A Prediction Model of Foreign Aid Projects Funded by the United States: An Analysis Using Data from Mexico

By

Ahmed Khaled Saif Al-Shaibani

Submitted to

Central European University

Department of Economics

In partial fulfillment of the requirement for the degree of Master of Arts in Economic Policy in Global Markets

Supervisor: Professor Gábor Békés

Budapest, Hungary

2019

Abstract

The purpose of this research is to provide policy-makers with a tool to alleviate some of the burden resulting from policy making in Foreign Aid. Using data from the foreign aid projects funded by the United States in Mexico in 2017, the purpose is to predict the probability of project completion given certain characteristics. The analysis is carried out using a Logistic Regression with 5-fold cross validation. The model has an AUC of 93.10%, which means that when the model is shown two randomly selected projects, one completed and one not, 93.10% of the time it will assign a higher probability of project completion to the project that is actually completed. The findings show that a typical project that is likely to be completed is done in the area of Peace and Security with funding arriving in the second quarter of the fiscal year. Due to security reasons, the organization type is retracted in the data. The share of completed projects with such characteristics is 94%. In contrast, a project that is typically less likely to be completed is done in the Environmental sector with funding arriving in the third quarter of the fiscal year. The organization type is in the NGO sector. The share of completed projects with such characteristics is 18%. Findings generated from this analysis can potentially be used by policy-makers in prioritizing which policies to allot for projects financed by foreign aid.

Acknowledgements

I would like to express my deep gratitude to my family, my aunt Samra Al-Shaibani, and Mr. Alwan Al-Shaibani for their relentless support throughout my studies, my supervisor Professor Gábor Békés for his guidance, patience, encouragement and useful critiques in providing direction of this research work. I would also like to thank my friends and classmates Jona-Marie Ang, Benedikt Urban, and Zsófi Zádor for taking the time to provide valuable feedback on my overall thesis work.

Table of Contents

List of Figures

List of Tables

1 Introduction

1.1 Foreign Aid definition

Foreign Aid is the assistance provided to developing countries as funding to increase the development level of the country or reduce poverty. According to Britannica Encyclopedia, Foreign Aid is

> the international transfer of capital, goods, or services from a country or international organization for the benefit of the recipient country or its population. Aid can be economic, military, or emergency humanitarian (e.g., aid given following natural disasters) (“Britannica Academic, s.v. ‘Development Assistance Committee (DAC),’” n.d.)

The United States is among the biggest donors in the world, with an estimated amount of 30 Billion dollars in 2017 alone. Chief among the categories of assistance is, aid directed towards health, which accounts for approximately 23.4% of total aid spent in 2017 globally (“Foreignassistance.Gov” n.d.). This thesis will focus on foreign aid provided by the United States only, using data of the most recent, publicly available records from the United States government found on foreignassistance.gov. Figure 1 shows the total aid spent by the United States in 2017, broken-down by category.

Figure 1: Worldwide United States aid spent in 2017

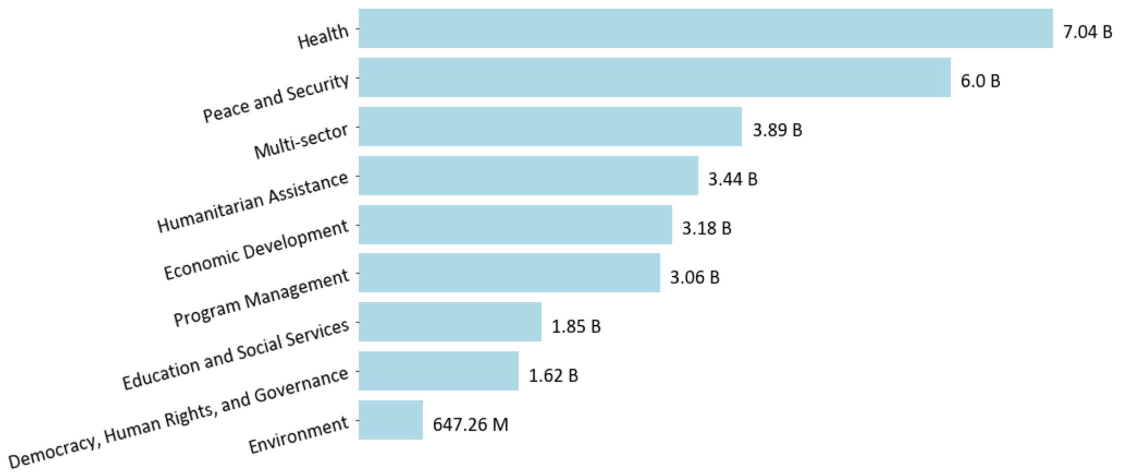

The analysis is restricted to one country only since countries differ in their political regime and culture, and thus respond differently to granted aid (Feeny 2005). This research focuses on Mexico because its shared borders with the United States renders the aligned political and security interests of the two countries, particularly in combating narcotics through the “Merida Initiative”[1] (Cook, Rush, and Seelke 2008; “Foreignassistance.Gov” n.d.). According to Cook, Rush, and Seelke (2008), the Merida Initiative would increase the number of bilateral aid projects, which means aid would be provided by the United States to Mexico. This can be seen in the data, where approximately 99% of the projects between the US and Mexico are bilateral. Having bilateral aid is vital to the analysis as it excludes other multilateral aid appropriations such as those provided by the World Bank. A full exploration of the data can be found in the Appendix

1.2 Contribution and motivation

Some researchers, such as Cordella and Dell’Ariccia (2007) and Burnside and Dollar (2000), tested the effectiveness of aid through the use of cross-country panel data. They found results that support the argument that aid is effective when good policies are present. This motivates the research of this paper to expand on current research by creating a tool for policy-makers to implement sound policies that further facilitates the effectiveness of aid. The research aims to produce a prediction model using a Logistic Regression Classifier and a Random Forest Classifier. The prediction outcome of the Logistic Regression is the probability of project completion. The Random Forest Classifier provides only classification but no predicted probabilities, so it can be used as a robustness check. Predicted probabilities serve as a tool to show the likelihood of the project being completed given certain features of the project: the foreign assistance category in which the aid is funding, the implementing organization, and

[1] The Mérida Initiative is a security cooperation agreement between the governments of Mexico and the United States aiming to address drug trafficking, money laundering, and transnational organized crime.

the fiscal quarter of the award transaction. The metric used to measure the performance of the model is the AUROC (Area Under the Receiver Operating Characteristic) Curve, or AUC for short, with 5-fold cross validation. The cross validation (CV) is implemented to reduce the risk of overfitting the model to the data. The AUC score with 5-fold CV is 93.10%. This means that when the model is to compare two randomly selected projects (one completed and one not), 93.10% of the time, it will assign a higher probability of project completion to the project that is actually completed.

The current chapter serves as an introduction and motivation to the research. Chapter 2 presents the review of literature on the topic. Chapter 3 discusses the data and its preprocessing, as well as model implementation. Chapter 4 provides the concluding remarks, policy application, and the limitations of the data.

2 Review of Literature

This chapter will present the review of literature on foreign aid. The first part will discuss briefly the history of foreign aid and the change in motivation to provide it on the global scale. This thesis does not go extensively into the history of foreign aid since the focus is on the policies derived from the model. The second section will present the different opinions about foreign aid, with some arguing for and some against it. In the rest of the paper, foreign aid is presented as good, given relevant sound policies. The third section will highlight the type of research done in foreign aid effectiveness and how this paper will contribute to the current research by taking a different direction.

2.1 History

The concept of foreign aid was introduced in the period between World War I and World War II, with the objective of financially assisting poor countries. However, the provision of aid was also politically motivated. According to LaL (1996), the current form of aid is a result of the post-World War II Marshall Plan, which was "an American economic aid program for Europe. Meant both to relieve suffering resulting from World War II's economic devastation and to contain the Soviet Union by strengthening Western Europe's ability to resist Soviet expansion, the plan was largely successful on both counts" (Archer 2019). While the Marshall Plan implementation was underway, a new frontier of war was upon the world. The Cold War between the United States and the Soviet Union changed the strategy of the US. The focus shifted towards supplying countries with aid to align the recipient countries' policies with anti-communism, thus gaining political allies (LaL 1996). After the Cold War ended, establishing political allies became less popular and it gave way to aid targeted toward humanitarian and economic development causes. Such a shift introduced many new agencies around the globe with the ones established since 1945 still operational (Harford, Hadjimichael, and Klein 2004). These new agencies rendered the share of foreign aid as high as 30 percent of the Gross National Product of a large number of developing countries (Lensink and White 1992). The increase of dependency on aid made the governments of recipient countries more liable for the allocation of policies pertaining to aid.

2.2 Effectiveness of aid and its measurements

While the expansion of aid is quite fascinating, a more essential literature on foreign assistance crystallize in the debatability of its effect on the development of recipient countries. In general, some of the voices claim that foreign assistance does in fact work and hence is favorable, while others denounce it. One of the strongest voices opposing aid is Friedman (1995), who argues in his book "Foreign Economic Aid: Means and Objectives" that foreign assistance hinders development in the sense that it does not promote democracy and does not contribute to the well-being of the citizens of the recipient countries, hence it should be eliminated. Barber and Bauer (1974) share the same argument with Friedman (1995). Other scholars further claimed that the reason aid is ineffective can be explained through proper economic incentives. Easterly (2002) claims that economic incentives have not been taken into account at the time of policy making, thus rendering aid ineffective.

Another group of scholars takes a position in support of foreign aid. According to the Hansen and Tarp (2001) analysis, the use of foreign aid leads to growth in GDP per capita based on a regression analysis from cross-country data. While Hansen and Tarp (2001) show in their research that growth rate generated from aid does not depend on sound policies, some scholars like Burnside and Dollar (2000)[2] showed that aid is effective given the recipient country has supportive policies, be it fiscal, monetary, or trade. It follows that recipient countries with inadequate policies do not benefit from foreign aid. Guillaumont and Chauvet

[2] The suggested citation is 1997, which is the date the paper was submitted as a working paper to the World Bank. Later in 2000, it was published in the American Economic Review. The most recent version is cited for comprehensiveness. The original working paper is included in the reference list.

(2001) conducted an econometric analysis on a "sample of developing countries on two pooled 12-year period." Their results indicated that in addition to the existence of good policies, environmental factors influence the level of aid effectiveness. They define environmental factors as "exogenous (mostly external)…[such as] terms of trade trend and real value of exports instability, climatic shocks, etc." (Guillaumont and Chauvet 2001).

In lieu of the aforementioned literature in favor of foreign aid, the rest of this literature review will present the different methods used to measure the effectiveness of foreign aid. However, the aim of the paper and the model is to use publicly available data to predict the success of foreign aid projects funded by the United States in Mexico. There are two reasons for limiting the research focus to one country and using a prediction model. The first and foremost important reason concentrates on the fact that countries tend to react differently depending on their political climate, culture, and level of development (or underdevelopment). Feeny (2005) asserts that omitting useful information on country-specific data by aggregating aid by country, in contrast with sector-specific data in one country, may contribute negatively to the research. The second reason is not to give a policy recommendation, but to alleviate some of the burden laid upon policy-makers. Thus, when policy-makers are equipped with the knowledge of the likelihood of the success of a project, they can implement policies with such information in mind because as shown above, sound policies have an impact on the effectiveness of foreign aid.

Burnside and Dollar (2000) examined the effectiveness of aid through the significance of policies using multiple-country panel data. They measured these policies by constructing an index of three variables based on Sachs and Warner (1997) policy indicators. The first variable measures trade openness, which is defined by Alotaibi and Mishra (2014) as "the sum of imports and exports normalized by GDP." The second variable measures monetary policy through inflation, while the third variable measures fiscal policy through budget surplus normalized to GDP (Feeny 2005). Burnside and Dollar (2000) ultimately found a positive relationship between foreign aid and good policies. Their research contributed significantly to the literature to the extent that the World Bank used their work as the main source of its 1998 report "Assessing Aid;" other established institutions, notably the British Department for International Development and the Canadian International Development Agency, similarly referenced this research (Vathis 2013). Isham and Kaufmann (1999) found similar results in their research suggesting that sound policies have a positive effect on growth of the recipient country from foreign aid.

However, the approach of Burnside and Dollar (2000) has also been criticized by scholars like Roodman (2007), who showed that the former's results are data-dependent. Other critics also point the issue with the Burnside and Dollar (2000) measurement of policy in Randen (2012). The openness measure can be represented by different variables and the results are dependent on the choice of such variables; hence, variable choice can render inconsistent results. Inflation can be caused partly by the government and partly by external factors. Additionally, inflation does not have a linear effect as a small increase in inflation is good for growth while a substantial increase can inflict harm on the economy. Using inflation to measure monetary policy is then a poor option. Finally, Randen (2012) claims that there are other measurements of development, such as privatization and market liberalization.

There are other measures of the effectiveness of aid, such as the type of aid received. Cordella and Dell'Ariccia (2007) estimated the effect of foreign aid by separating it into budget support and project aid. In their research, Cordella and Dell'Ariccia (2007) argued that budget support has more impact when awarded to wealthier states in relative terms with states that lean less towards development. They argue that poor countries are more likely to benefit from a project-type aid when such assistance is large enough because governments of these recipient countries will be less likely to use their own revenue to complete the project. In contrast, budget support would be deemed more useful where the recipient government is development-oriented such that foreign aid would grant more resources to complete projects initiated by the recipient's country government. Cordella and Dell'Ariccia (2007) suggested that such separation of measurement finds a positive effect of GDP given the presence of sound policies in recipient countries.

