BUSINESS SOLUTIONS PROJECT

Community FOOD PANTRY

SLOATSBURG • SUFFERN • HILLBURN • TUXEDO

JOY YIN AND AYUSHI MEHROTRA

Suffern High School

49 Viola Road, Suffern, NY 10901

March 2026

CONTENTS

I. EXECUTIVE SUMMARY

ORGANIZATION BACKGROUND

The **Community Food Pantry** is a "client-choice" food pantry in Rockland County, NY. Founded in 1990, the pantry serves Sloatsburg, Suffern, Tuxedo, and Hillburn, providing essentials such as fresh produce, cereal, pasta, rice, and frozen meat to families in need. Staff order this produce biweekly from the Regional Food Bank of Northeastern New York. In an area with a high immigrant population, population surges and recent economic and demographic changes have led to an increase in the number of visitors and frequent guests. In recent years, the organization has moved locations to accommodate this increase and has undergone a rebranding to represent the diverse community it serves.

*in 2024

13,367	**43,352**	**390,168**
Families Served	Individuals Served	Meals Provided

PROBLEM

Pantry demand is volatile and depends on factors such as weather and time of year. Thus, staff have had trouble predicting the amount of food to sufficiently serve all families in need. Recently, both stockouts, when groceries run out before the ordering cycle is over, leaving some families without enough items, and leftover products, which expire, wasting important food and funds, are increasingly common. Moreover, volunteer count is often inconsistent, either over- or understaffed.

SOLUTION: DEMAND PREDICTION MODEL

A machine learning model predicting the number of households and individuals served per distribution cycle will ensure the pantry can make informed decisions about its orders. The model will estimate the total amount of food needed (lbs), including a category breakdown and ideal volunteer count. After inputting factors and characteristics of the cycle, a simple dashboard will display the predictions and a weekly "Prep Plan."

MAIN GOAL Ensure that at least 80% of distribution cycles maintain adequate stock through at least 75% of the cycle, while keeping food waste under 30 lbs

BENCHMARKS

Forecast Accuracy

The first benchmark is accuracy of our model; effective code is needed to reliably predict the attendance of clients during a specific distribution cycle. Our model will accurately estimate the amount of food and volunteers needed and present the results in a comprehensive way to pantry staff.

Resource Efficiency

From excess food to oversaturated volunteer logs, our project is also aiming to reduce inefficiencies and waste in the Community Food Pantry. The predictive model aims to predict the exact amount of groceries needed to be ordered and the number of volunteers needed to serve the estimated turnout, reducing food waste and volunteer overstaffing, keeping the pantry running smoothly and efficiently.

MANAGEMENT TEAM


Joy Yin and Ayushi Mehrotra
Project Managers

Elizabeth Takacs
HR and Operations Coordinator
Provided the necessary Community Food Pantry historical data, final approval of proposal

Marisol Roman
Director
Primary and initial contact with the pantry, aided with implementation

Terry Deasy
Food Pantry and Recovery Coordinator
Ensured accurate data collection and inventory logs

KEY METRICS

- Longevity of Orders during Cycles of Distribution
- Amount of Perishable Food Waste
- Predictive Model's Forecast Accuracy
- Volunteer Efficiency and Utilization

RISK MANAGEMENT

- Data Quality and Consistency
- Model Error and Demand Surges
- Inventory and Supply Variability
- Staff Adoption and Utilization
- Client Privacy and Data Protection

PROPOSED BUDGET

Category	Cost
Hosting & Deployment (6 months)	$30
Backups & Security (6 months)	$120 ($20/mo)
Staff Training & Handoff	$40
Contingency	$100
TOTAL	**$290**

