Team Cedar - Executive Summary

Project Description and Goals

Our project focused on analyzing a dataset of claim billing data for prescription medications.

The goals for our project were twofold:

1. Predict whether a prescription medication request would be accepted or rejected by insurance.
2. If accepted, predict the amount of money the patient will likely be required to pay for the prescription.

Stakeholders for this project include CoverMyMeds, as well as patients and clinicians. By being able to know whether a medication is likely to be accepted and how much the patient will have to pay, medical professionals can make informed decisions at the point of prescription, especially when multiple medications may be available.

Goal 1: Predicting likelihood of acceptance/rejection

We used the XGBClassifier from xgboost to predict whether a medication will be accepted or rejected by insurance. Our model did not perform very well at the default 0.5 threshold. We, therefore, explored other thresholds that would minimize false negatives, as these are what would be most costly for the patients. We decided on a threshold of 0.17 as the optimal trade-off between specificity and recall.

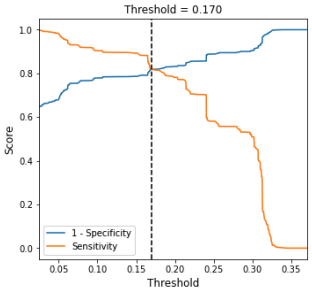

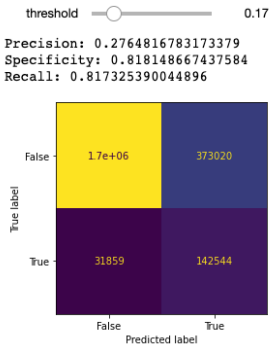

Goal 2: Predicting required patient pay

We used the XGBRegressor model from xgboost to predict how much patients will have to pay for their medications. Our model performed well and predicted co-pay accurately to within a little over $1 on average. Mean patient pay: $26.10 Baseline mean absolute error: $21.80 Model mean absolute error: $1.17.

Future steps

As a next step we would like to build a tool that allows providers to see likelihood of acceptance/rejection, cost of the medication to the patient, and potential alternative prescriptions for the diagnosis with better odds of being accepted and/or a lower cost to the patient.