level ablation confirms that the liquidity set as a whole remains predictively important: removing all DEX variables lowers the test ROC-AUC by 0.019. Among these, *dex_share_max* shows the strongest marginal contribution ($\beta = -0.738$, OR ≈ 0.48).

The supply-growth variable (H2) enters positively and significantly in the baseline specification ($\beta = +0.050$, $p = 0.027$), but, as shown in the winsorization robustness check, its sign is sensitive to extreme right-tail observations. This is why H2 is treated as weak support rather than a fully confirmed structural effect. Market-return variables (H3) are individually significant under the canonical estimation run: the 7-day BTC return is significantly negative as predicted ($\beta = -0.339$, $p < 0.001$), while the 7-day ETH return is significantly positive ($\beta = +0.196$, $p = 0.014$), opposite to the hypothesized sign. The two returns are highly correlated in the sample, and ablating the entire market block lowers the test ROC-AUC by 0.185 (the largest single-block drop in Table 13). The ETH sign asymmetry is therefore best interpreted as a collinearity artefact plus an ETH-specific risk-on channel rather than a direct rejection of H3.

Overall, coefficient evidence suggests that liquidity deterioration is the central structural warning signal in the baseline model.

5.4.2 Feature Group Contribution (Ablation)

Table 13. Feature Group Ablation Results.

Group removed	ROC-AUC	Δ vs full
DEX (tvl, vol, share)	0.833	-0.019
Market (vol, dom, btc, eth)	0.667	-0.185
Price (abs_dev)	0.851	-0.001
Supply 7d (H2)	0.849	-0.003
CEX vol 30d (H1, aux)	0.846	-0.005
Token type	0.823	-0.029

Feature ablation investigates how much predictive power is lost when entire blocks of variables are excised. Taking away market controls and token-type information leads to the biggest drop in ROC-AUC while DEX liquidity variables also add significant value to model discrimination. The removal of the supply block, however, only changes AUC marginally, in keeping with the weak-support interpretation of H2.

This means that the model is not driven by an isolated predictor but liquidity and market-structure variables form the dominant information set.