Various researchers also utilized disaggregated data to measure the effectiveness of aid by analyzing certain sectors. A four-year data approach was selected to examine the effectiveness of foreign aid using regression analysis. By following such approach, the risk of reaching insignificant results may be dependent on the period of the analysis as the effect of some project might take more time to be apparent. Even if the analysis is run on an extended period, there might be other factors such as noise (unrelated to the measured variables) that can distort the results (Clemens, Radelet, and Bhavnani 2008).

In all the cases presented, measuring the effectiveness of aid can be done through different metrics. Finding model specifications that can explain the effect of aid on growth has proven to be a challenging task. However, it can be seen from the literature presented in this chapter that policy influences the effectiveness of aid. Hence, the purpose of this thesis is to create a policy tool which will show the likelihood of completion of a project and thus assist policy makers in implementing policies that will contribute to the development of their respective countries. The methodology, data pre-processing, and analysis (along with policy applications) will be presented in the next chapters.

3 Methodology

In this chapter, the approach of conducting the analysis will be discussed. First, the data collection process will be briefly discussed. Second, the different techniques applied in data pre-processing will be presented. This section will only show the methods used, while further details of implementation can be found in the Appendix**Error! Reference source not found.**. The next section will show the justification for the prediction models used, their approach, and interpretation. Finally, one model will be selected as a primary prediction model.

3.1 Data Collection and Description

To perform a prediction analysis of the completion of a project funded by foreign aid, a dataset with relevant information must be acquired. The United States Agency of International Development (USAID) reports data on projects funded by the United States globally. While their data is invaluable, a new initiative in accordance with the International Aid Transparency Initiative (IATI), signed by the United States government in 2011, created a new data source that is publicly available on foreignassistance.org. IATI is a "voluntary, multi-stakeholder initiative that includes donors, partner countries, and civil society organizations whose aim is to make information about aid spending easier to access, use, and understand" ("Foreignassistance.Gov" n.d.)[3]. The foreginassistance.org website also shows that the United States government is "committed to making information on foreign assistance programs more transparent, accessible, and compatible with international standards" ("Foreignassistance.Gov" n.d.)[4]. This dataset is crucial to the analysis as it is the most recent data available and, more importantly, the data shows the completion status of each project. The analysis will attempt to develop a model to predict the completion of a project given some characteristics of the project.

[3] More information can be found on https://www.foreignassistance.gov/learn/IATI
[4] More information can be found on https://www.foreignassistance.gov/learn/about-aid-transparency

The data is limited to one time period and one country; with limitation on the time period arising from the data structure itself. The data reports approximately 95% of the projects was funded in 2017, with the rest distributed among the remaining years, namely between 2009 and 2018. The data is limited to one country as well, Mexico, because each country differs in its response due to unique factors, such as political climate, culture, and development level (Fenny 2005).[5] The final dataset contains projects implemented in Mexico and funded by the United States foreign aid in the year 2017.

The raw data contained 56 variables and approximately 25,000 observations. After limiting the data for projects done in 2017, the total number of observations was reduced to approximately 19,000. Even so, the dataset contained an ample number of observations to conduct a prediction analysis. Since the magnitude of variables was quite high and not meaningful for the analysis in many cases, certain variables were selected. These variables, along with their types, are presented below in Table 1. Further details and definitions about the variables can be found in the Appendix:

Table 1: Important Variables for analysis

Variable	Type	Unit Example
Award Status	Binary	Completion or not
Implementing Organization Type	Categorical	NGO
Foreign Assistance Aid Category	Categorical	Peace and Security
Award Transaction Fiscal Quarter	Numeric (Quarter)	1
United States Treasury Account Start of Fiscal Year	Numeric (year)	2013
United States Treasury Account End of Fiscal Year	Numeric (year)	2018
Award Transaction Value in USD	Numeric (continuous)	10,000

[5] To generate a model based on multiple countries would mean to also aggregate the data representing aid as total amounts. By using one country, the type of aid is broken down into categories, allowing for more variation. Variation provides more specific information about each project, hence a better prediction.

3.2 Descriptive Statistics

The model will attempt to predict the Award Status outcome (whether it is completed or not) using the remaining variables as predictors. The aim is to produce a tool that facilitates the policy-making process by providing policy-makers with probabilities of project completion based on some features of the project. Thus, the first step is to prepare the data before attempting any analysis or use of models, such as Logistic Regression, Random Forest Classifier, or K-nearest Neighbor, to mention a few. The variables serve as classifiers for Recipients and Funders. Recipients are classified through Implementing Organization Type, while the rest of the variables categorize Funders.

3.2.1 Implementing Organization Type

Implementing Organization Type classifies projects into a finite set of legal bodies through which the project is done. The variable has a substantial number of missing values. However, this variable is of high significance to the model as it is the only variable that identifies the recipient. Thus, it is crucial to conduct imputation optimally. Using Neural Network Supervised Machine Learning, the variable can be imputed with high confidence based on variables with high variability. The reason of using such approach for imputation stems from the data structure. Most of the variables are text-data, which is challenging to convert to numerical form. Neural Network Supervised Machine Learning is able to handle text-variables and impute missing values with high accuracy. The imputed values are predicted with an average accuracy of 86%. More details can be found in the Appendix.

The type of organization classified as “Other” is correlated with an implementing organization under the name “Miscellaneous Dollar”, which is not addressed in the original metadata. Multilateral projects were dropped as the focus of the analysis is only on bilateral aid. International, Regional, and National non-government organizations (NGOs) were merged to take the value of “NGO.” “Academic Training and Research” was dropped due to low number of observations. Finally, Figure 2 shows the distribution of implementing organization type prior to the grouping while Figure 3 shows the grouping inputted in the model. The reason for such transformation is to simplify interpretation and reduce the dimensionality of the data. A prediction model with categorical variables transforms each category into a binary indicator, since the categories do not follow a hierarchical ranking (e.g. high, medium, low).

Figure 2: Implementing Organization Type Frequency Distribution Before Grouping

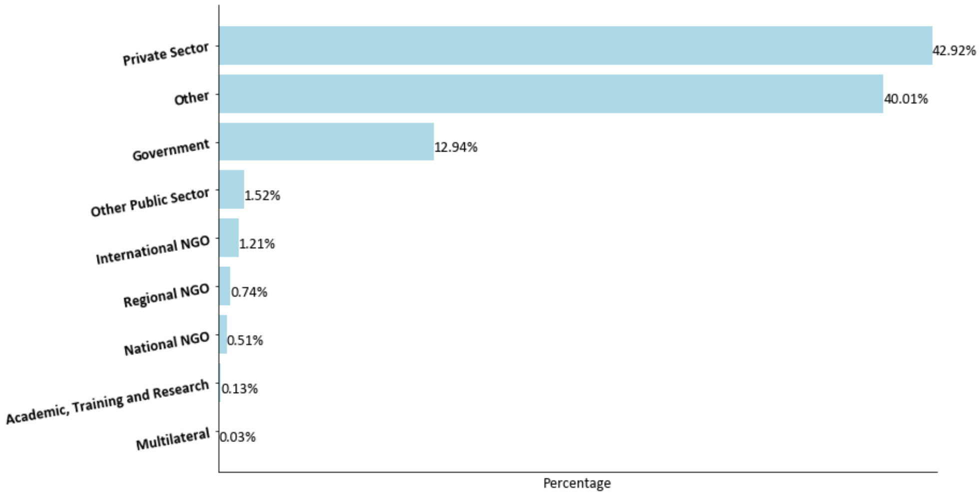

Figure 3:Implementing Organization Type Frequency Distribution After Grouping

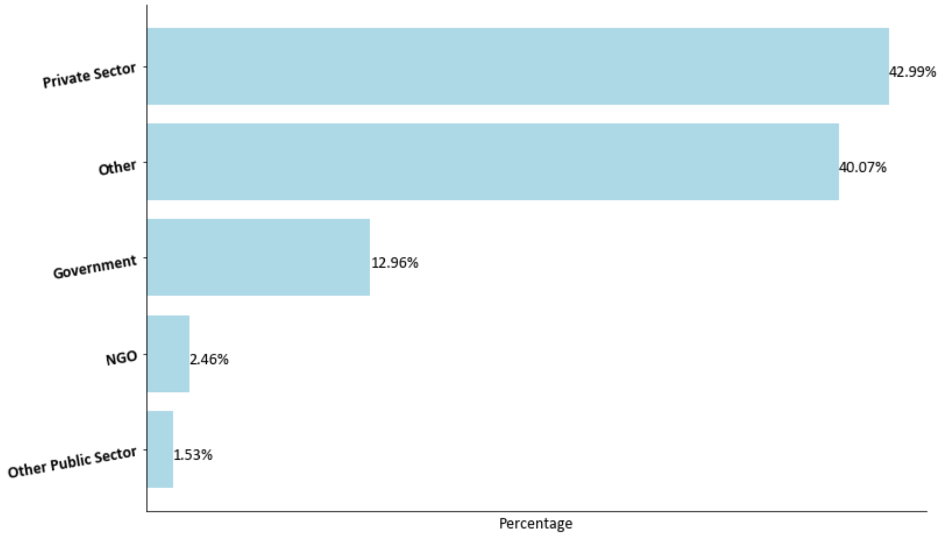

The following variables classify Funders (and the funding) into four different categories: the Foreign Assistance Aid Category, Award Transaction Fiscal Quarter, United States Treasury Accounts Fiscal Start and End of Year incurring funding, and the Award

Transaction, presented as reimbursements or obligations. The following subsections will provide descriptive statistics and figures of the aforementioned variables.

3.2.2 Foreign Assistance Aid Category

Funders are classified by the area in which the aid is allocated. Foreign Assistance Aid Category represents the category which a project falls under. The Foreign Assistance Committee decides which sector receives aid (more details and definitions are presented in the Appendix). Figure 4 shows the various categories of aid, which are also the groups reported by the Foreign Assistance Committee. Categories with a small number of observations such as Economic Development, Education and Social Services, Humanitarian Assistance, and Health, were grouped together in a new category named "Social and Economic." This grouping is done to reduce the number of groups that have a very small representation in the data. Figure 5 shows the groups after creating the Social and Economic category.

Figure 4: Foreign Assistance Aid Frequency Distribution Before Grouping

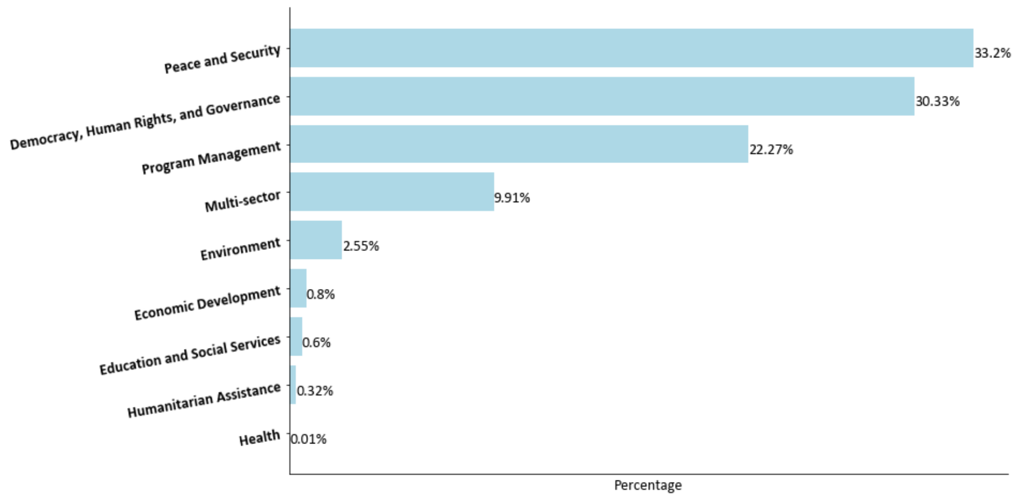

Figure 5: Foreign Assistance Aid Frequency Distribution After Grouping

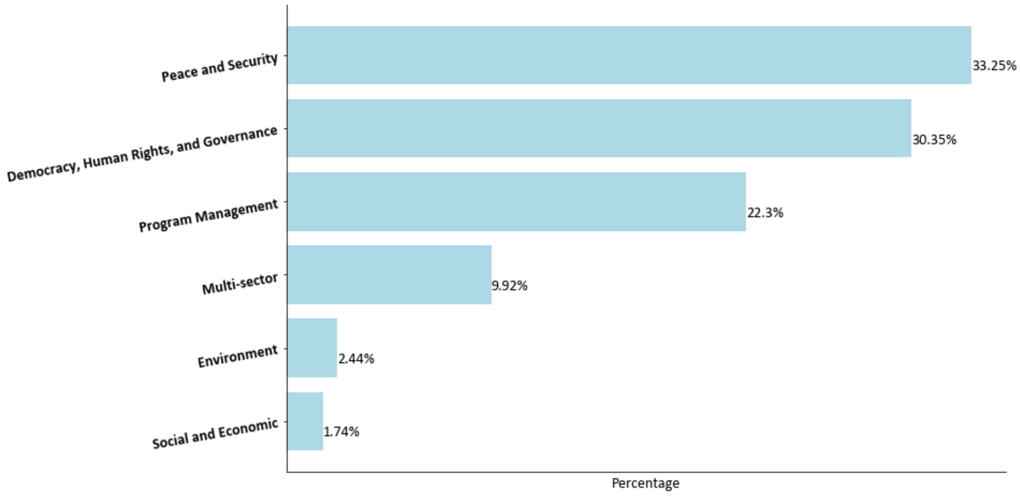

3.2.3 Award Transaction Fiscal Quarter

Award Transaction Fiscal Quarter defines the quarter in which the aid was given in that year. The variable does not require any further transformation, hence Figure 6 shows the frequency distribution of the four quarters for each observation.

