

Introduction & Background

A sovereign debt crisis occurs when a national government is on the verge of defaulting on the money it has borrowed from other countries, international investors, and major banks. These crises are far from being rare and have in fact been a recurring feature of the global economy for centuries, impacting nations from every part of the world.

A country facing a debt crisis loses the trust of its lenders which triggers a cascade of devastating consequences. This typically includes severe economic recessions, crippling unemployment, and considerable cuts to essential public services like education and healthcare. Needless to say, the effects of sovereign debt crises are felt by millions of a country's citizens and frequently impact the broader global economy as well.

Problem Definition & Motivation

The problem we are trying to address is about countries at risk of a debt crisis and how they can prevent such a crisis from happening. In order to avert it, the country needs to be able to recognize that it is in danger early on, so we want to try to figure out what the main risk factors are that can indicate a country is in danger of a debt crisis in the near future.

The need for a solution to this problem is motivated by the fact that the consequences to a country that reaches a debt crisis can be devastating to both its economy and reputation. The debt will hurt the country internally with the economic recession that this leads to, which the citizens of the country will feel. It will also decrease the opinion the rest of the world has of the country in crisis, as other countries and investors will not trust it enough to be able to repay borrowed money. The best way to avoid this is to avoid reaching the crisis point in the first place, which is why a country needs to be able to tell when it is at risk.

Literature Review

After reviewing the literature, we found that Sovereign debt crises have long been linked to macroeconomic fundamentals. While Manasse and Roubini [2] point to inflation, slow growth, and weak institutions as important risk factors, Reinhart and Rogoff [1] demonstrate that high debt loads and external imbalances frequently precede defaults. These indicators are also used by the IMF's Early Warning Systems [3], despite the fact that complex, nonlinear relationships are difficult for traditional econometric models to handle. Recent research shows that predictive accuracy can be increased through machine learning. Tree-based approaches perform better than regressions, according to Ciarlone and Trebeschi [4], while Barrell et al. [5]

demonstrate that ensemble approaches improve crisis forecasting. These findings support our application of supervised machine learning to improve prediction and systematically assess macroeconomic indicators.

Data Preprocessing Overview

Dataset Description

The dataset for this project consist of a selection of macroeconomic and financial indicators from the [World Bank Open Data](#) portal. The indicators we track for each country are GDP growth rate, inflation rate, domestic credit to the private sector as a percentage of GDP, current account balance as a percentage of GDP, and total reserves. The data ranges from 1961 to 2024 and includes coverage of 68 different countries. For labelling our data, we used Harvard Business School's "[Global Crisis Data by Country](#)", particularly using their external debt crisis data.

Our final dataset ended up having 3526 complete rows. In the screenshot below, you can see our data sorted by country, with the label column on the far right signifying if that country did or did not experience a debt crisis in the next 3 years.