5.4.3 SHAP Feature Importance (XGBoost)

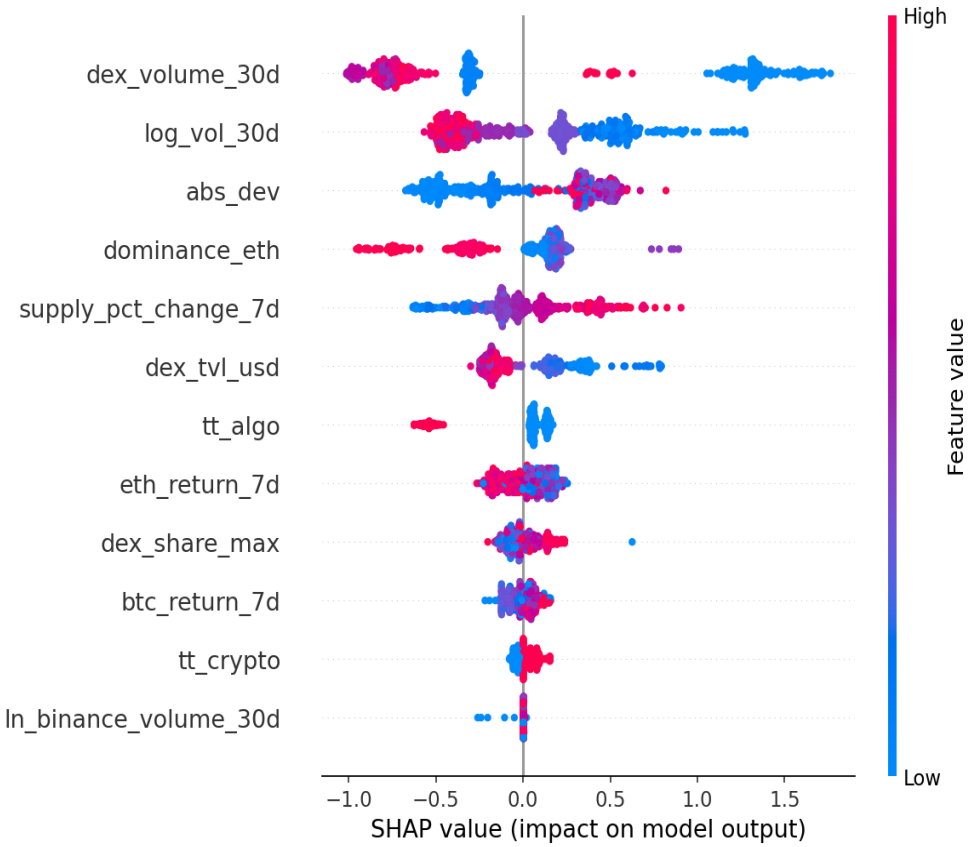


Figure 15. SHAP summary: mean absolute impact of each feature on XGBoost predictions.

To check if similar drivers are detected by machine learning, we report in Figure 15 the SHAP feature importance for XGBoost. The prevailing features mostly correspond to the Logit findings: *dex_share_max*, token type, trading volume, and liquidity-related variables are also most significant predictors of the crisis risk.

The consistency between interpretable and flexible models strengthens confidence that key signals are structural rather than model-specific artifacts.

5.4.4 Event-Study Illustration

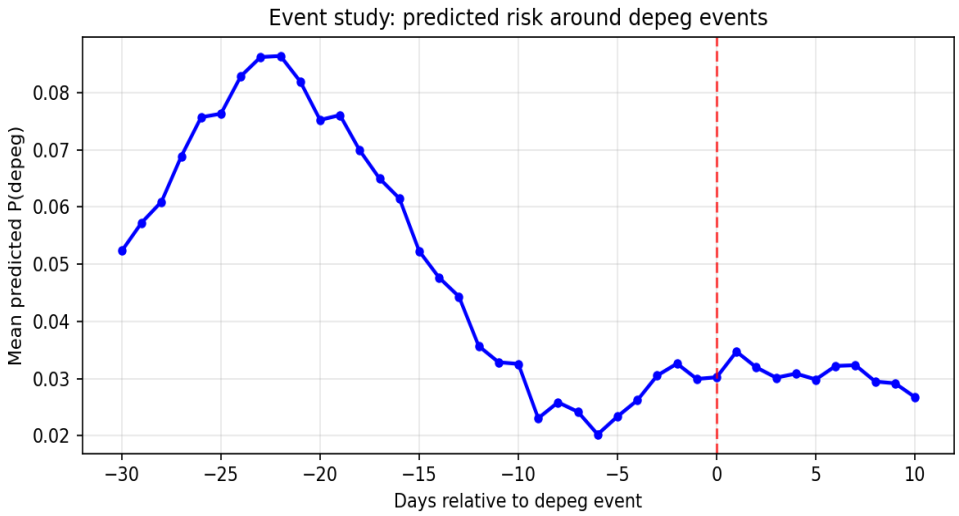


Figure 16. Event study: mean predicted P(depeg) around depeg days.

As a further interpretability check, the event-study figure plots average predicted risk around depeg starts. Predicted probabilities increase prior to crisis onset and peak around the event window, suggesting that the system operates more like an early warning mechanism rather than merely reacting contemporaneously.

Aggregated, the evidence shows that liquidity deterioration is the most important depeg-risk driver, and flexible machine-learning models enhance risk ranking by concentrating these risks more strongly in the right tail rather than uncovering an entirely different signal structure.

5.5 Robustness

The robustness checks revealed that the primary findings are consistent under other sample restrictions, feature treatments, threshold policies, and crisis-label perturbations. Specifically, burn-in exclusions, supply winsorization, train-test

threshold sensitivity, and the perturbations of depeg-rule parameters do not affect the central inferences: liquidity variables are still the best predictors, the ranking result XGBoost > Random Forest > Probit ≈ Logit is maintained, and H2 correctly is treated as weak support because of tail sensitivity.

Further, a sensitivity check around the depeg-labeling rule shows moderate variation in threshold bounds, volatility weights, MAD weights, and volume-ratio triggers mechanically influencing crisis incidence but does not change qualitative empirical interpretation.

A detailed robustness evidence including depeg-rule parameter sensitivity is described in Appendix B since it is known that these checks will not require any revision of reported test-set results, and thus will only be considered as supporting validation, rather than direct empirical results.