Figure 6: Award Transaction Fiscal Quarter

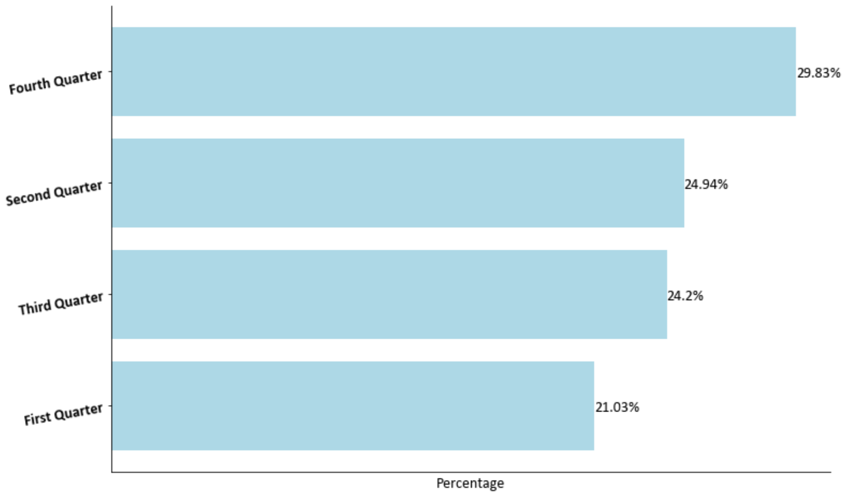

Figure 7 and Table 2 show the frequency distribution of the fiscal quarter by the status of the project.

Figure 7: Fiscal Quarter Frequency Distribution by Project Completion

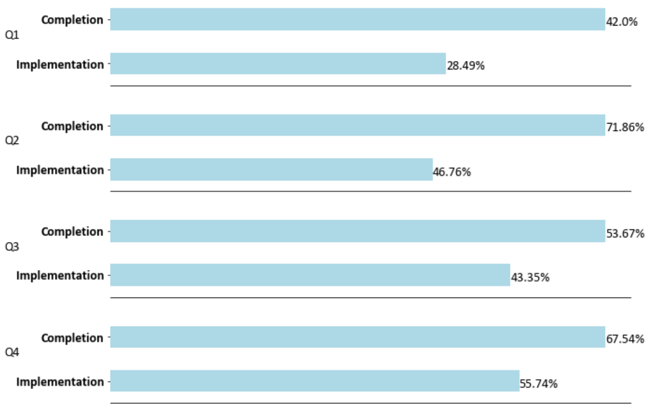

Table 2: Frequency Distribution Numbers

	Completion	Implementation	Total
Q1	2,365	1,604	3,969
Q2	2,852	1,856	4,708
Q3	2,527	2,041	4,568
Q4	3,085	2,546	5,631
Total	10,829	8,047	18,876

3.2.4 Treasury Account Fiscal Year

United States Treasury Account "Start of Fiscal Year" and "End of Fiscal Year" show when the Treasury accounts incurred costs for the project and when it stopped incurring costs. Specific dates on the start and the end of the project are missing at large because according to IATI, donors as well as recipients are not required to supply those dates, making Treasury Accounts start and end year an approximation of these dates but not a substitute. Values shown as zeros indicate that the account has no date reported. Figure 8 and Figure 9 show the frequency distribution in the data of the fiscal years in which the Treasury Account started and stopped incurring costs.

Figure 8: Frequency Distribution of Start Fiscal Year

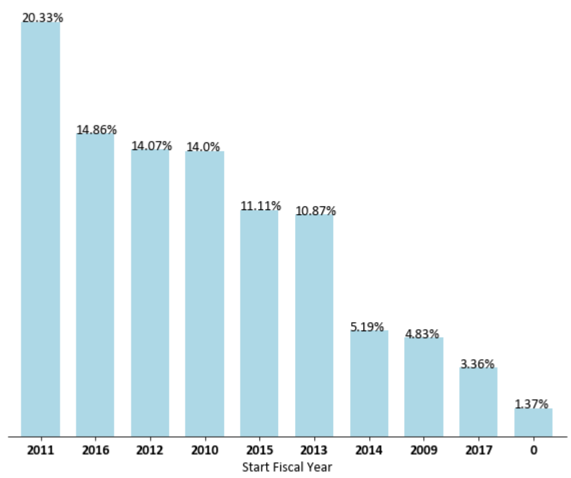

Figure 9: Frequency Distribution of End Fiscal Year

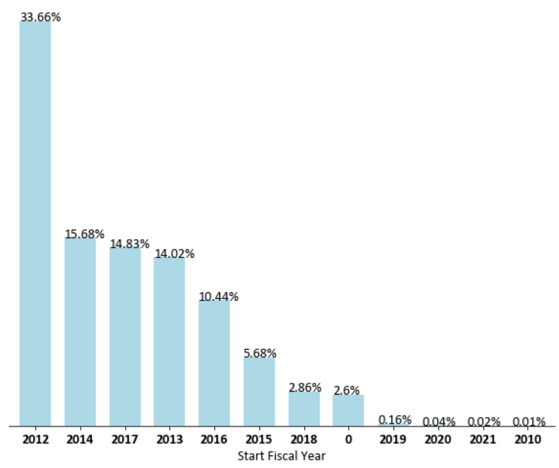

3.2.5 Award Transaction Value

The final explanatory variable is Award Transaction Value, which is the amount of aid transferred to the implementing organization to execute the projects in the recipient country. The amounts are reported in United States dollars. The values of this variable contain negative numbers. According to the official website of Foreign Assistance:

> negative obligations or disbursements result when adjustments are made in the current quarter to a previous transaction. Most commonly, a negative obligation represents a downward adjustment to an obligation made in a prior period. The downward adjustment or "de-obligation" may result from a correction to an erroneous posting made in a prior period or the cancelation of a prior award. A negative disbursement is commonly a refund or reimbursement of unused funds, or a correction to an erroneous posting made in a prior period.[6]

According to the definition, negative values contain information that is meaningful to the analysis. One interpretation of negative numbers can mean that a project with negative numbers can disrupt the project completion. One solution could be to group the observations by the reported ID. However, the reported ID could be erroneous and misleading because according to the official website of Foreign Assistance, the ID is

> [t]he unique identifier for each discrete activity found on the implementing mechanism document (e.g. grant, contract, personnel services contract, etc.). Notes: This field can contain text and numeric values. If field was submitted with blank values, a system generated ID was created.[6] ("Foreignassistance.Gov" n.d.)

This will affect the status of the project since there is no indicator of the originality of the ID (whether the ID was reported by the implementing organization or generated by the system). Hence, the ID's of projects were not used to group the projects and each project was treated as an individual project.

The values were rescaled to a logarithmic scale for two reasons. The first reason is to look at percentage changes. The second reason is to capture a better fit because there are extreme values in both the low end and the high end. Table 3, Table 4, and Table 5 show

[6] More information can be found on https://www.foreignassistance.gov/learn/faqs

summary statistics of Award Transaction Values, the positive values, and the negative values respectively. Figure 10 shows the histogram of Award Transaction and it can be seen that the distribution of values is concentrated in the middle which is due to extreme values. Dropping extreme values is not favorable as these variables may contain important information. This can be seen when the values are split into positives and negatives, then transformed into a logarithmic style. The corresponding distribution of the logarithmic transformation of positive and negative Award Transaction is shown in Figure 11 and Figure 12 respectively. These figures show an approximately normal distribution of Award Transactions.

Table 3 : Summary Statistics of Award Transaction

Metric	Value
Number of Observation	18,876
Average Value	-3,637.36
Standard Deviation	799,578.01
Minimum Value	-33,680,000.00
25th percentile	-711.41
50th percentile (median)	20.61
75th percentile	2,033.52
Maximum Value	44,489,000.00

Figure 10: Histogram of Award Transaction

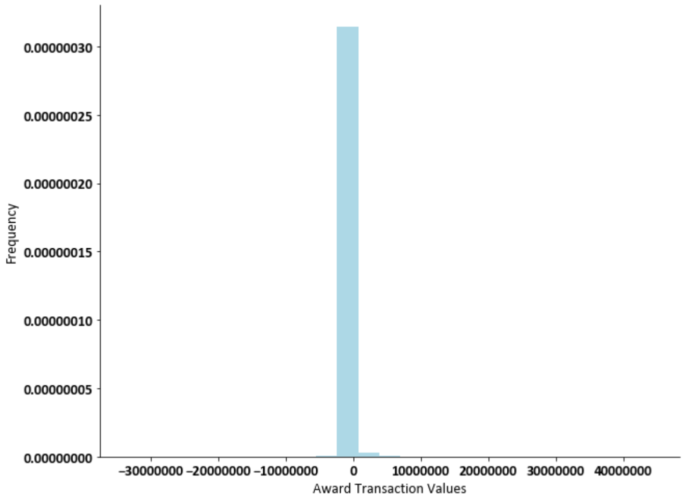

Table 4: Summary Statistics of Positive Award Transaction

Metric	Value
Number of Observation	10,871
Average Value	73,392.85
Standard Deviation	656,835.20
Minimum Value	0.03
25th percentile	88.76
50th percentile (median)	1,265.39
75th percentile	10,501.91
Maximum Value	44,489,000.00

Figure 11: Histogram of Positive Award Transaction

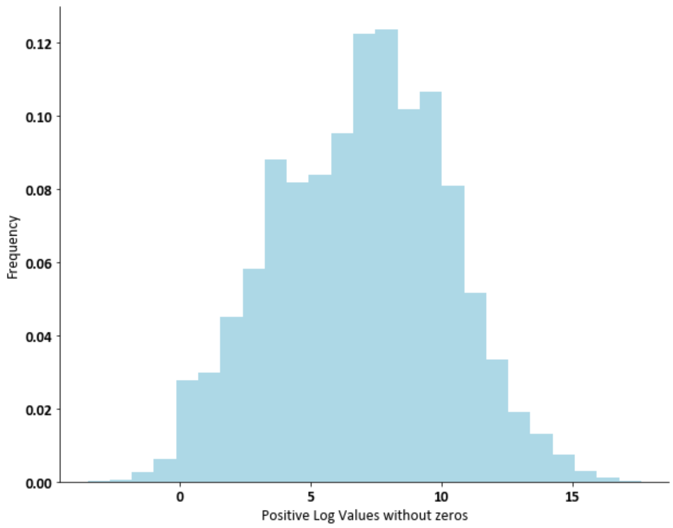

Table 5: Summary Statistics of Negative Award Transaction

Metric	Value
Number of Observation	8,123
Average Value	106,672.62
Standard Deviation	943,285.26
Minimum Value	0.010
25th percentile	122.03
50th percentile (median)	13,26.70
75th percentile	10,350.00
Maximum Value	33,680,000.00

Figure 12: Histogram of Negative Award Transaction

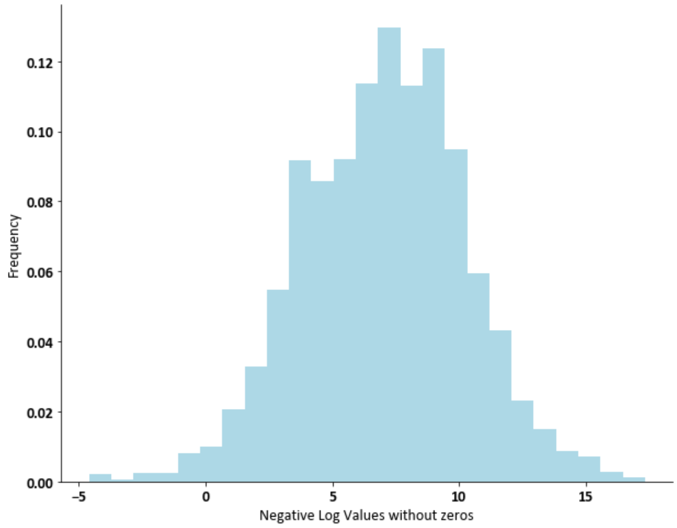

The final dataset thus has approximately 19,000 data points and 6 predictors, with no missing values. Next, we look at the different models that will be used to perform the analysis. First, a logistic regression was used to estimate the likelihood of the success; then, a Random Forest Model was done.

3.3 Modeling: Overview and Estimation

3.3.1 Models Overview: Logistic Regression vs. Random Forest Classifier

The estimated models will attempt to perform a binary classification, namely, predict the completion of foreign aid projects funded by the United States to Mexico in 2017. To reiterate, the predictions serves as a policy tool for policy-makers to use in the process of decision-making of policies pertaining to foreign aid facilitation in Mexico. This subsection will discuss the models estimated for prediction and the advantages and disadvantages of each model.

The first model estimated is the Logistic Regression (Sperandei 2014). This model estimates a likelihood measure, also known as predicted probabilities of the dependent variable taking the value of 1, which is the project completion. The prediction is done using the independent variables as features of the project. The advantage of using a Logistic Regression Classifier is the ability to see the coefficients of each feature, and more importantly, the marginal effects of each feature to the project completion on average which shows the increase or decrease in the probability of project completion for a given feature when all other features held equal. The Logistic Regression also provides predicted probabilities of the project completion, which are important for policy applications. The predicted probabilities provide the likelihood of completion of a project given certain characteristics of the project. This helps identify what projects typically have a high probability of being completed and what projects typically have a low probability of completion. Policy-makers can use such information to implement policies accordingly.