Country	Year	gdp_growth_l	gdp_growth_l	gdp_growth_l	inflation_mei	inflation_std	inflation_trei	domestic_cri	domestic_cri	domestic_cri	current_acc	current_acc	current_acc	total_reserve	total_reserve	total_reserve	target_crisis_next_3y
Algeria	1964	-0.6080664	7.87601586	3.80976307	-0.0856371	-0.0831471	-0.061596	-1.060097	-0.6438892	-0.1783234	0.9271243	-0.8052171	-0.0267048	-0.2196815	-0.1993985	-0.1212902	0
Algeria	1965	-0.4513762	7.90636408	4.8090035	-0.0856371	-0.0831471	-0.061596	-1.0601106	-0.643902	-0.1781515	0.92708614	-0.8052064	-0.0267307	-0.2197778	-0.1993351	-0.1214285	0
Algeria	1966	0.18057435	7.31686308	0.14402604	-0.0856371	-0.0831471	-0.061596	-1.0600974	-0.6439362	-0.178618	0.9270589	-0.8052792	-0.0263844	-0.2198235	-0.1997611	-0.1197064	0
Algeria	1967	2.27293736	5.04151938	-4.4124674	-0.0856371	-0.0831471	-0.061596	-1.0599631	-0.6443016	-0.1790605	0.92719307	-0.8049374	-0.0260374	-0.2195197	-0.1979209	-0.1179824	0
Algeria	1968	0.58417632	1.42514077	0.98939058	-0.0856371	-0.0831471	-0.061596	-1.0599536	-0.6441869	-0.1772477	0.92732933	-0.8046861	-0.0255863	-0.2191715	-0.1965147	-0.1159048	0
Algeria	1969	0.77043863	1.48734859	1.49590514	-0.0856371	-0.0831471	-0.061596	-1.0598964	-0.6442674	-0.1774627	0.92743112	-0.8047514	-0.0258018	-0.2189168	-0.1969685	-0.1169417	0
Algeria	1970	0.9611453	1.54394128	1.95593893	-0.0856371	-0.0831471	-0.061596	-1.0598502	-0.6445223	-0.1782214	0.92751048	-0.804971	-0.0263747	-0.2187184	-0.1981342	-0.1196864	0
Algeria	1971	0.49246781	2.79571886	-3.173561	-0.0856371	-0.0831471	-0.061596	-1.0598048	-0.6451079	-0.1794979	0.92763413	-0.8050758	-0.0265052	-0.2184064	-0.1986237	-0.1202355	0
Algeria	1972	1.78294294	4.71349476	0.10369023	-0.0856371	-0.0831471	-0.061596	-1.059713	-0.6440402	-0.1781069	0.92772883	-0.8048842	-0.026165	-0.2181736	-0.1976718	-0.1186707	0
Algeria	1973	1.28149813	4.75902961	0.70737148	-0.0856371	-0.0831471	-0.061596	-1.0594974	-0.6416368	-0.1749857	0.92819594	-0.8029529	-0.024032	-0.2189998	-0.1874655	-0.1084731	0
Algeria	1974	1.21411551	4.75451283	0.93426684	-0.0856371	-0.0831471	-0.061596	-1.0590155	-0.6381091	-0.1709615	0.92916716	-0.8007775	-0.0213696	-0.2145627	-0.1760126	-0.0957541	0
Algeria	1975	0.93998984	4.75203837	0.96480448	-0.0856371	-0.0831471	-0.061596	-1.058663	-0.6386183	-0.1719544	0.92989306	-0.8011032	-0.0220387	-0.2127404	-0.1777287	-0.0986521	0
Algeria	1976	2.35595758	2.95378192	-2.6842069	-0.0856371	-0.0831471	-0.061596	-1.058152	-0.6398418	-0.1734108	0.93081981	-0.8014102	-0.0224987	-0.2104205	-0.179302	-0.1010971	0
Algeria	1977	0.76428961	-0.4102511	0.30046713	-0.0856371	-0.0831471	-0.061596	-1.0577506	-0.6417296	-0.1756049	0.93173532	-0.8030222	-0.0243891	-0.2081175	-0.1878458	-0.1101739	0
Algeria	1978	1.15217589	-0.4183177	0.29094717	-0.0856371	-0.0831471	-0.061596	-1.0573372	-0.6416553	-0.1755358	0.93253745	-0.8029755	-0.0243415	-0.2061064	-0.1875997	-0.1099464	0