5.6 Hypothesis Evaluation

Hypotheses are judged using an ordinary binary-response decision rule: support needs estimated coefficients to exhibit the theorised sign and to be statistically significant given robust inference (Wooldridge 2010). If either scenario doesn't work, the evidence is rated as weak rather than full support. Consistent with early warning system philosophy, economic interpretation is based on marginal effects rather than raw coefficients (Bussière and Fratzscher 2006).

Table 14. Hypothesis evaluation summary.

Hypothesis	Variables	Expected Sign	Empirical Result	Conclusion
H1	*log_vol_30d, dex_tvl_usd, dex_volume_30d, dex_share_max, ln_binance_volume_30d*	negative	3 of 5 significant with correct sign in Logit (same in Probit); 2 remaining NS – 1 correct sign (*dex_tvl_usd*), 1 wrong sign (*dex_volume_30d,* collinear)	supported
H2	*supply_pct_change_7d*	positive	+0.050 (p = 0.027) in baseline; sign flips under 1–99% winsorization	weak support (tail-sensitive)
H3	*btc_return_7d, eth_return_7d*	negative	BTC −0.339 *** supports H3; ETH +0.196 ** opposite sign; block AUC drop 0.185	partial support (BTC channel only)

H1 Supported. Three out of five liquidity indicators have anticipated negative and significant coefficients; ablation suggests that liquidity factors materially exert predictive power.

H2 Weak support. Supply growth is positive in baseline estimates but may be sensitive to tail treatment, meaning the evidence is not robust.

H3 Partial support. The evidence for BTC return is consistent with our hypothesis; ETH shows reverse direction, which is indicative of collinearity and asymmetric transmission. Overall, the evidence for the hypothesis suggests that the signal of liquidity drop is the most precise early-warning signal, while market-return variables

are secondary but useful predictors and supply growth is another more indirect but less stable indicator.

Chapter 6

CONCLUSIONS AND POLICY RECOMMENDATIONS

In this thesis, we develop an empirical early warning framework for stablecoin depeg risk to aid monitoring decisions in turbulent digital-asset markets. In this forward-looking multi-token design, the findings suggest a clear signaling hierarchy. Liquidity proxies turn out to be the most stable predictors of depeg risk while supply and other market indicators provide weaker and more context-dependent effects.

As for predictions, the machine-learning models perform best in ranking where XGBoost scores test ROC-AUC 0.927 (versus logit and probit) that is roughly 0.85. In some of the findings, however, one important distinction between predictive ranking and operational usability has been suggested: models that deliver better discrimination do not necessarily produce reliable alert signals at set threshold rates. This leads to more adaptive tail risk management strategies, such as Top 1% and Top 5% risk policies, and also more relevant rules for deployment in rare-event conditions.

This thesis contributes to an integrated approach to constructing an early warning system in stablecoins, which integrates liquidity, supply, and market signals, clearly separating model rank from alert-rule performance. A more general indication of this is provided by the observation that deterioration of liquidity is the most telling early warning in stablecoin stress episodes, and that setting of operational threshold matters as much as model design.

At the same time, we want to rationalize the findings against several constraints like extreme rare-event bias (about 4% of observations in crisis), threshold heterogeneity between samples, and shortcomings in token coverage and data availability. These limitations suggest broader data coverage in the future, stronger calibration, as well as more frequent monitoring with more frequent input.

6.1 Limitations and Future Research

The common limitations associated with the empirical design are the following: 1. The dataset presents an unbalanced panel because stablecoins differ by launch histories, market coverage, and distribution between CEX and DEX venues; however, various variables do have occasional gaps and uneven coverage. It is important here that the structures of fiat-backed stablecoins, crypto-backed stablecoins, and algorithmic stablecoins vary, generating more heterogeneity; the scant number of depeg episodes also reflects the rare-event nature of the problem. The depeg label is, for one, also based in part on calibrated thresholds and heterogeneous data sources, which introduces some sensitivity in the identification of the crisis as well as calibrating operational alerts. Even more generally, certain depeg episodes might be triggered or enhanced by exogenous shocks, e.g. pressures from the banking sector, regulatory triggers, or idiosyncratic confidence gyrations, which cannot be entirely encompassed in market-derived factors. These limitations reflect numerous potential lines of future research. Wider token coverage, longer samples over more stress episodes, and ongoing data collection should additionally improve external validity of results. An obvious logical evolution would be to move from static estimation to one that continuously recalibrates the real-time monitoring framework, which could require the additional signals to be high in frequency, with richer on-chain variables and external stress indicators. Future works may also explore differences through stablecoin models and external validation on newer issuances and new depeg episodes.