The disadvantage of the Logistic Regression is that it should be supplied with independent variables that are not highly correlated with each other, also known as multicollinearity (Ranganathan, Pramesh, and Aggarwal 2017). The Appendix shows a correlation matrix using a heatmap in Figure 15. The figure shows that the features are not highly correlated with each other on average, thus undermining the disadvantage of the Logistic Regression.

The second model estimated is a Random Forest Classifier (Breiman 2001). This model generates many decision trees that are sampled from the full set of data with different predictors and different thresholds of decision split. A Decision Tree Classifier uses the features to split the decision of classification, the project completion, based on a threshold. One disadvantage of A decision Tree is that it is prone to overfitting. A Random Forest Classifier reduces this overfitting by creating many trees randomly (Breiman 2001) (more details in section 3.3.3 below). But because of such random tree generation, the relationship between the project features and the project status is clouded, which means that the specific contribution of each feature (e.g. the type of the project or the implementing organization) to the project completion is not known. Additionally, Couronné, Probst, and Boulesteix (2018) found that the Random Forest Classifier outperforms Logistic Regression in terms of classification accuracy.

While the Random Forest can also yield predicted probabilities, they are not quite meaningful. The reason is that these predicted probabilities are coming from the different trees in the ensemble, where each tree provides a classification for a given observation (for example, tree 1 predicts completion, tree 2 predicts non-completion, and so on). These votes are the predicted probabilities and hence cannot be used because they do not come from the features of the projects but from the random trees in the ensemble generated by the Random Forest Classifier.

In summary, the predictive power of classification of a Random Forest Classifier is higher on average relative to the Logistic Regression (Couronné, Probst, and Boulesteix 2018). However, a Logistic Regression Classifier has the power to show the contribution of each feature to the project completion, whereas the Random Forest Classifier only extracts the features that were important to the Decision Trees in the splitting and not probability of project completion. In the next subsection, a Logistic Regression Estimation and a Random Forest Classifier Estimation are presented with an explanation and performance measures.

3.3.2 Logistic Regression Estimation

Logistic regression models are useful because they report coefficients through which marginal effects can be derived to show which feature (e.g. Foreign Assistance Aid Category such as Peace and Security) increases or decreases the likelihood of completing a project on average. Table 6 shows the coefficients along with the marginal effects.

Table 6: Logistic Regression Output (coefficients & marginal effects)

Outcome: Completion of Project Variables	(1) Logit coefficients	(2) Logit marginals
Foreign Assistance Category **Reference Group: Democracy, Human Rights, and Governance**		
Environment	-1.644***	-0.167***
	(0.513)	(0.051)
Multi-sector	0.286***	0.030***
	(0.098)	(0.010)
Peace and Security	0.404***	0.043***
	(0.074)	(0.008)
Program Management	-1.421***	-0.145***
	(0.124)	(0.013)
Social and Economic	-2.080***	-0.210***
	(0.422)	(0.041)
Implementing Organization Type **Reference Group: Private Sector**		
Government	0.269***	0.038***
	(0.080)	(0.012)
NGO	-2.924***	-0.313***
	(0.506)	(0.034)
Other (not specified)	3.561***	0.484***
	(0.064)	(0.007)
Other Public Sector, omitted	-	-
Award Transaction Fiscal Quarter **Reference Group: First Quarter**		
Second Quarter	0.070	0.007
	(0.075)	(0.007)
Third Quarter	0.634***	0.063***
	(0.074)	(0.007)
Fourth Quarter	0.379***	0.037***
	(0.072)	(0.007)
Constant	2.207***	
	(0.204)	
Treasury Accounts Start Year Dummies	YES	YES
Treasury Accounts End Year Dummies	YES	YES
Observations	18,575	18,575

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In Table 6, a typical project that is likely to be completed on average is conducted in the Peace and Security category with funding arriving in the second quarter of the fiscal year and by a non-specified implementing organization type. The share of completed projects with the aforementioned characteristics is 94%.

The type is not specified because it can mean that the information is retracted for security purposes. According to the official website of Foreign Assistance, projects implemented under the category of Peace and Security aim "[t]o help nations effectively establish the conditions and capacity for achieving peace, security, and stability; and for responding effectively against arising threats to national or international security and stability" ("Foreignassistance.Gov" n.d.). It can be deduced from the definition that the information about the implementing organization type is likely to be retracted.

In contrast, projects done by an NGO in the Environmental sector with funding arriving in the third quarter of the fiscal year exemplify projects that are less likely to be completed on average. The share of completed projects with the aforementioned characteristics is 18%. Given this, policy-makers can choose to direct their policy efforts towards facilitation of the work of Environmental NGOs. This can be done through "further reform and support" in many areas, including "legal framework(s) to recognize NGOs," "high-level support/endorsement from local figureheads; and engaging NGOs in policy development and implementation." ("The Role of NGOs in Tackling Environmental Issues" 2011) Furthermore, policy-makers can implement policies that promote inter-sectoral collaboration (especially between civil society organizations and local/national levels of government), capacity-building of NGOs, and increased NGO access to state information on environmental management.

Table 7 shows the confusion matrix, which reports the ratio of the number of correctly classified labels (in bold) given their true labels. Table 8 provides more details about the prediction rates, namely the sensitivity[7], specificity[8], and precision (positive predictive values)[9] of the prediction. Additional metrics for Table 9 are included in the Appendix for completeness.

Table 7: Confusion Matrix

Classified Labels	True Labels: Completion	Implementation	Total
Completion (+)	$\mathbf{9,524}$	933	10,457
Implementation (-)	1,304	$\mathbf{6,814}$	8,118
Total	10,828	7,747	$\mathbf{18,575}$

Table 8: Classification Report

Metric	Conditional probability	Rate	Predicted Label	Actual Label	Equation
Sensitivity[7]	Pr (+|completion)	87.96%	Completion	Completion	$\frac{\mathbf{9,524}}{\mathbf{18,575}}$
Specificity[8]	Pr (-|implementation)	87.96%	Implementation	Implementation	$\frac{\mathbf{6,814}}{7,747}$
Positive predictive value[9]	Pr (completion|+)	91.08%			$\frac{\mathbf{9,524}}{10,457}$
Negative predictive value[10]	Pr (implementation|-)	83.94%			$\frac{\mathbf{6,814}}{8,118}$
Correctly classified		87.96%			$\frac{\mathbf{9,524+6,814}}{\mathbf{18,575}}$

The model correctly classifies 87.96% of the cases as presented in Table 8 To further test the model, a cross-validation with five folds was implemented. Cross validation starts with a holdout method, which splits the data into a training set and a test set. The training set is used to fit the prediction model, whereas the test set serves as unseen data through which the model can predict the outcomes, in this case the project completion. This process is done five times

[7] correctly classifying true positives. Out of all the project that were completed, 87.96% were predicted as completed. Also known as Recall.

[8] correctly classifying true negatives. Out of all the project that are under implementation, 87.96% were predicted as being under implementation.

[9] Of the total projects that were predicted as completed, 91.08% of them were actually completed. Also known as Precision.

[10] Of the total projects that were predicted to be under implementation, 83.94% were actually under implementation.

with different portions of the data partitioned into a training and a test set. The main reason for using cross-validation is to reduce the risk of overfitting or underfitting the data. The average accuracy score of the 5-fold cross validation is 86.6%, which shows that the model is not overfitting the data. However, this prediction is done using a 0.5 threshold cut-off, which means that if the predicted probability of the project completion is greater than 0.5 (50%), then the project is classified as completed. In order to see how the model predicts when the threshold changes, a ROC (Receiver Operating Characteristic) curve is graphed (Hanley 1989). The ROC curve shows the trade-off between True-positives and False-positives for different thresholds between 0.0 and 1.0. The ROC curve is preferred over the Precision-Recall curve because the outcome variable, that is the completion of a project, is somewhat balanced[11]. The True-positives rate is plotted on the y-axis and the False-Positives rate (1 – Specificity) on the x-axis. The ROC curve is shown in Figure 13 below.

Figure 13: ROC Curve

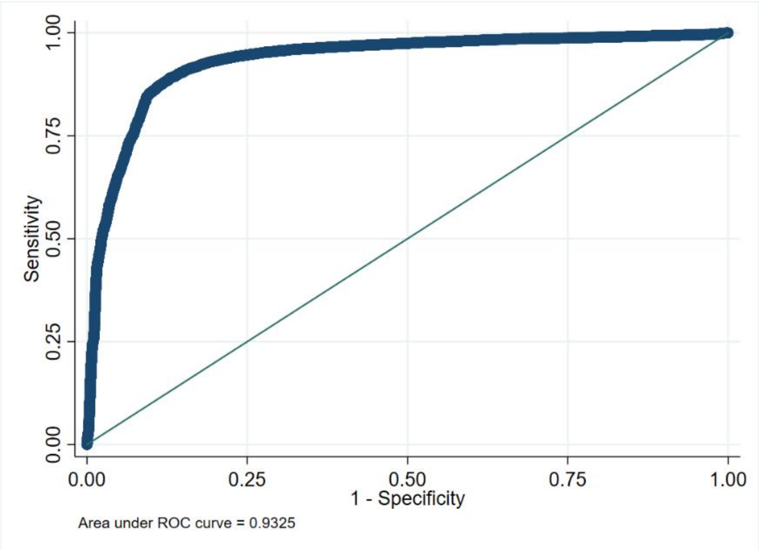

The Area Under the Curve (AUC) is of interest here. AUC shows the model's power in assigning higher probabilities of project completion to projects that are actually completed. The higher the score, the better the model is. Figure 13 indicates an AUC of 93.25%, which means

[11] In the case of imbalanced outcome (e.g. high number of completion and small number of non-completions), Precision-Recall curve is used.

that when the model is to compare two random projects (one completed and one not), 93.25% of the time, the model will assign a higher probability of the project being completed to projects that are actually completed.

Figure 14 shows the ROC curve and the average AUC after 5-fold cross validation using STATA (Luque-Fernandez, Maringe, and Nelson 2019). Average AUC is 93.10%, which means that when the model is to compare two random projects (one completed and one not), 93.10% of the time, the model will assign a higher probability of the project being completed to projects that are actually completed on average.

Figure 14: ROC Curve with 5-fold Cross Validation

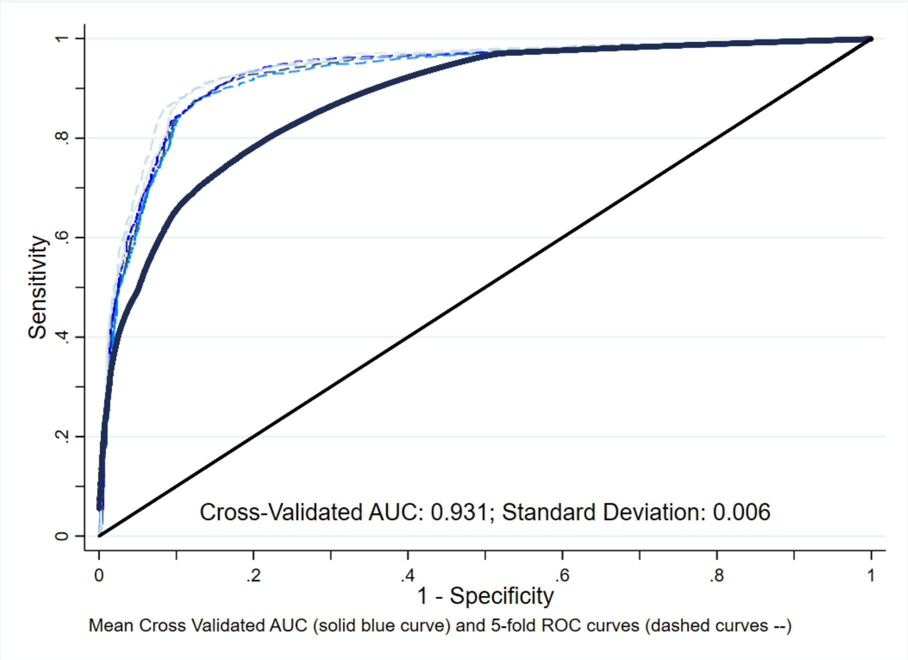

3.3.3 Random Forest Classifier Estimation

One more model is used to predict if the project is completed or not; for this purpose, a Random Forest Classifier is selected. Before making a Random Forest Classifier, a Decision Tree Classifier (Breiman 1998) is estimated. This classifier generates a set of rules based on the predictor variables. These rules are thresholds that define the target variable outcome, which is in this case the project completion. These thresholds split the decision process recursively into different nodes to arrive at a classification (completion of the project or not).

The advantage of this classifier comes from its ability to be visualized, but it suffers from overfitting issues (Jadhav and Channe 2016).