Algeria	1979	1.1509486	-0.4187216	0.44082173	-0.0856371	-0.0831471	-0.061596	-1.0566234	-0.6336246	-0.1670602	0.9339692	-0.7980266	-0.0186879	-0.2025141	-0.1615299	-0.0829319	0
Algeria	1980	0.84542152	0.23551575	-0.9305594	-0.0856371	-0.0831471	-0.061596	-1.0643265	-0.5118449	-0.2725409	0.93900409	-0.322778	-0.4205617	-0.1964122	-0.1443425	-0.0638193	0
Algeria	1981	0.45860511	0.23442588	-0.9282341	-0.0856371	-0.0831471	-0.061596	-1.0718724	-0.4835203	-0.3232755	0.70610795	-0.0813961	-0.7796556	-0.1923832	-0.1481704	-0.0711648	0
Algeria	1982	0.5405689	0.25496674	-0.7200291	-0.0856371	-0.0831471	-0.061596	-1.0793023	-0.4826253	-0.3250642	0.55507259	0.01266672	-0.9370075	-0.1895707	-0.1615159	-0.1020806	0
Algeria	1983	0.26663062	-0.0558374	0.12936448	-0.0856371	-0.0831471	-0.061596	-1.0868777	-0.5090657	-0.277304	0.43241076	-0.1267466	-0.6582892	-0.1886491	-0.169243	-0.1441961	0
Algeria	1984	0.13178603	-0.2267791	0.90556758	-0.0856371	-0.0831471	-0.061596	-1.094929	-0.6383531	-0.1715905	0.32034618	-0.5660571	-0.0686203	-0.1914023	-0.1580824	-0.1655437	0
Algeria	1985	0.34070565	-0.6121076	0.06663989	-0.0856371	-0.0831471	-0.061596	-1.0941637	-0.6396697	-0.1726715	0.35963671	-0.3999251	0.37817024	-0.1942587	-0.1721976	-0.1405532	0
Algeria	1986	0.15400331	-0.1904115	-0.983968	-0.0856371	-0.0831471	-0.061596	-1.0927743	-0.6261372	-0.1611269	0.24591746	-0.0142806	-0.2867103	-0.1967052	-0.1812627	-0.1278964	0
Algeria	1987	-0.3558388	0.02436485	-1.2558023	-0.0856371	-0.0831471	-0.061596	-1.0916838	-0.6243081	-0.1579643	0.28912529	0.0080609	-0.2492465	-0.1974487	-0.1854038	-0.1140785	0
Algeria	1988	-0.8154152	0.0338926	-1.2704963	-0.0856371	-0.0831471	-0.061596	-1.0900642	-0.6216645	-0.1545533	0.16283151	0.18982318	-0.7094228	-0.1984165	-0.1823585	-0.1216404	1
Algeria	1989	-0.9015853	-0.1339194	0.02255909	-0.0856371	-0.0831471	-0.061596	-1.0883197	-0.6244957	-0.1579309	0.08667918	0.1581372	-0.5758017	-0.1985341	-0.1816403	-0.1379391	1
Algeria	1990	-1.1098304	-0.2917855	0.45602613	-0.0856371	-0.0831471	-0.061596	-1.0869283	-0.6343393	-0.1685654	0.1055214	0.230483	0.68188009	-0.2008259	-0.1825879	-0.136842	1
Algeria	1991	-1.2247244	-0.2134129	0.08133408	-0.0856371	-0.0831471	-0.061596	-1.0854972	-0.6245496	-0.1582996	0.42113833	0.60049056	1.15087142	-0.2012776	-0.183669	-0.1308346	1
Algeria	1992	-1.0452026	-0.2320479	0.02255979	-0.0856371	-0.0831471	-0.061596	-1.0717036	-0.4315667	-0.0079523	0.50767151	0.66630981	1.45392831	-0.2024885	-0.1928307	-0.1173334	1
Algeria	1993	-1.1241919	-0.1148312	-0.85907	-0.0856371	-0.0831471	-0.061596	-1.06042	-0.4044805	0.04193076	0.69777812	0.23261378	0.54116636	-0.2019393	-0.1905925	-0.1120588	1
Algeria	1994	-1.5047781	-0.5393898	-0.2933576	-0.0856371	-0.0831471	-0.061596	-1.0491559	-0.412349	0.0268026	0.61230262	0.61709332	-1.2721145	-0.1998993	-0.1792976	-0.0995874	1
Algeria	1995	-1.2893523	-0.1665878	0.55888228	-0.0856371	-0.0831471	-0.061596	-1.0423362	-0.4641739	-0.1170881	0.34161712	1.04499045	-2.1197433	-0.1981745	-0.1837605	-0.1066673	1