6.2 Policy Recommendations

The empirical results lead to some suggestions of utility. For risk management practitioners and institutional investors, monitoring systems need to concentrate on signs of liquidity deterioration and concentration risks across critical trading pools rather than just deviations in price. For stablecoin issuers and DeFi protocols, the results support resilient secondary-market liquidity and transparent monitoring

of liquidity and reserve conditions. And the results favour including market-based liquidity indicators in stablecoin supervision as a complement to issuer-based disclosures for supervisory authorities.

The second implication relates to the operational application of these signals. Model outputs should work as risk prioritization tools (not as mere mechanical trigger rules; threshold instability is questionable under distribution shifts). In order to enhance practical supervision, the findings suggest the use of adaptive tail-based policies in particular top risk buckets, such as Top 1% and Top 5%, along with periodic recalibration and contingency monitoring rules. More generally, supervisory monitoring could benefit from a more dynamic early warning architecture, rather than being a static one-off approach.

Taken all together, the results indicate stablecoin fragility can be monitored through empirical data, however early warning systems within these contexts are probably more appropriately conceived as adaptive risk-monitoring tools rather than static 'crisis prediction' instruments.

WORKS CITED

Ahmed, Rashad, Iñaki Aldasoro, and Chanelle Duley. 2024. "Public Information and Stablecoin Runs." *BIS Working Papers* No. 1164. Basel: Bank for International Settlements. https://www.bis.org/publ/work1164.htm.

Amihud, Yakov. 2002. "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets* 5, no. 1: 31–56. https://linkinghub.elsevier.com/retrieve/pii/S1386418101000246.

Anadu, Kenechukwu E., Pablo D. Azar, Marco Cipriani, Thomas M. Eisenbach, Catherine Huang, Mattia Landoni, Gabriele La Spada, Marco Macchiavelli, Antoine Malfroy-Camine, and J. Christina Wang. 2024. "Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?" *Federal Reserve Bank of New York Staff Reports* No. 1073. New York: Federal Reserve Bank of New York. https://www.newyorkfed.org/research/staff_reports/sr1073.html.

Andersen, Torben G. 1996. "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility." *Journal of Finance* 51, no. 1: 169–204. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1996.tb05206.x.

Andreessen Horowitz. 2025. *State of Crypto Report 2025*. San Francisco, CA: Andreessen Horowitz. https://a16zcrypto.com/posts/article/state-of-crypto-report-2025/.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Ante, Lennart, Ingo Fiedler, and Elias Strehle. 2021. "The Influence of Stablecoin Issuances on Cryptocurrency Markets." *Finance Research Letters 41*: 101867. https://doi.org/10.1016/j.frl.2020.101867.

Barboza, Flavio, Herbert Kimura, and Edward Altman. 2017. "Machine Learning Models and Bankruptcy Prediction." *Expert Systems with Applications* 83: 405–417. https://linkinghub.elsevier.com/retrieve/pii/S0957417417302415.

Barthélémy, Sylvain, Virginie Gautier, and Fabien Rondeau. 2024. "Early Warning System for Currency Crises Using Long Short-Term Memory and Gated Recurrent Unit Neural Networks." *Journal of Forecasting* 43, no. 5: 1235–1262. https://doi.org/10.1002/for.3069.

Berg, Andrew, and Catherine Pattillo. 1999. "Predicting Currency Crises: The Indicators Approach and an Alternative." *Journal of International Money and Finance 18*, no. 4: 561–586. https://linkinghub.elsevier.com/retrieve/pii/S0261560699000248.

Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31, no. 3: 307–327. https://linkinghub.elsevier.com/retrieve/pii/0304407686900631.

Breiman, Leo. 2001. "Random Forests." Machine Learning 45, no. 1: 5–32. https://link.springer.com/article/10.1023/A:1010933404324/fulltext.html.