To reduce the risk of overfitting, a Random Forest Classifier (Breiman 2001) will be fitted to a subset of the data, the training set, and will create many Decision Tree Classifiers with different threshold splits and randomly selected subsamples and features[12] for prediction. The model then takes a test set to predict the outcome, which is the project completion. The process is done by taking the outcome of all the Decision Trees and storing it[13]. The decision of which outcome will be predicted is based on the majority votes of each tree. This process reduces the issue of overfitting the data that Decision Tree Classifiers have by considering many trees for the prediction. The downside of a Random Forest is that it is a black box, meaning that its process of choosing trees – where each tree receives a random subset of the training data and a random number of features – is unknown. The model is estimated to check how the data fits other models. Furthermore, the model can be used to extract the important features on average that contributed to the recursive splitting for each Decision Tree in the Random Forest (Strobl et al. 2007). However, it can inflate the importance of the continuous variables[14]. Figure 16 in the Appendix shows the top 5 important features and indeed the continuous variables are inflated. Moreover, the Logistic Regression is more reliable in showing the feature importance using the coefficients of the regression (reported in Table 6). The magnitude of each feature contributes its importance to the predicted probability in absolute terms (ignoring the sign of the coefficient) of project completion.

[12] Features are the predictors, like “Foreign Assistance Category: Peace and Security” or “Implementing Organization type: NGO.”

[13] If 100 trees are estimated, with 45 trees predicting the outcome to be project completion and the rest otherwise, the Random Forest will compare each predicted outcome with the actual outcome in the training set and will provide a classification on the majority of the classification given by each tree.

[14] These features do not correspond to their importance to project completion; they merely represent the most useful features for the trees.

The Decision Tree Classifier yields a prediction accuracy of 88% with 5-fold cross validation. As mentioned above, to avoid overfitting, a Random Forest Classifier is specified with 100 random trees. The Classifier predicts the completion of the project with 86.7% accuracy. After performing a 5-fold cross validation, the Model predicts with 86.0% accuracy compared to the Logistic Regression accuracy of 86.6%.

The Logistic Regression provides a slight improvement in prediction over the Random Forest in terms of average prediction accuracy. For the purposes of policy-making, correct classification of projects is not as important as the probability of completion of the project. The next section discusses model choice and its importance to policy-makers.

3.4 Model Choice

The model chosen after the analysis is the Logistic Regression because it has a minor improvement in average accuracy of 86.6% compared to the Random Forest accuracy of 86.0% after 5-fold cross-validation. Moreover, the Logistic Regression is able to calculate predicted probabilities of project completion where the Random Forest only provides classification. The average AUC score of the Logistic Regression is quite high at 93.10% after a 5-fold cross validation. A high score means that when the model is presented with two randomly selected projects (one completed and one not), it will assign a higher probability of the project completion to the project that is actually completed.

Predicted probabilities of project completion are more informative as they provide policy-makers with a tool that shows the likelihood of the completion of foreign aid projects funded by the United States. Policy-makers can use this tool along with other policy recommendations to implement policies that can assist in the completion of foreign aid projects. Policy-makers can also use the tool to allocate fewer policies to projects that are typically likely to be successful, since these projects are predisposed to do well even without policies facilitating their implementation.

4 Concluding Remarks and Policy Application

The current literature available on research on Foreign Aid aims to estimate the effectiveness of aid through different metrics or measurement. This thesis attempts to add to the existing literature by estimating a prediction model that calculates the probability a given project is going to completed. This approach can alleviate some of the burden on policy-makers in decision-making about passing laws or policies in the area of projects funded by Foreign Aid. Policy-makers can use the predicted probabilities provided by the predictive model, among other policy tools available, to create policies to facilitate the completion of projects that are typically less likely to be completed.

In order to carry out the abovementioned prediction model, a Logistic Regression and a Random Forest Classifier were used for the analysis. The models scored 86.6% and 86.0%, respectively, in terms of average accuracy. The Random Forest Classifier was used as well to see how the data fits other predictive models. The model selected for analysis was the Logistic Regression because it has the power to calculate predicted probabilities whereas the Random Forest Classifier does not. For policy-makers, the predicted probabilities of project completion can assist in policy-making decisions along with other policy decision making tools by providing more knowledge on the chances of completion of a new project given some of its particular characteristics. The Logistic Regression model has an average AUC of 93.10%, which means that when it is presented with two randomly selected projects, one completed and one not, it will assign a higher probability of completion 93.10% of the time to the project that is actually completed.

From the Logistic Regression results, a project that is highly likely to be completed possesses the following features: it is done in Peace and Security, receiving funding in the second quarter of the fiscal year, and implemented by a non-specified organization type. The share of completed projects with such characteristics is 94%. The “Merida Initiative,” which aims to combat narcotics in Mexico, is a good example of an agreement that falls within the scope of security and thus does not require as much attention from policy-makers to be facilitated.

The model can also show the projects with low probability of completion. According to the findings from the selected model, a typical project that is less likely to be completed is done in the Environmental sector with funding received in the third fiscal quarter and is implemented by an NGO. The share of completed projects with such features is 18%. Policy-makers can choose to direct their policy-making efforts to these projects. For example, it would be more optimal for policy-makers to allocate more policies to facilitate the work of Environmental NGOs through such measures as creating legal frameworks to improve local/national government support to the latter.

The limitations of this research stem from the size of the data. Ideally, a more comprehensive set of data on foreign aid projects would improve the predictive power of the model. It will do so with additional project features, such as the much-needed start and end dates. According to the official website of foreignassistance.org (n.d.), the data is going to be improved in the future, but no clear indication of specific improvements were communicated.

Reference List

Alotaibi, Abdullah R., and Anil V. Mishra. 2014. "Determinants of International Financial Integration of GCC Markets." In *Emerging Markets and the Global Economy*, 749–71. Elsevier. https://doi.org/10.1016/B978-0-12-411549-1.00031-4.

Archer, Stanley. 2019. "Marshall Plan." In *Salem Press Encyclopedia*. Salem Press. ers. https://ceuedu.sharepoint.com/sites/itservices/SitePages/vpn.aspx.

Barber, William J., and P. T. Bauer. 1974. "Dissent on Development: Studies and Debates in Development Economics." *American Journal of Agricultural Economics* 56 (4): 845. https://doi.org/10.2307/1239326.

Breiman, Leo, ed. 1998. *Classification and Regression Trees*. Repr. Boca Raton: Chapman & Hall [u.a.].

———. 2001. "Random Forests." *Machine Learning* 45 (1): 5–32. https://doi.org/10.1023/A:1010933404324.

"Britannica Academic, s.v. 'Development Assistance Committee (DAC),.'" n.d. In *Britannica Academic, s.v. "Development Assistance Committee (DAC),."* Accessed May 15, 2019. https://academic.eb.com/levels/collegiate/article/Development-Assistance-Committee/600894.

Burnside, Craig, and David Dollar. 2000. "Aid, Policies, and Growth." *American Economic Review* 90 (4): 847–68. https://doi.org/10.1257/aer.90.4.847.

Clemens, Michael A., Steven Radelet, and Rikhil Bhavnani. 2008. "Counting Chickens When They Hatch: The Short Term Effect of Aid on Growth." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.1112709.

Cook, Colleen W, Rebecca G Rush, and Clare Ribando Seelke. 2008. "Merida Initiative: Proposed U.S. Anticrime and Counterdrug Assistance for Mexico and Central America," March, 7.

Cordella, Tito, and Giovanni Dell'Ariccia. 2007. "Budget Support Versus Project Aid: A Theoretical Appraisal." *The Economic Journal* 117 (523): 1260–79. https://doi.org/10.1111/j.1468-0297.2007.02082.x.

Couronné, Raphael, Philipp Probst, and Anne-Laure Boulesteix. 2018. "Random Forest versus Logistic Regression: A Large-Scale Benchmark Experiment." *BMC Bioinformatics* 19 (July). https://doi.org/10.1186/s12859-018-2264-5.

Easterly, William. 2002. *The Elusive Quest for Growth: Economists' Adventures and Misadventures in the Tropics*. 1. paperback ed. Cambridge, Mass.: MIT Press.

Feeny, Simon. 2005. "The Impact of Foreign Aid on Economic Growth in Papua New Guinea." *Journal of Development Studies* 41 (6): 1092–1117. https://doi.org/10.1080/00220380500155403.

"Foreignassistance.Gov." n.d. Accessed May 14, 2019. https://www.foreignassistance.gov.

Friedman, Milton. 1995. *Foreign Economic Aid: Means and Objectives*. 1. print. Essays in Public Policy 60. Stanford: Hoover Inst. on War, Revolution and Peace, Stanford Univ.

Guillaumont, P., and L. Chauvet. 2001. "Aid and Performance: A Reassessment." *Journal of Development Studies* 37 (6): 66–92. https://doi.org/10.1080/713601083.

Hanley, J. A. 1989. "Receiver Operating Characteristic (ROC) Methodology: The State of the Art." *Critical Reviews in Diagnostic Imaging* 29 (3): 307–35.

Hansen, Henrik, and Finn Tarp. 2001. "Aid and Growth Regressions." *Journal of Development Economics* 64 (2): 547–70. https://doi.org/10.1016/S0304-3878(00)00150-4.

Harford, Tim, Bita Hadjimichael, and Michael Klein. 2004. "Aid Agency Competition : A Century of Entry, But No Exit." Viewpoint: Public Policy for the Private SectorAid Agency Competition : A Century of Entry, But No Exit. Viewpoint: Public Policy for the Private Sector. https://openknowledge.worldbank.org/handle/10986/11251.

Isham, J., and D. Kaufmann. 1999. "The Forgotten Rationale for Policy Reform: The Productivity of Investment Projects." *The Quarterly Journal of Economics* 114 (1): 149–84. https://doi.org/10.1162/003355399555972.

Jadhav, Sayali D., and Hemlata Channe. 2016. "Comparative Study of K-NN , Naive Bayes and Decision Tree Classification Techniques." In .

LaL, Deepak. 1996. "Foreign Aid: An Idea Whose Time Has Gone." *Economic Affairs* 16 (4): 9–13. https://doi.org/10.1111/j.1468-0270.1996.tb00549.x.

Lensink, Robert, and Howard White. 1992. "Is There an Aid Laffer Curve?," 42.

Luque-Fernandez, Miguel Angel, Camille Maringe, and Paul Nelson. 2019. *CVAUROC: Stata Module to Compute Cross-Validated Area Under the Curve for ROC Analysis after Predictive Modelling for Binary Outcomes*. Boston College Department of Economics. https://ideas.repec.org/c/boc/bocode/s458324.html.

Randen, Jaap Willem van. 2012. "An Assessment of the Aid Effectiveness Literature and an Analysis of Aid Instrumentation."

Ranganathan, Priya, C. S. Pramesh, and Rakesh Aggarwal. 2017. "Common Pitfalls in Statistical Analysis: Logistic Regression." *Perspectives in Clinical Research* 8 (3): 148–51. https://doi.org/10.4103/picr.PICR_87_17.

Roodman, David. 2007. "The Anarchy of Numbers: Aid, Development, and Cross-Country Empirics." *The World Bank Economic Review* 21 (2): 255–77. https://doi.org/10.1093/wber/lhm004.

Sachs, J. D., and A. M. Warner. 1997. "Sources of Slow Growth in African Economies." *Journal of African Economies* 6 (3): 335–76. https://doi.org/10.1093/oxfordjournals.jae.a020932.

Sperandei, Sandro. 2014. “Understanding Logistic Regression Analysis.” *Biochemia Medica*, 12–18. https://doi.org/10.11613/BM.2014.003.

Strobl, Carolin, Anne-Laure Boulesteix, Achim Zeileis, and Torsten Hothorn. 2007. “Bias in Random Forest Variable Importance Measures: Illustrations, Sources and a Solution.” *BMC Bioinformatics* 8 (1): 25. https://doi.org/10.1186/1471-2105-8-25.

“The Role of NGOs in Tackling Environmental Issues.” 2011. Middle East Institute. April 2011. https://www.mei.edu/publications/role-ngos-tackling-environmental-issues.

Vathis, Orestis. 2013. “Aid Effectiveness: A Literature Review,” January, 28.

Appendix

This appendix will first show additional tables and graphs related to the main body of the text. The pre-processing phase of the data used for modeling will then be addressed to predict the values of the Award Status column. First, the data will be described. Then, the award status followed by the dates will be discussed. Next, the funders will be addressed in variables “Award Transaction DAC purpose Code and Name” and “Award Transaction US Foreign Assistance Category.” Afterwards, the discussion will move on to the receivers in the variable “Implementing Organization Type.” Finally, Other features that was engineered to extract useful information but nonetheless was not quite useful on their own.