We elaborate on the dataset features in the following section.

Feature explanation

GDP (Annual % growth) measures the annual change in the total value of goods and services produced by a country. A low or negative growth rate shrinks government tax revenue and makes its debt burden (relative to its income) larger and more difficult to service. **Inflation (average consumer prices, index)** tracks the rate at which consumer prices are rising. High and volatile inflation is a classic sign of economic mismanagement that erodes investor confidence and can signal that a government is printing money to pay its bills. **Domestic**

[credit to private sector \(% of GDP\)](#) shows the total amount of loans given by banks to individuals and private businesses, relative to the size of the economy. A rapid, unsustainable credit “bubble” can lead to a banking crisis, which often forces the government to take on massive debts to bail out the financial system, triggering a sovereign debt crisis. [Current Account Balance \(% of GDP\)](#) measures the difference between a country’s total income from abroad (mainly exports) and its payments to the rest of the world (mainly imports), expressed as a share of its economy. A large and persistent deficit (a negative number) signals that a country is living beyond its means by borrowing heavily from foreigners, which builds up external debt and makes it highly vulnerable to a sudden stop in lending. [Total reserves including gold \(current USD\)](#) is the country’s rainy day fund of foreign currency and gold held by its central bank. Since external debt must be paid in foreign currency (like US dollars), a falling reserve level is a critical danger sign that a country is running out of the hard cash needed to pay its bills.

We summarized each of the above indicators over a 5 year window to provide the model with a dynamic, contextual picture of a country’s economic health rather than just a single-year snapshot. The mean establishes the average level (e.g., “is this a high-debt country?”), the standard deviation measures volatility (e.g., “is the economy stable or erratic?”), and the trend reveals momentum (e.g., “is the situation stable, improving, or rapidly getting worse?”). This allows our model to learn the crucial difference between a country with high-but-stable debt and a country whose debt level is suddenly and dangerously exploding.