Brier, Glenn W. 1950. "Verification of Forecasts Expressed in Terms of Probability." *Monthly Weather Review* 78, no. 1: 1–3. https://journals.ametsoc.org/view/journals/mwrev/mwrev/78/1/1520-0493_1950_078_0001_0001.xml.

Briola, Antonio, David Vidal-Tomás, Yuanrong Wang, and Tomaso Aste. 2023. "Anatomy of a Stablecoin's Failure: The Terra-Luna Case." *Finance Research Letters* 51: 103358. https://linkinghub.elsevier.com/retrieve/pii/S1544612322005359.

Bucevska, Vesna. 2015. "Currency Crises in EU Candidate Countries: An Early Warning System Approach." *Panoeconomicus* 62, no. 4: 493–511. https://ideas.repec.org/a/voj/journl/v62y2015i4p493-510id59.html.

Bullmann, Dirk, Jonas Klemm, and Andrea Pinna. 2019. "In Search for Stability in Crypto-Assets: Are Stablecoins the Solution?" *ECB Occasional Paper* Series No. 230. Frankfurt am Main: European Central Bank. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3444847.

Bussière, Matthieu, and Marcel Fratzscher. 2006. "Towards a New Early Warning System of Financial Crises." *Journal of International Money and Finance* 25, no. 6: 953–973. https://linkinghub.elsevier.com/retrieve/pii/S0261560606000532.

Carmona, Pedro, Francisco Climent, and Alexandre Momparler. 2019. "Predicting Failure in the U.S. Banking Sector: An Extreme Gradient Boosting Approach." *International Review of Economics & Finance* 61: 304–323. https://linkinghub.elsevier.com/retrieve/pii/S1059056017306950.

Chen, Tianqi, and Carlos Guestrin. 2016. "XGBoost: A Scalable Tree Boosting System." In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. New York: Association for Computing Machinery. https://dl.acm.org/doi/10.1145/2939672.2939785.

Cintra, Thomas N., and Maxwell P. Holloway. 2023. "Detecting Depegs: Towards Safer Passive Liquidity Provision on Curve Finance." *arXiv preprint arXiv:2306.10612*. https://arxiv.org/abs/2306.10612.

Clements, Ryan. 2021. "Built to Fail: The Inherent Fragility of Algorithmic Stablecoins." *Wake Forest Law Review Online* 11: 131–156. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3952045.

Comelli, Fabio. 2014. "Comparing the Performance of Logit and Probit Early Warning Systems for Currency Crises in Emerging Market Economies." *IMF Working Paper* No. WP/14/65. Washington, DC: International Monetary Fund. https://www.imf.org/en/Publications/WP/Issues/2016/12/31/Comparing-the-Performance-of-Logit-and-Probit-Early-Warning-Systems-for-Currency-Crises-in-41489.

d'Avernas, Adrien, Vincent Maurin, and Quentin Vandeweyer. 2022. "Can Stablecoins Be Stable?" *Becker Friedman Institute Working Paper* No. 2022-131. Chicago, IL: University of Chicago. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4226027.

Davis, E. Philip, and Dilruba Karim. 2008. "Comparing Early Warning Systems for Banking Crises." *Journal of Financial Stability* 4, no. 2: 89–120. https://linkinghub.elsevier.com/retrieve/pii/S1572308908000144.

DeFiLlama. 2026. "Stablecoins Dashboard." Accessed February 11, 2026. https://defillama.com/stablecoins.

Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91, no. 3: 401–419. https://www.journals.uchicago.edu/doi/abs/10.1086/261155.

Drehmann, Mathias, and Mikael Juselius. 2014. "Evaluating Early Warning Indicators of Banking Crises: Satisfying Policy Requirements." *International Journal of Forecasting* 30, no. 3: 759–780. https://linkinghub.elsevier.com/retrieve/pii/S0169207013001428.

Duan, Kun, and Andrew Urquhart. 2023. "The Instability of Stablecoins." *Finance Research Letters* 52: 103573. https://linkinghub.elsevier.com/retrieve/pii/S1544612322007498.

Duffie, Darrell, and Jun Pan. 1997. "An Overview of Value at Risk." *Journal of Derivatives* 4, no. 3: 7–49. https://www.pm-research.com/content/iijderiv/4/3/7.