Additional Tables and graphs

Table 9: Additional Metrics for Classification Report

Metric	Conditional probability	Rate	Predicted Label	Actual Label	Equation
false-positive rate given true negative	Pr (+|implementation)	12.04%	Completion	Implementation	$\frac{933}{7,747}$
false-negative rate given true positive	Pr (-|completion)	12.04%	Implementation	Completion	$\frac{1,304}{10,828}$
false-positive rate given classified positive	Pr (implementation|+)	8.92%	Implementation	Completion	$\frac{933}{10,457}$
false-negative rate given classified negative	Pr (completion|-)	16.06%	Completion	Implementation	$\frac{1,304}{8,118}$

Figure 15: Correlation Heatmap of Features

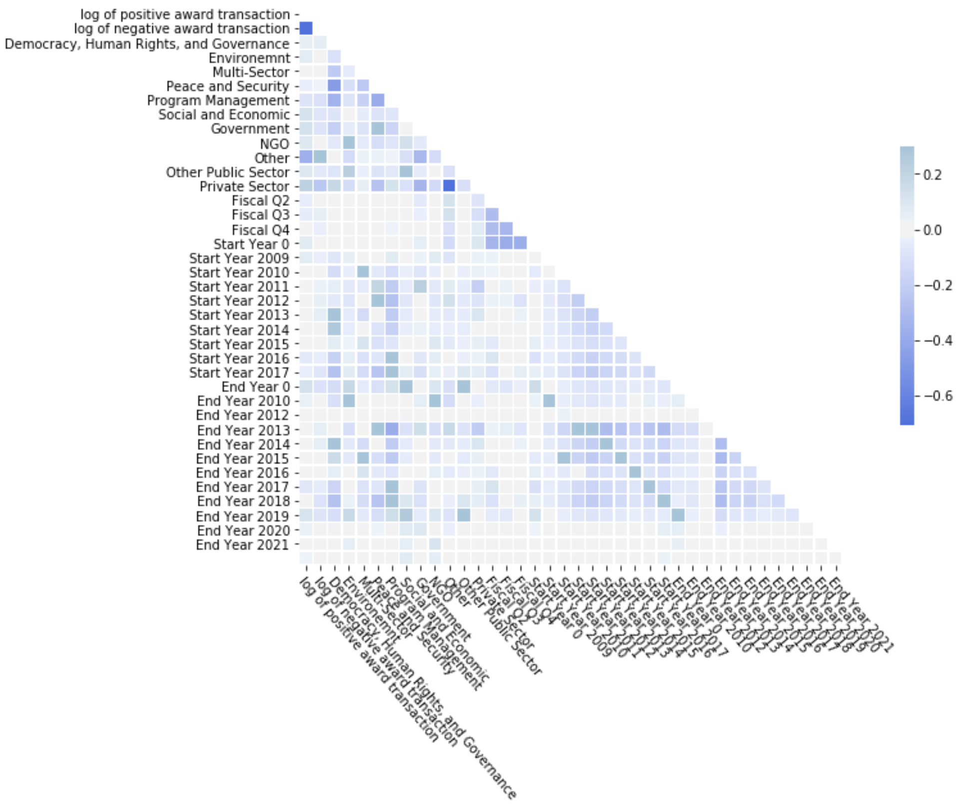

Figure 16: Feature Importance From Random Forest

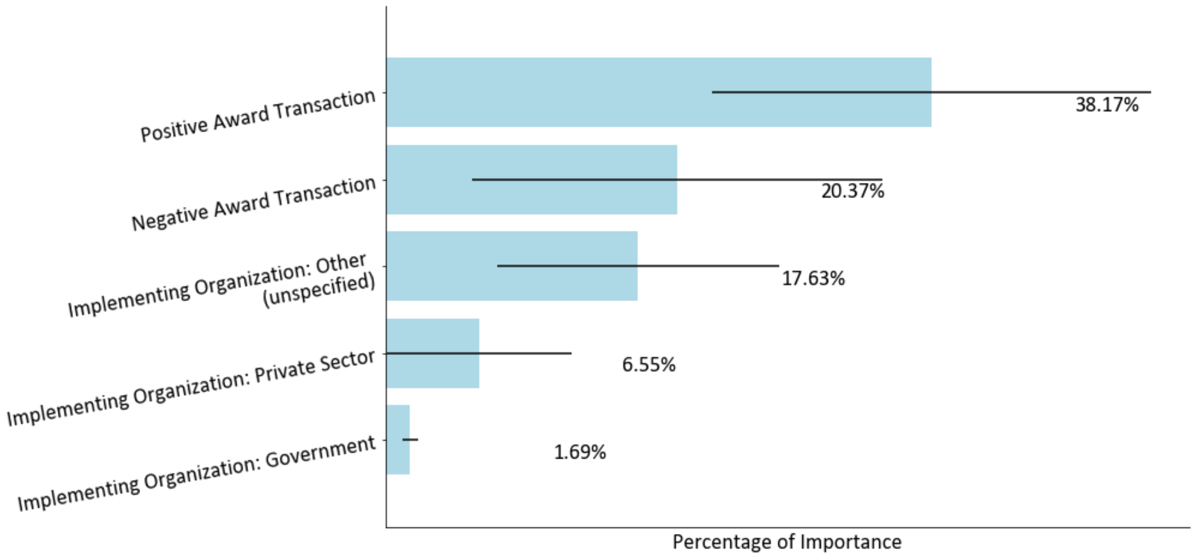

Data Exploration

The data was acquired from the official website the United States Foreign Assistance (“Foreignassistance.Gov” n.d.). The data contains awards and grants from foreign aid provided by the United States and the United States Agency for International Development (USAID) to all the countries in the world. The country of interest for this model is Mexico and the year of interest is the 2017 fiscal year. The row data contains 18,981 observations of grants and awards. Grants and awards have a unique ID represented in the dataset under the variable name “Award Identifier.” This variable has 6,924 unique IDs, which means that this ID may represent the project ID and not the award ID. The data contains different award transactions for the same ID, which could mean that these transactions are part of the same project. However, according to the data dictionary found on the Foreign Assistance official website (“Foreignassistance.Gov” n.d.), implementing organizations are partially compliant to provide an ID for the reports, while unreported IDs are generated by the system. Hence, grouping the projects by IDs is not a good practice as there is no flag for which ID is randomly generated and which are reported.

The name of the project is under the variable “Award Title.” However, this variable has many (about 7,179) missing values, and the values are not quite informative in some cases.

The table below shows the number of variables and a brief description. Subsequent discussion will feature variables of interest, including preprocessing and variation within the variable, their utility, and the method used to impute missing values (when needed).

Table 10: Variable Names and Description

Variable name	Type
Award Identifier	The award ID, which is represented as the unique number of the project. Multiple IDs are present, and they signify the number of transactions made for that project.
Extending Organization	The name of the organization or agency that received the fund.
Extending Organization Office	The name of the office or division of the agency that received the fund (mostly missing)
Accountable Organization	The name of the agency that is responsible for managing the transaction of the fund in the *Implementing organization* variable.
Accountable Organization Office	The name of the office or division of the agency that received the fund.

Implementing Organization	The name of the organization that acquired the fund from the Accountable organization to implement the foreign assistance work.
Implementing Organization Type	The type of the organization that received the fund (mostly missing) (imputed). a. Academic, Training and Research b. Government c. International NGO d. Multilateral e. National NGO f. Other Public Sector g. Private Sector h. Regional NGO
Implementing Organization Country	The country where the implementing organization is based.
Implementing Organization DUNS number	A unique 9-digit number of the implementing organization (about half is missing).
Award Title	The official title of the award if reported.
Award Description	A brief description of the implementation of the award and its impact (can sometimes be keywords)
Award Status	The stage of the award at the reported time.
Award Collaboration Type	Either Bilateral or Multilateral (mostly Bilateral)
Award Total Estimated Value	The maximum estimated value of the award (many zeros, could mean missing)
Award Inter-agency Transfer Status	Direct appropriation of funds to the implementing agency from the extending agency. When specified as 632(a), the reporting of funds is done by the implementing organization. When specified as 632(b), the reporting is done by the USAID or the Department of State.
Award Start Date	Date of starting the award identified. (Mostly Missing)
Award End Date	Date of finishing the award identified. (Mostly Missing)
Recipient Location	Location of the country in which the funds are allocated.
Award Transaction Description	The description included in the transaction of the award. The definition provided by Foreign Assistance metadata is not elaborate. Using a word count vectorizer model, keywords have been identified with frequencies (see below).
Award Transaction Value	The value of the transaction in US Dollars. Foreign Assistance publishes all the transactions that the implementing organizations report (very small to very

	large), including *zeros.* Zeros denote that the award has been budgeted but not transferred yet. *Negative Obligation* can be corrections or downward changes from the last quarter of the fiscal year or refunds. *Negative Disbursements* can be refunds or reimbursements of unused funds from the Implementing Organizations back to the Extending Organizations.[15] More details are found below.
Award Transaction Type	The status of the money transfer to the implementing organization. a) Disbursement: the transaction of the money to the implementing organization has been completed b) Commitment: the transaction of the money to the implementing organization is scheduled but has not been completed yet.
Award Transaction Date	The long format of the date of the transaction given in *(MM/DD/YYYY).* All the transaction dates are more informative when using fiscal year and quarter.
Award Transaction Fiscal Year	The Fiscal Year of the transaction. If the transaction occurred beforehand, the transaction will be reported in the accounting system in the current? fiscal year.
Award Transaction Fiscal Quarter	the transaction of the money to the implementing organization has been completed
Award Transaction Aid Type	Adopted from the International Aid Transparency Initiative, or IATI: In the dataset, the types are not reported robustly; most of the activities are reported as Project-type Interventions, which include a variety of types.
Award Transaction Tied Status	Signifies the purchasing of goods and services that will be used to conduct the project. Ranges from Tied to Untied, with Tied being only able to purchase required goods and services from the donor country(s) and Untied being the other end of the spectrum (with freedom to purchase from third parties). Most of the projects are Tied.

[15] More information can be found in https://www.foreignassistance.gov/learn/faqs

Award Transaction Flow Type	Type of the Flow, either Official Development Assistance (ODA) or Other Official Flow (OOF). OOF has been withdrawn by IATI.
Award Transaction Finance Type	Finance type as categorized by IATI. Since most of the awards are grants, the vast majority of this column is one category called "aid grants."
Award Transaction DAC purpose Code	The OECD's Development Assistance Committee categorized aid projects based on a 5-digit code. The first three digits are more general groups (or supergroups) like Education and Health, while the next two indicate the subgroups. For example, within the Education group, Higher Education would be a subgroup
Award Transaction DAC purpose Code Name	The name of the category of the 5-digit DAC purpose code
Award Transaction US Foreign Assistance Code	As reported by Foreign Assistance website: "A 4-digit code used to identify the US sector, sub-sector, and/or element that the requested funds intend to foster based on the sector framework known as the Standardized Program Structure and Definitions (SPSD). This field will be used to generate the website visualizations on the country pages."
Award Transaction US Foreign Assistance Category	As reported by Foreign Assistance website: "The description of the US Foreign Assistance Code that the requested funds intend to foster based on the sector framework known as the Standardized Program Structure and Definitions (SPSD)."
Treasury Accounts Codes	As reported by Foreign Assistance website: "Main and regular account codes assigned by the Department of Treasury, representing the account and agency to which the funds were appropriated."
Treasury Accounts Title	As reported by Foreign Assistance website: "The title of the account to which the requested funds were appropriated."
Treasury Account Start Fiscal Year	As reported by Foreign Assistance website: "Identifies the first fiscal year of availability under law that an appropriation account may incur new obligations."
Treasury Account End Fiscal Year	As reported by Foreign Assistance website: "Identifies the last fiscal year of funds availability under law that an appropriation account may incur new obligations. Note: Zero values signify "No-year (X) accounts."

Data Submission Date	Date when the data was submitted to ForeignAssistance.gov.

Award Status

Table 11 and Figure 17 shows the distribution of the current award status. "Post-completion" has been depreciated by IATI[16] and can either be dropped or merged with "Completion" but it was dropped due to low number of observations. Pipeline/identification denotes the planning phase and thus will not be used.

Table 11: Award Status Frequency Distribution

status	Frequency	Percent
Completion	10,838	57.10
Implementation	8,067	42.50
Pipeline/identification	2	0.01
Post-completion	74	0.39
Total	18,981	100.00

Figure 17: Award Status Frequency Distribution

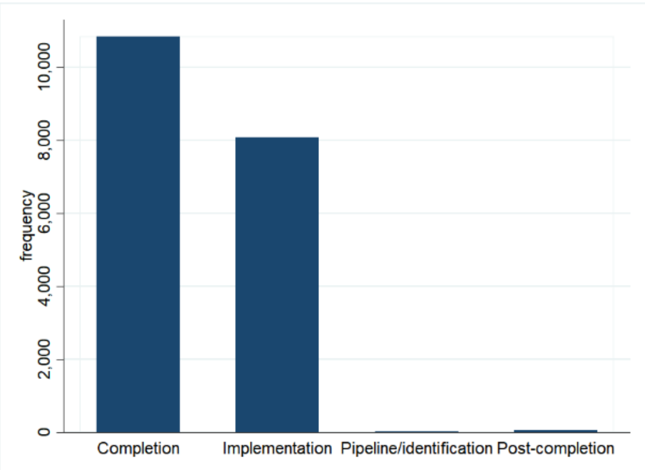

The Award Status indicates the current stage of the project. It can either be under implementation or in completion. The International Aid Transparency Initiative (IATI)[17] has not been clear on the definitions of implementation and completion. From their current

[16] https://iatistandard.org/en/

[17] More information can be found in http://reference.iatistandard.org/203/codelists/ActivityStatus/

definition, "Implementation" indicates that an activity is currently being implemented. "Completion" has been removed and replaced with "Finalization." Their definition of Finalization entails that "[p]hysical activity is complete, or the final disbursement has been made, but the activity remains open pending financial sign off or M&E"[18] (M&E stands for Monitoring and Evaluation). Based on this information, the Implementation category pertains to projects that are not completed yet while the Completion category refers to those already completed.

Dates of Award Transaction

To understand how this abovementioned information can be useful, the dates which are reported tell the story of the project life cycle. There are three variables which can be used to indicate that a project should have been completed or not. The Foreign Assistance Data[19] reports the starting and ending dates of the award. The official website of Foreign Assistance published the IATI implementation plan[20], which seeks to report all the information about foreign assistance provided by the United States foreign aid agencies to the public domain. The published implementation schedule for data collection provided by the International Aid Transparency Initiative (IATI) (refer to Excel file, tab "Activity Data", cells 29C to 29I) dictates that Extending Organizations are *"partially compliant"*[21] to publish the start and end date of the activity. The dates will provide a better understanding of the interval of the start and end date of the project.