Preprocessing Methods

Our data preprocessing began with data integration. We started by merging 9 candidate indicator datasets, and we then merged this with the sovereign crisis label dataset using an inner merge on (Country, Year). To mitigate the curse of dimensionality and address high sparsity, we performed an initial feature selection step, removing the four indicators with the poorest data coverage, leaving us with five core indicators, which are explained in the feature explanation section. The remaining NaN values were then handled using multivariate imputation (IterativeImputer), which was applied on a per-country basis to estimate missing data by modeling the inter-relationships between the indicators.

Following imputation, we executed feature summarization by transforming the five raw time-series indicators into a 15-dimensional feature space. For each indicator, we calculated three new features based on a 5-year rolling window: the mean, the standard deviation, and the linear trend/slope. Our motivations behind this are explained more within the feature explanation section. Finally, after dropping any rows still containing NaNs, the entire engineered feature set was normalized using StandardScaler to ensure all features had a mean of zero and variance of 1, making them fully prepared for model training.

Model

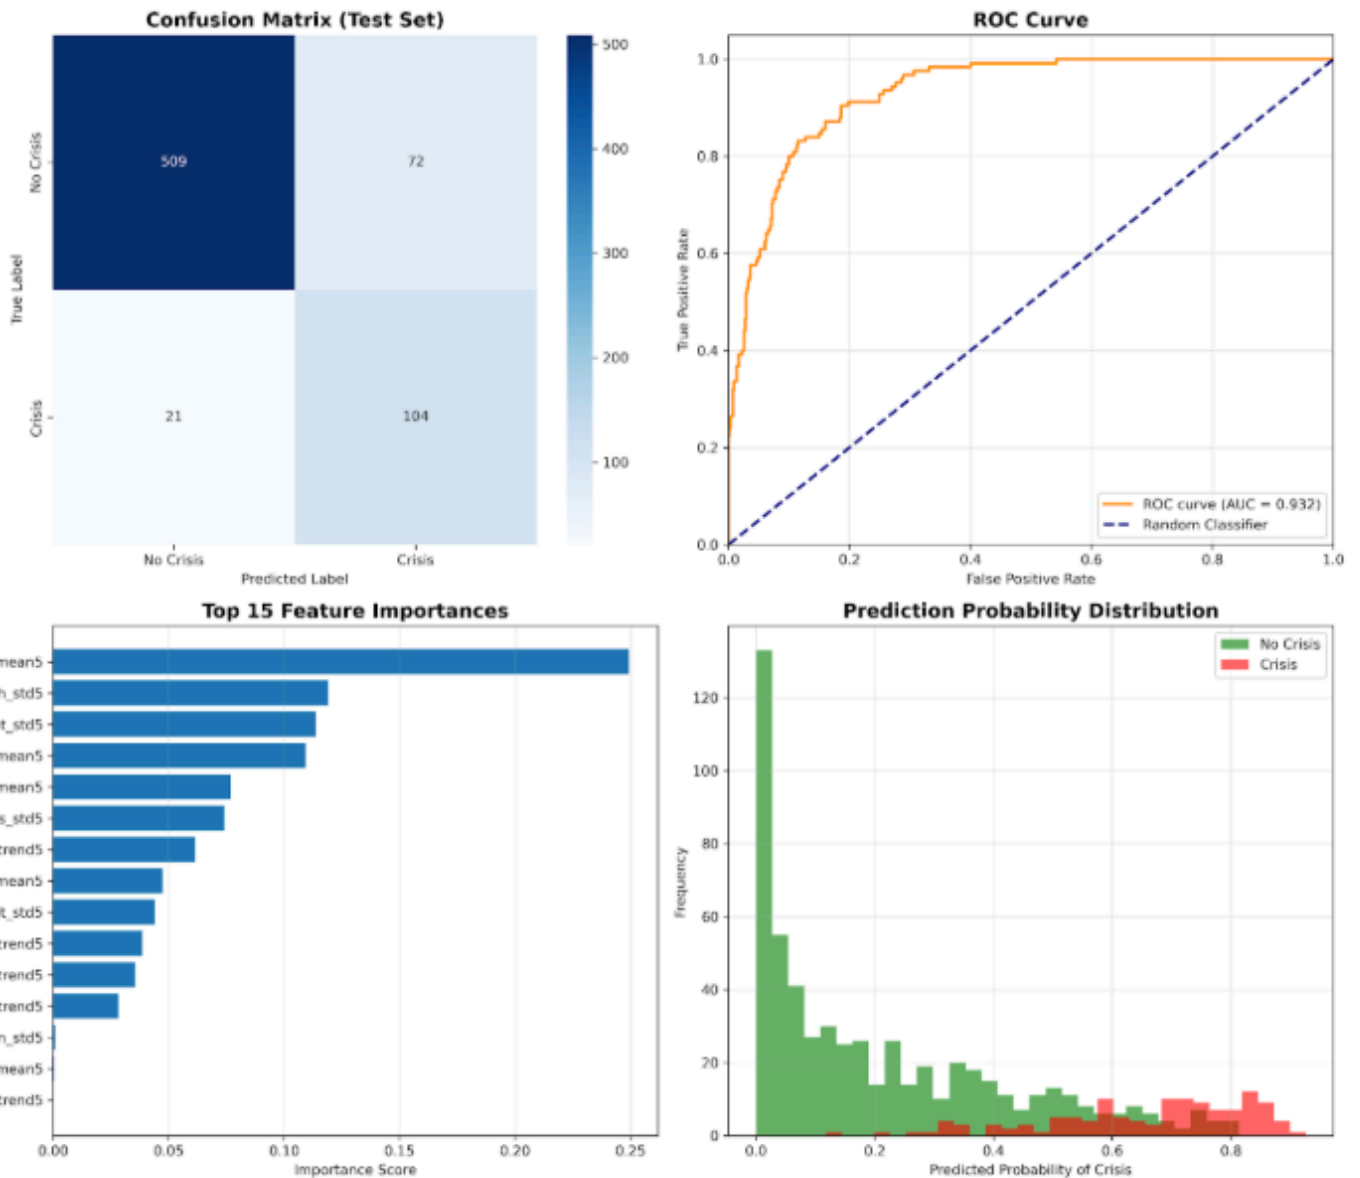
Random Forest

The Random Forest Classifier was chosen because it effectively captures the complex, nonlinear relationships among macroeconomic indicators that often precede sovereign debt crises. It performs well on structured data without requiring strict statistical assumptions and is robust against overfitting due to its

ensemble nature. Additionally, it handles class imbalance using balanced weighting and provides interpretable feature importance scores that highlight key predictors of crisis risk. Prior research also supports tree-based ensemble methods as superior to traditional econometric models for early crisis detection.

Analysis & Discussion of Models

Random Forest Metrics & Visualization



Analysis of Random Forest

After analyzing our Random Forest Classifier, we found that the model performed very well in predicting whether a country would experience a debt crisis within the next three years. It achieved an ROC-AUC score of 0.932, which shows that it could clearly distinguish between crisis and non-crisis periods. Looking at the

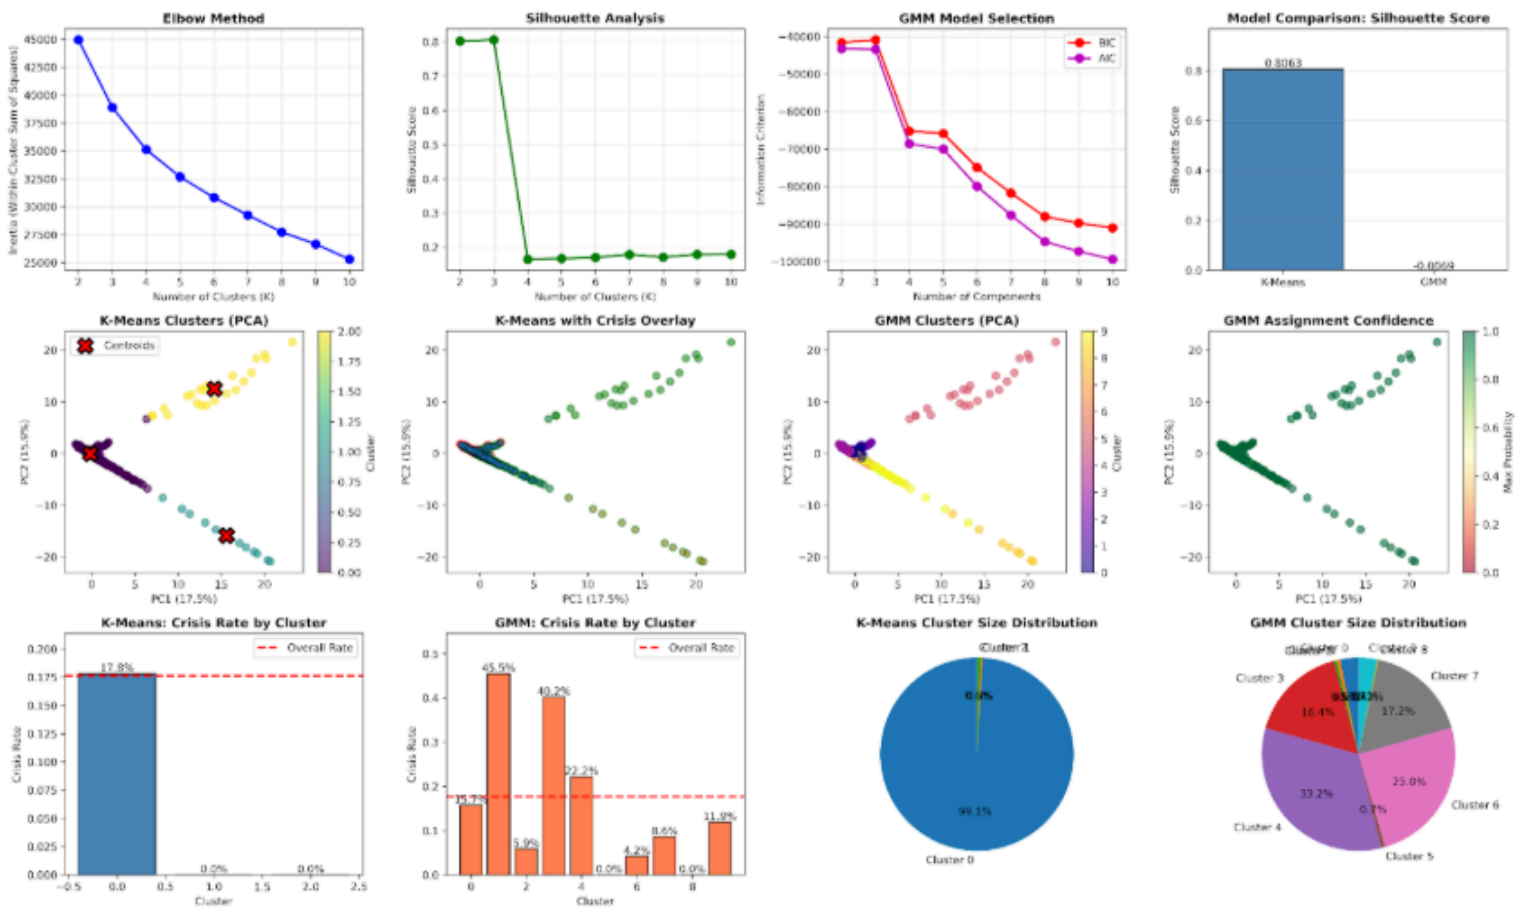
confusion matrix, most of the predictions were correct, with only a few cases where the model misclassified a year as stable when a crisis actually occurred.