Eichengreen, Barry, My T. Nguyen, and Ganesh Viswanath-Natraj. 2025. "Stablecoin Devaluation Risk." *The European Journal of Finance* 31, no. 11: 1469–1496. https://www.tandfonline.com/doi/full/10.1080/1351847X.2025.2505757.

Eichengreen, Barry, Andrew K. Rose, and Charles Wyplosz. 1995. "Exchange Market Mayhem: The Antecedents and Aftermath of Speculative Attacks." *Economic Policy* 10, no. 21: 249–312. https://www.jstor.org/stable/1344591.

Engle, Robert F. 1982. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50, no. 4: 987–1007. https://www.jstor.org/stable/1912773.

Florez-Lopez, Raquel. 2007. "Modelling of Insurers' Rating Determinants: An Application of Machine Learning Techniques and Statistical Models." *European Journal of Operational Research* 183, no. 3: 1488–1512. https://linkinghub.elsevier.com/retrieve/pii/S0377221706011891.

Frankel, Jeffrey A., and Andrew K. Rose. 1996. "Currency Crashes in Emerging Markets: An Empirical Treatment." *Journal of International Economics* 41, no. 3–4: 351–366. https://linkinghub.elsevier.com/retrieve/pii/S0022199696014419.

Gadzinski, Gregory, Alessio Castello, and Florian Mazzorana. 2023. "Stablecoins: Does Design Affect Stability?" *Finance Research Letters* 53: 103611. https://linkinghub.elsevier.com/retrieve/pii/S1544612322007875.

Huber, Peter J. 1981. *Robust Statistics*. New York: John Wiley & Sons.

Hyndman, Rob J., and George Athanasopoulos. 2021. Forecasting: Principles and Practice. 3rd ed. Melbourne: OTexts. https://otexts.com/fpp3/.

Jorion, Philippe. 2007. *Value at Risk: The New Benchmark for Managing Financial Risk*. 3rd ed. New York: McGraw-Hill.

Kaminsky, Graciela, Saul Lizondo, and Carmen M. Reinhart. 1998. "Leading Indicators of Currency Crises." *IMF Staff Papers* 45, no. 1: 1–48. https://www.jstor.org/stable/3867328.

Kazemian, Sam, Kedar Iyer, and Jason Huan. 2020. Frax: Fractional-Algorithmic Stablecoin Protocol. *Frax Finance White Paper,* July 2020. https://github.com/FraxFinance/frax-solidity/blob/master/frax_whitepaper_v1.pdf.

Klages-Mundt, Ariah, and Andreea Minca. 2021. "(In)Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks." *Cryptoeconomic Systems* 1, no. 1. https://www.researchgate.net/publication/333640394_InStability_for_the_Blockchain_Deleveraging_Spirals_and_Stablecoin_Attacks.

Kosse, Anneke, Marc Glowka, Ilaria Mattei, and Tara Rice. 2023. "Will the Real Stablecoin Please Stand Up?" *BIS Papers* No. 141. Basel: Bank for International Settlements. https://www.bis.org/publ/bppdf/bispap141.htm.

Krause, David. 2025. "Algorithmic Stablecoins: Mechanisms, Risks, and Lessons from the Fall of TerraUSD." Working paper. *SSRN Electronic Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5092827.

Krugman, Paul. 1979. "A Model of Balance-of-Payments Crises." *Journal of Money, Credit and Banking* 11, no. 3: 311–325. https://www.jstor.org/stable/1991793.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53, no. 6: 1315–1335. https://www.jstor.org/stable/1913210.

Lee, Yi-Hsi, Yu-Fen Chiu, and Ming-Hua Hsieh. 2025. "Stablecoin Depegging Risk Prediction." *Pacific-Basin Finance Journal* 90: 102640. https://doi.org/10.1016/j.pacfin.2024.102640.

Liao, Gordon Y., and John Caramichael. 2022. "Stablecoins: Growth Potential and Impact on Banking." *Board of Governors of the Federal Reserve System International Finance Discussion Papers* No. 1334. Washington, DC: Federal Reserve. https://www.federalreserve.gov/econres/ifdp/files/ifdp1334.pdf.