The dates of the treasury accounts can be utilized to get an idea about the interval of the projects. The data contains for each activity the Treasury Account Start Fiscal Year and

[18] See footnote 17

[19] More information can be found in www.foreignassistance.gov

[20] More information can be found in https://www.foreignassistance.gov/learn/IATI/

[21] More information can be found in https://www.foreignassistance.gov/assets/iati/IATI%20Implementation%20Schedule_Final.xlsx

Treasury Account End Fiscal Year. According to the Foreign Assistance official website ("Foreignassistance.Gov" n.d.), these dates identify the first and the last fiscal year of "availability under law that an appropriation account may incur new obligations" ("Foreignassistance.Gov" n.d.). According to their definition, this can translate to the period when an account can provide money to a project. Furthermore, the Data reports the Award transaction fiscal year and quarter as well as the full date for each transaction (usually 3 to 6 dates per year).

Funders

In the examination of variables, the extending organization and the accountable organization have low variance, with "U.S Department of State" as the main funder. The extending organization office has a substantial amount of missing values. For those reasons, the variables above are not meaningful. The implementing organization variable suffers from the same symptom, which is low variance.

Award Transaction DAC purpose Code and Name

Since it is not possible to group extending organizations and offices into meaningful groups due to low variance, an alternative identification must be found. A good approach is to use the DAC code to group the activities based on the category for which the funding is allocated. In the data, each DAC code consists of a 5-digit number. The first three represent the super category of funding, like Education, Health, and others[22].

The first three digits were extracted, and each 3-digit DAC code was assigned to a group. The resulting groups were saved in a new variable named "dac5_group" and their distribution are presented below:

[22] More information can be found in http://www.oecd.org/dac/stats/documentupload/CRS_BI_VOLUNTARY_purpose_codes2016flows_en_July17.pdf

Table 12: Frequency Distribution of DAC 5-digit code groups

DAC 5-digit groups	Frequency	Percent
Administrative Costs	4,130	21.76
Agriculture	16	0.08
Banking & Finance	4	0.02
Development	2,022	10.65
Disaster Aid	1	0.01
Education	113	0.60
Emergency Response	1	0.01
Environmental	484	2.55
General Budget Support	960	5.06
Government & Civil Society	8,648	45.56
Health	2	0.01
Industrial	6	0.03
Other Economic Infrastructure	17	0.09
Other Social Infrastructure	2,568	13.53
Transport	9	0.05
Total	18,981	100.00

Figure 18: Frequency Distribution of DAC 5-digit code groups

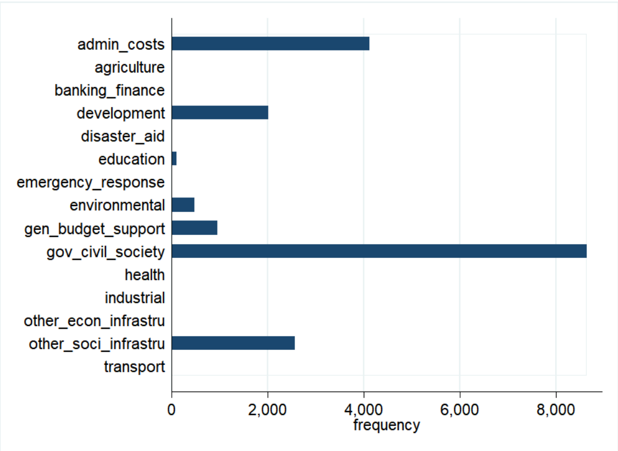

Table 12 and Figure 18 show that the frequency of each variable is not equally distributed. Based on this, the frequencies do not serve the purpose as one of the major categories is "Other Social Infrastructure." After some investigation on other data columns, it appeared that "Other Social Infrastructure" is highly associated with "counter narcotics" and thus has been renamed accordingly. The next step was to investigate other groups that have very few observations, namely: Education, Health, Industrial, Agriculture, and Economic. Most of these activities are related to the development of the country when financed by foreign aid, hence they were merged with the Development category. The final result of the categorization is displayed in Figure 19 and Table 13 (Humanitarian Aid was later dropped due to the small number of observations):

Table 13: DAC Groups Frequency Distribution

DAC group	Frequency	Percent
Administrative Costs	4,130	21.76
Counter Narcotics	2,568	13.53
Development	2,189	11.53
Environmental	484	2.55
General Budget Support	960	5.06
Government & Civil Society	8,648	45.56
Humanitarian aid	2	0.01
Total	18,981	100.00

Figure 19: DAC Groups Frequency Distribution

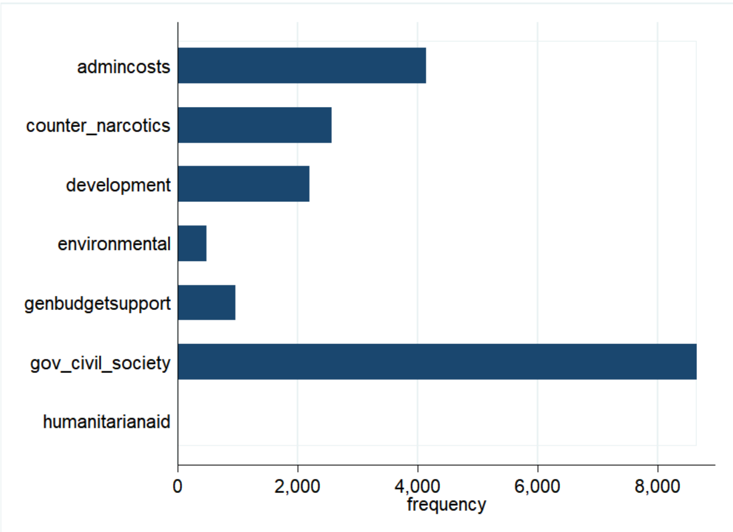

Award Transaction US Foreign Assistance Category

There is another method for categorizing the projects. The US Foreign Assistance website provides a similar grouping of the data which is displayed in Table 14 and Figure 20. Their methodology of categorization including supergroups (called categories) and specific groups (called sectors) is different but overlaps with the DAC's OECD categorization.

Table 14: Foreign Assistance Category Frequency Distribution

Award Transaction US Foreign Assistance Category	Frequency	Percent

Democracy, Human Rights, and Governance	5,734	30.21
Economic Development	154	0.81
Education and Social Services	114	0.60
Environment	483	2.54
Health	2	0.01
Humanitarian Assistance	61	0.32
Multi-sector	1,944	10.24
Peace and Security	6,279	33.08
Program Management	4,210	22.18
Total	18,981	100.00

Figure 20: Foreign Assistance Category Frequency Distribution

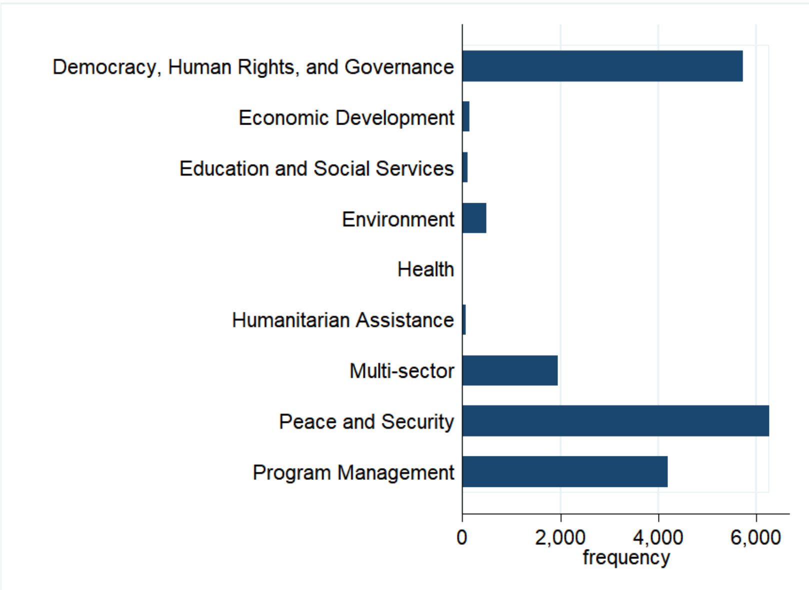

When running the models, high collinearity was found between the two variables, more specifically, between observation in "counter narcotics" category from the "DAC group" variable and "Peace and Security" category from the Foreign Assistance variable This is the reason why only the Award Transaction US Foreign Assistance Category was used.

Receivers

Implementing Organization Type

The receiver of the award is only represented in the dataset by the implementing organization. This variable has much more variance because it consists of all the legal bodies that received the funds to carry out the project. For the purposes of the analysis, it can create an overfitting issue and thus it needs to be categorized. One variable of interest to conduct such categorization is the Implementing Organization Type, but it has many missing values and may not be very useful. The Implementing Organization Type can be imputed nonetheless using Neural Networks machine learning algorithm via Python. This process uses columns with few or no missing data to create features for the non-missing observation of the Implementing Organization Type. The columns which are fed to the algorithm to extract features are:

1. Implementing Organization
2. Accountable Organization Office
3. Award Description
4. Award Transaction Description
5. Award Transaction Value
6. DAC 5-digit Groups (from the earlier step)
7. Treasury Account Title
8. Award Transaction DAC Purpose Code Name
9. Award Status
10. Award Transaction US Foreign Assistance Sector
11. Award Transaction US Foreign Assistance Category

The next step is to impute the missing data of the Implementing Organization Type column using the loaded features in the algorithm. The algorithm reports the accuracy (between 0.0 and 1.0) of each predicted imputation in a separate column. The average accuracy is 0.865 or 86.5%. An issue that was present during the process was one unique value *"MISCELLANEOUS DOLLARS,"* with 7,564 observations (39.87%) and no corresponding type. It is, however, highly associated with observations that have the value of *"Bureau of International Narcotics And Law Enforcement Affairs"* in the Implementing Organization Office variable. For this reason, it was dropped from the algorithm but not from the data as it was not clear what type of category it belonged to and therefore was categorized as other. Note that it cannot be categorized as “Other Public Sector” because this category is highly associated with Peace Corps, which is a volunteer-based initiative.

Other Features

In order to further understand the awards and create more features for the prediction analysis whenever possible, the categories of Award Description and the Award Transaction description were used. These two variables do not have missing observations and may thus contain meaningful information. Due to the sheer amount of reading, a word count vectorizer was used to extract the most used single word, 2-word phrases, and 3-word phrases in both the award description and the award transaction description. The figures below show the frequency distribution of the most common words.

Award Description

Figure 21, Figure 22, and Figure 23 show that the Award Description can be categorized into demand reduction (of drugs as it is correlated with the counter-narcotics), “program development”, and “DEA training”, while the rest can be categorized as “other.” However, this categorization was not used in the main model because it cannot be interpreted in a meaningful way. The categorization was used as part of the imputation process of the implementing organization type.