We did notice that the model tended to predict "no crisis" slightly more often than "crisis." This makes sense because, in our dataset, actual crises are much rarer than stable years. Even with that challenge, using balanced class weights helped reduce this bias and improved how well the model detected real crisis situations.

A big reason for the model's success was the feature engineering we did. Instead of only looking at single-year values, we summarized each economic indicator over a five-year window, adding information about averages, volatility, and trends. This gave the model a clearer picture of how a country's economy was behaving over time. It also helped the model pick up on realistic warning signs, like rising private credit or falling reserves, which often signal financial trouble.

Overall, the Random Forest worked well because it could handle complex, nonlinear relationships between economic factors and didn't overfit to the data. The results are not only accurate but also make sense from an economic perspective, showing that the model can be a valuable tool for predicting and preventing sovereign debt crises in the future.

Kmeans/GMM Metrics & Visualization



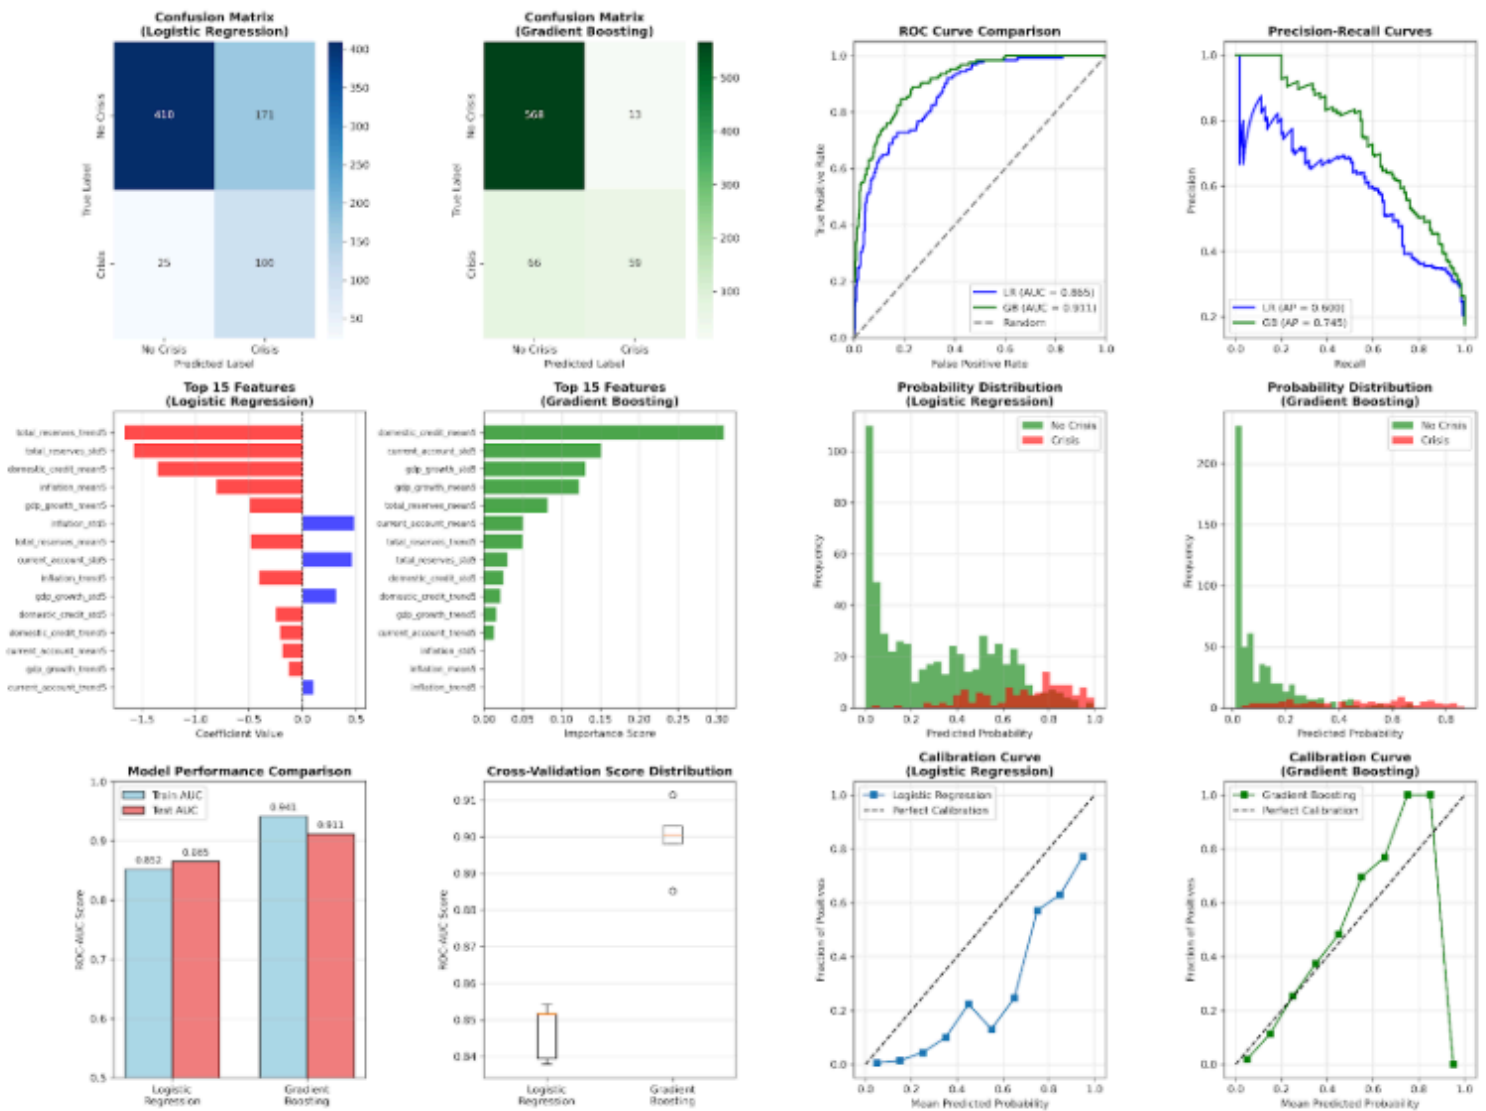
Analysis of Kmeans/GMM

After analyzing our clustering results, we found that both K-Means and GMM helped us uncover natural groupings in the macroeconomic data that were not visible from supervised models alone. Using the Elbow Method and Silhouette Analysis, we identified that $K=3$ was the optimal number of clusters for K-Means, with a silhouette score of 0.806 showing very strong separation between groups. The PCA visualizations confirmed this, as the clusters appeared cleanly separated with clear centroids. When we overlaid the crisis labels, one cluster showed a much higher crisis rate around 17.8% indicating that K-Means successfully captured a meaningful high-risk group.

GMM, on the other hand, revealed a more detailed structure in the data. The BIC and AIC charts suggested roughly 8–9 components as optimal, and the PCA plots showed smoother, overlapping cluster boundaries that better represent economic uncertainty. GMM assigned probabilities rather than fixed labels, allowing us to see how confidently each observation belonged to a specific cluster. The crisis rate plots showed several clusters with very high vulnerability, some exceeding 40% suggesting that GMM was able to detect subgroups of countries with higher systemic fragility that K-Means could not isolate. The cluster size distribution further highlighted this difference: while K-Means produced one dominant cluster, GMM distributed countries more evenly, capturing more nuanced economic patterns.