Liu, Jiageng, Igor Makarov, and Antoinette Schoar. 2023. "Anatomy of a Run: The Terra Luna Crash." *NBER Working Paper* No. 31160. Cambridge, MA: National Bureau of Economic Research. https://www.nber.org/papers/w31160.

Lyons, Richard K., and Ganesh Viswanath-Natraj. 2023. "What Keeps Stablecoins Stable?" *Journal of International Money and Finance* 131: 102777. https://linkinghub.elsevier.com/retrieve/pii/S0261560622001802.

MacDonald, Cameron, and Laura Zhao. 2022. "Stablecoins and Their Risks to Financial Stability." *Bank of Canada Staff Discussion Paper* 2022-20. Ottawa: Bank of Canada. https://www.bankofcanada.ca/2022/11/staff-discussion-paper-2022-20.

Moody's. 2023. "Moody's Launches New Digital Asset Monitor to Track Risk." *Moody's Insights*, November 6, 2023. https://www.moodys.com/web/en/us/insights/banking/moody-launches-new-digital-asset-monitor-to-track-risk.html.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88, no. 3: 587–597. https://www.jstor.org/stable/116850.

Obstfeld, Maurice. 1996. "Models of Currency Crises with Self-Fulfilling Features." *European Economic Review* 40, no. 3–5: 1037–1047. https://linkinghub.elsevier.com/retrieve/pii/0014292195001115.

Platt, John C. 1999. "Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods." In *Advances in Large Margin Classifiers*, edited by Alexander J. Smola, Peter L. Bartlett, Bernhard Schölkopf, and Dale Schuurmans, 61–74. Cambridge, MA: MIT Press.

Rousseeuw, Peter J., and Christophe Croux. 1993. "Alternatives to the Median Absolute Deviation." *Journal of the American Statistical Association* 88, no. 424: 1273–1283. https://www.tandfonline.com/doi/abs/10.1080/01621459.1993.10476408.

S&P Global. 2023. "Stablecoin Stability Assessment." *S&P Global Ratings*. https://www.spglobal.com/ratings/en/products/stablecoin-stability-assessment.

Sarlin, Peter. 2014. "On Policymakers' Loss Functions and the Evaluation of Early Warning Systems." *Economics Letters* 119, no. 1: 1–7. https://linkinghub.elsevier.com/retrieve/pii/S0165176513000025.

Shrivastav, Lokesh Kumar, and Ravinder Kumar. "An Ensemble of Random Forest Gradient Boosting Machine and Deep Learning Methods for Stock Price Prediction," Journal of Information Technology Research (JITR) 15, no.1: 1-19. https://doi.org/10.4018/JITR.2022010102.

Uhlig, Harald. 2022. "A Luna-tic Stablecoin Crash." *NBER Working Paper* No. 30256. Cambridge, MA: National Bureau of Economic Research. https://www.nber.org/papers/w30256.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* 2nd ed. Cambridge, MA: MIT Press.

Xu, Yunsong, Jiaqi Li, and Anqi Wu. 2024. "An XGboost Algorithm Based Model for Financial Risk Prediction." *Tehnički vjesnik* 31, no. 6: 1898–1907. https://doi.org/10.17559/TV-20231021001043.

Yildirim, Yusuf, and Anirban Sanyal. 2022. "Evaluating the Effectiveness of Early Warning Indicators: An Application of Receiver Operating Characteristic Curve Approach to Panel Data." *Scientific Annals of Economics and Business* 69, no. 4: 557–597. https://saeb.feaa.uaic.ro/index.php/saeb/article/view/1529.

Appendix A: Threshold Policy Sensitivity

Table 15. Alternative Threshold Sensitivity Results.