Figure 21: Award Transaction Uni-grams

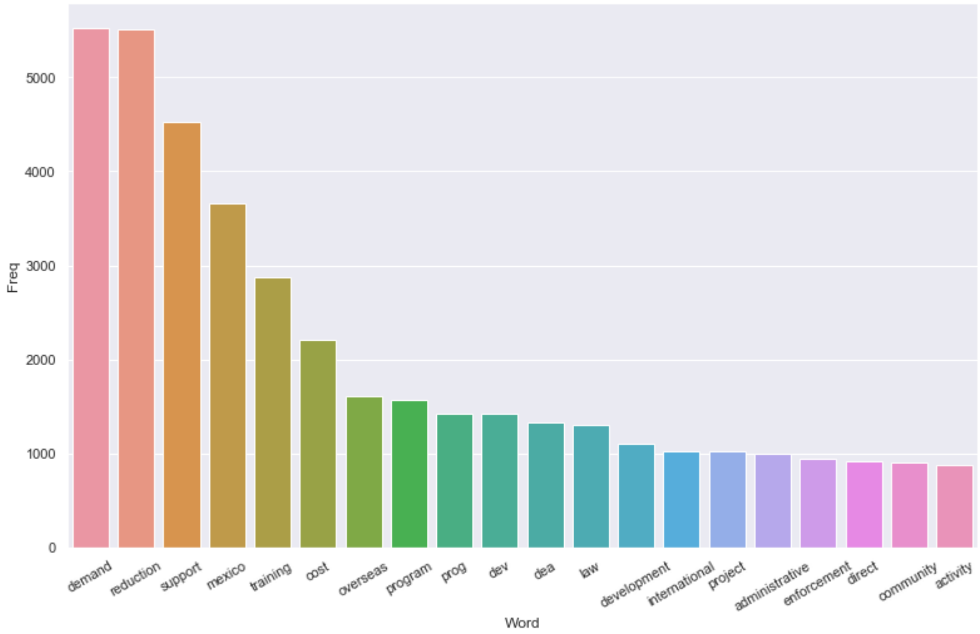

Figure 22: Award Transaction Bi-grams

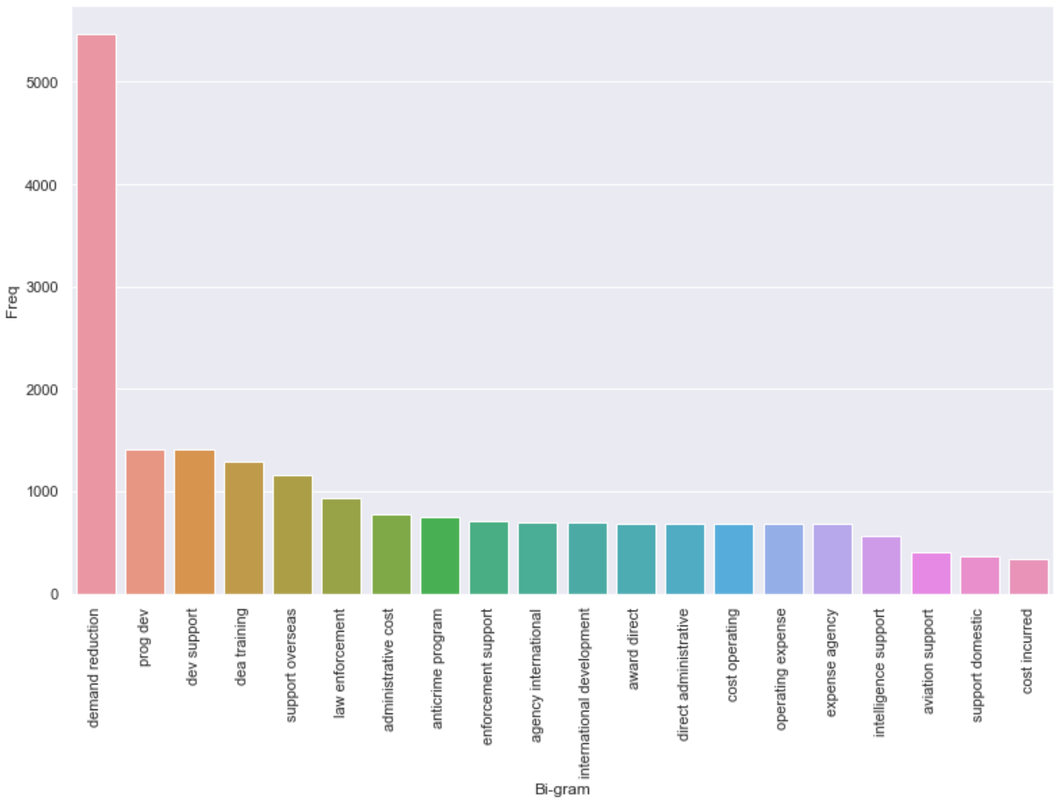

Figure 23: Award Transaction Tri-grams

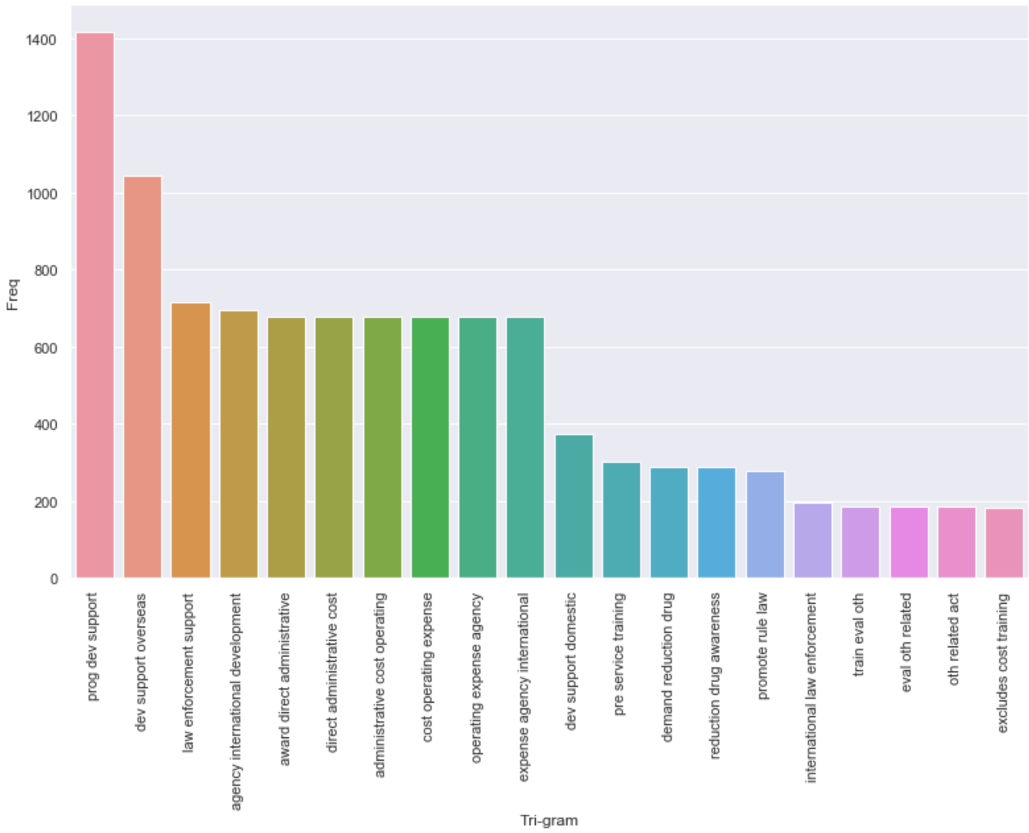

Award Transaction Description

Figure 24: Award Transaction Description Uni-grams

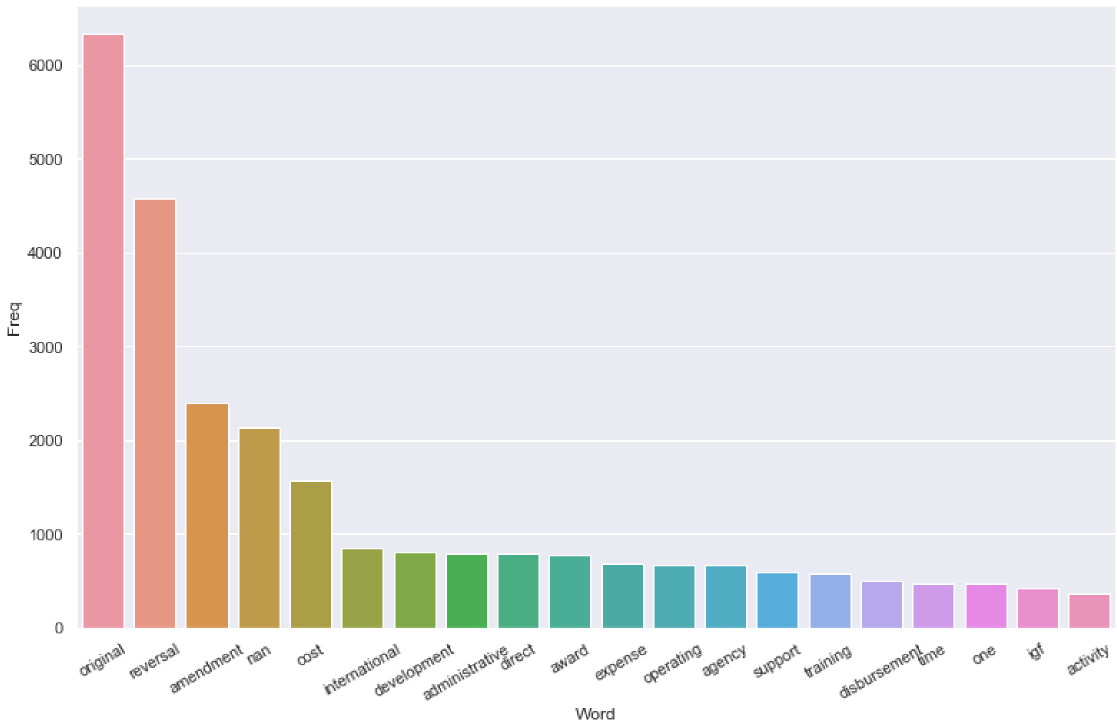

Figure 25: Award Transaction Description Bi-grams

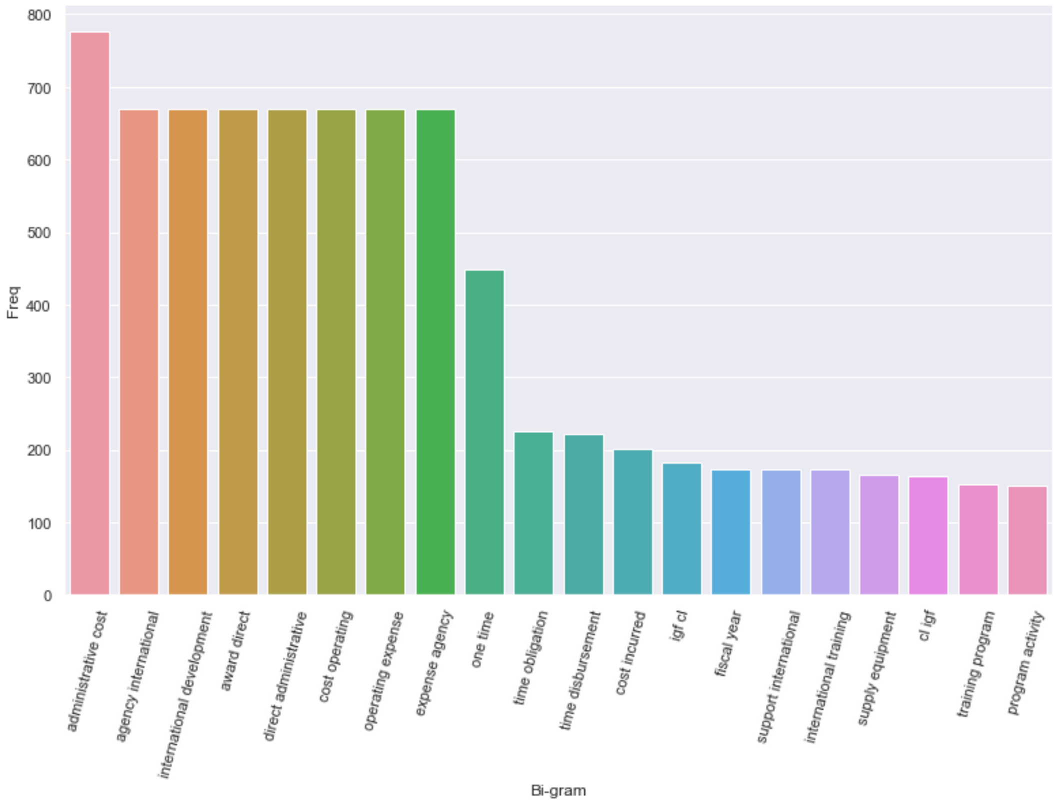

Figure 26: Award Transaction Description Tri-grams

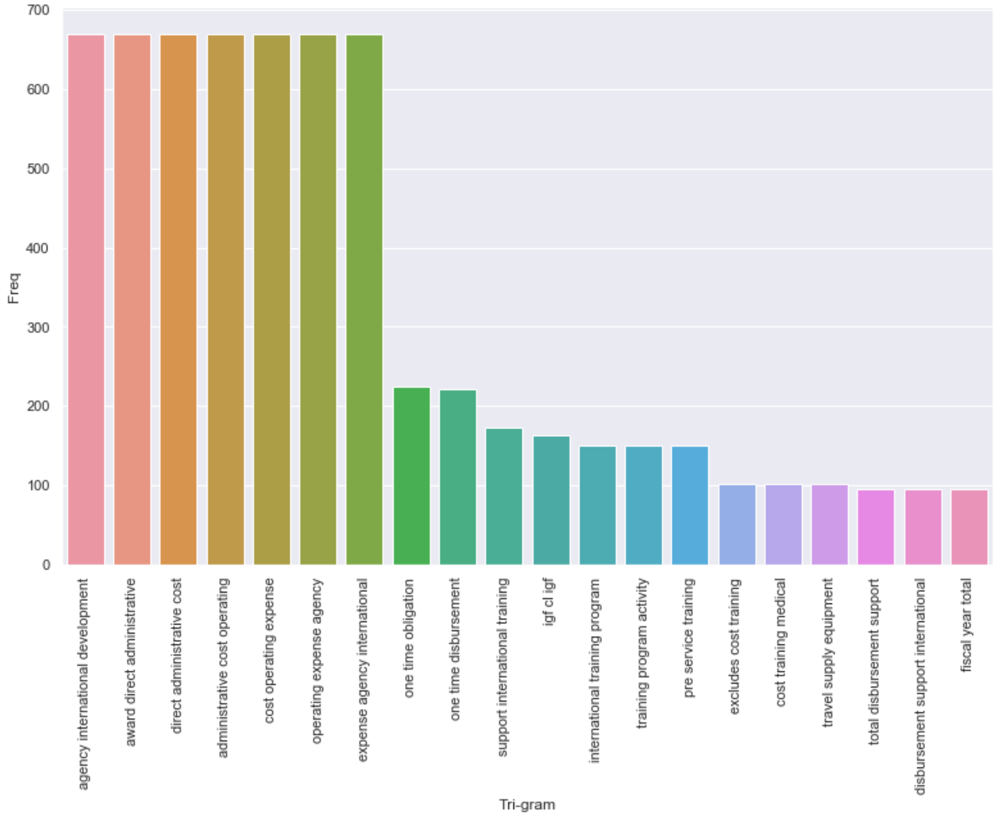

In Figure 24, Figure 25, and Figure 26, the Award Transaction Description has "original", "reversal", "amendment", "administrative costs", and "agency international development" as the most frequent words and phrases. Reversal was highly correlated with the outcome variable and hence was not included in the model. The categorization was used as part of the imputation process of the Implementing Organization Type.

It is worth noting that the Award Title has 7,179 missing observations and 2,750 unique observations. Using the same method of analysis used above, we get the word count. They are, however, not very meaningful in the sense they do not provide enough information for categorization of the data point and so the graphs are not reported here.