Overall, our clustering analysis provided an unsupervised perspective on sovereign risk segmentation. K-Means gave us clear, interpretable broad groups, while GMM uncovered deeper layers of crisis-prone subpopulations. Together, these methods helped validate the risk patterns identified by our supervised models and offered new insight into how countries group themselves based on shared macroeconomic vulnerabilities.

Linear Regression/Gradient Boosting Metrics & Visualization



Analysis of Linear Regression/Gradient Boosting

After reviewing the results of both Logic Regression and Gradient Boosting, we found clear differences in how each model handled the challenge of predicting sovereign debt crises. Logic Regression performed reasonably well in identifying stable years, but its confusion matrix showed that it missed many crisis cases, which is expected from a linear model facing complex nonlinear economic patterns. Its ROC-AUC score of 0.865 confirms decent performance, but not at the level needed for reliable early-warning. The probability distribution and calibration curve also showed that Logic Regression struggled to assign high, confident probabilities to true crisis years, leading to overlapping predictions and weaker crisis detection.

Gradient Boosting, on the other hand, delivered much stronger predictive accuracy. Its confusion matrix showed far better crisis detection, while still maintaining high precision for non-crisis years. With an ROC-AUC of 0.911, Gradient Boosting demonstrated a much stronger ability to distinguish crisis from non-crisis periods. The model also achieved higher precision and recall values, which is important given the rarity of crisis events in the dataset. This improvement comes from the model's sequential learning process, which allows it to capture nonlinear interactions among credit trends, reserve movements, current account volatility, and growth patterns.

The feature importance comparisons also highlight that both models rely on similar key variables such as domestic credit, GDP growth volatility, and reserve levels but Gradient Boosting interprets these relationships in a more flexible and nuanced way. Cross-validation results show that Gradient Boosting generalizes better across folds, and its calibration curve closely aligns with the ideal diagonal, meaning its probability outputs better reflect true crisis risk. Overall, Gradient Boosting proved to be the more reliable and accurate model, offering a clear advantage in early-warning detection compared to Logic Regression.

Conclusions and Future Work

After comparing and analyzing the models, some things we noticed were that Random Forest delivered the strongest overall performance with an ROC-AUC of 0.932 and clear interpretability through feature importance, and Gradient Boosting showed improved crisis detection and probability calibration, making it the most reliable early-warning model. Meanwhile, Logistic Regression provided a useful baseline but struggled with nonlinear patterns and under-detected crisis events, making it less useful of a model. K-Means identified three well-separated macroeconomic clusters, including a distinct high-risk group with elevated crisis rates, while GMM improved on this by uncovering deeper, overlapping subgroups and revealing multiple high-vulnerability clusters exceeding 40% crisis probability.

The way we summarized features during preprocessing, using five-year averages, volatilities, and trends, contributed to the success of our models, especially the random forest, and significantly improved model accuracy across all supervised methods. Additionally, combining supervised and unsupervised approaches provided consistent insights into sovereign risk patterns and strengthened the robustness of our conclusions. Next steps can include training the models to predict risk of crisis further out into the future, adding more features for further analysis, or exploring alternative models.

References

- [1] C. M. Reinhart and K. S. Rogoff, *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press, 2009.
- [2] P. Manasse and N. Roubini, "Rules of thumb for sovereign debt crises," *Journal of International Economics*, vol. 78, no. 2, pp. 192–205, 2009.
- [3] International Monetary Fund, "The IMF's Early Warning Exercise," IMF Policy Paper, 2010.
- [4] A. Ciarlone and G. Trebeschi, "Sovereign risk and macroeconomic indicators: An assessment through machine learning," *Economic Notes*, vol. 48, no. 3, pp. 1–30, 2019.
- [5] D. Barrell et al., "Evaluating the probability of financial crises: A cross-country logit approach," *Journal of Banking & Finance*, vol. 34, no. 2, pp. 276–291, 2010.

Gantt Chart

https://gtvault-my.sharepoint.com/:x/g/personal/lzhou330_gatech_edu/EYdppBAM669KqQtGfymsDucBVD_pGnIE6iE52oYdiFAiCQ?e=eNgW5K

Contribution Table

https://gtvault-my.sharepoint.com/:x/g/personal/lzhou330_gatech_edu/Edlczq6ej9INri1kMjw76doB2vZNPE90uQKyzV-u-iHNAg?e=9kbVFe