Model	Policy	Threshold	Alerts/day (cal)	Alerts/day (token)	n_cal	n_token	Precision	Recall	F1	MCC
Logit	Youden (reference)	0.081	0.90	1.63	401	723	0.036	0.743	0.069	0.122
Logit	F1-train (reference)	0.142	0.63	0.99	279	439	0.048	0.600	0.089	0.135
Logit	Top 1% risk (operational)	0.354	0.06	0.06	26	26	0.000	0.000	0.000	−0.011
Logit	Top 5% risk (operational)	0.250	0.30	0.30	132	133	0.060	0.229	0.095	0.095
Logit	Target 0.10 alerts/day	0.313	0.10	0.10	45	45	0.089	0.114	0.087	0.087
Probit	Youden (reference)	0.074	0.98	1.74	437	772	0.045	1.000	0.087	0.181
Probit	F1-train (reference)	0.135	0.64	1.02	286	454	0.046	0.600	0.086	0.132
Probit	Top 1% risk (operational)	0.300	0.06	0.06	26	26	0.039	0.029	0.033	0.022
Probit	Top 5% risk (operational)	0.224	0.30	0.30	132	133	0.060	0.229	0.095	0.095
Probit	Target 0.10 alerts/day	0.267	0.10	0.10	45	45	0.089	0.114	0.087	0.087
Random Forest	Youden (reference)	0.457	0.33	0.33	145	146	0.048	0.200	0.077	0.074
Random Forest	F1-train (reference)	0.583	0.01	0.01	4	4	0.250	0.029	0.051	0.081

Model	Policy	Threshold	Alerts/day (cal)	Alerts/day (token)	n_cal	n_token	Precision	Recall	F1	MCC
Random Forest	Top 1% risk (operational)	0.509	0.06	0.06	26	26	0.077	0.057	0.066	0.056
Random Forest	Top 5% risk (operational)	0.461	0.30	0.30	132	133	0.053	0.200	0.083	0.080
Random Forest	Target 0.10 alerts/day	0.491	0.10	0.10	45	46	0.043	0.057	0.049	0.035
XGBoost	Youden (reference)	0.601	0.00	0.00	0	0	0.000	0.000	0.000	0.000
XGBoost	F1-train (reference)	0.832	0.00	0.00	0	0	0.000	0.000	0.000	0.000
XGBoost	Top 1% risk (operational)	0.145	0.06	0.06	26	26	0.115	0.086	0.098	0.089
XGBoost	Top 5% risk (operational)	0.050	0.30	0.30	133	133	0.075	0.286	0.119	0.125
XGBoost	Target 0.10 alerts/day	0.116	0.10	0.10	45	45	0.067	0.086	0.075	0.062

Appendix B: Robustness Checks Summary

This appendix summarizes robustness checks reported in Chapter 5 and consolidates evidence on the stability of the main findings. It does not introduce new experiments.

Table 16. Summary of Robustness Validation Results.

Robustness check	Purpose	Outcome
Threshold / tail-policy sensitivity (F1, Youden, Top 1%, Top 5%, target alerts/day) – Section 5.4.4	Test whether operational conclusions depend on one cut-off policy	Main conclusions unchanged; supports tail-based policies for stable alerting.
Burn-in exclusions(k = 0, 90, 180 launch days)launch days per token) - Appendix B	Remove launch-phase volatility and thin-data noise	Main findings unchanged; key signs and model ordering remain stable.
Winsorization sensitivity of *supply_pct_change_7d* (1–99%, train-based cutoffs) – Section 5.4.1	Check outlier sensitivity of H2	H2 sign is tail-sensitive, but broader ranking and core conclusions are unchanged.
Train–test comparison – Tables 2–3, Section 5.1	Assess generalization under distribution shift	Overfitting risk is confirmed, especially for tree models; reinforces conservative threshold policy.
Feature-group ablation (DEX, market, price, supply, CEX, token type) - Section 5.4.2	Test dependence on any single predictor block	Liquidity and structure variables remain central; no single block overturns the main interpretation.
Depeg-rule parameter sensitivity (bounds, σ weight, MAD weight, volume triggers) - Section 5.5	Check dependence on one crisis-definition calibration	Crisis frequency shifts mechanically, but ranking and liquidity-signal conclusions remain qualitatively stable.

Synthesis. The robustness audit reaffirms the same key message–that the findings are stable across model, sample, threshold, and labeling perturbations. Liquidity signals remain the strongest and most reliable predictors. The choice of threshold has implications for the intensity of the operational signaling but not the ultimate ranking conclusion. Combined, the checks show that results are not influenced by one crisis-definition context or model specification. This appendix condenses robustness checks presented in Chapter 5 and consolidates evidence on the stability of the main findings. It does not add new experiments.