TIMELINE

C
Create plan & obtain approval
May 20th to June 10th

A
Analyze data & create ML model
June 10th to July 12th

R
Run ML model & prepare for implementation
July 12th to July 30th

E
Execute plans & **e**valuate impact
July 30th to Dec 7th

MONITORING & CONTROLLING

MONITORING	CONTROLLING
Turnout, distribution, and waste were logged to compare predicted vs. actual outcomes.	Two-Tier Forecasting: standard plan for a typical week and Surge Plan to handle demand spikes
Unusual conditions were flagged for updates.	Model was refined with adjusted calculations and visuals for staff clarity
Staff feedback was collected to ensure the dashboard and Pantry Prep Plan were accessible.	Missing or inconsistent data was flagged and corrected

KEY FINDINGS

- **Highly accurate forecasts:** Model achieved 1.32% error across 9 live cycles
- **Stockouts reduced:** Before model, 41.7% of pantry cycles ran out of food. With the model, only the first live cycle had a full stockout; after retraining and adding buffers, 0/8 later cycles stocked out.
- **Food waste controlled:** In all 9 model cycles, perishable waste stayed under 30 lbs, with most cycles seeing only 10–20 lbs of unavoidable waste.
- **Volunteer use improved:** Households served per volunteer-hour increased by 30%, and days rated over-staffed/under-staffed decreased by 61.9%.
- **Staff-friendly system:** Dashboard + weekly "Pantry Prep Plan" gave a simple ordering and staffing guide, saving staff planning time and increasing confidence for both normal and high-demand weeks.

RECOMMENDATIONS

Expand Scope

- Create a model applicable to multiple food pantries or regions
- Create predictive model for other charities and causes

Community Engagement and Feedback

- Obtain feedback and suggestions from community members and volunteers

LESSONS LEARNED

Understand Staff Needs

- Volunteer recommendation predictions, comprehensive dashboard, Pantry Prep Plan
- **Continuous Improvement**
- Initial struggles with underestimation of demand, recent distribution cycles added to training data, improvement in accuracy recently

II. INITIATING

Statement of Problem

What is the client-choice model?

In client-choice foods pantries, such as the Community Food Pantry, clients select their own groceries like a grocery store instead of receiving pre-packed bags, promoting dignity and reducing food waste.

A food pantry serving Rockland County, NY, the Community Food Pantry offers food and hope to struggling members of the community. A 501(c)(3) not-for-profit organization, it allows clients to select their groceries based on their preferences and available selections. Since the early 1990s, this "client-choice" model has been successful. Recently, however, the Pantry's clientele has recently expanded dramatically, from a reported number of 5,189 households in 2020 to 13,367 in 2024. This is likely as a result of demographic changes in the area and widespread economic changes. Staff have reported that this client increase has caused the number of visitors per distribution cycle to become difficult to predict, especially as attendance and turnout are volatile and depend on multiple confounding factors, including weather and time of year. Moreover, volunteers are often unevenly distributed, with some days or weeks having too many and others being understaffed. The Community Food Pantry has experienced both stockouts and excessive products, especially in the past two years. Their biweekly orders from the Regional Food Bank of Northeastern New York have often not been able to accurately represent the needs of their clients, leading to some families without groceries or food being wasted. Given this unpredictability, from January through June 2025, the pantry ran out of food before the cycle was at least 75% complete during five distribution cycles and wasted at least ten lbs of food in three cycles out of twelve total cycles.

157.6%
increase in clients from 2020 to 2025

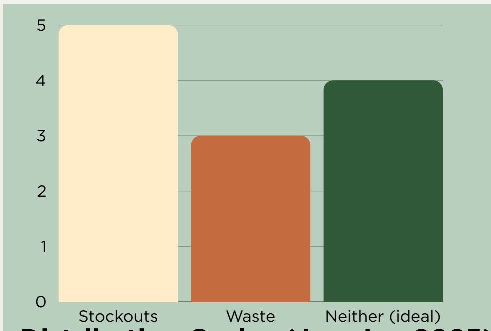

Distribution Cycles (Jan–Jun 2025)

Out of 12 distribution cycles in early 2025, five were stockouts, three had leftovers, and four were balanced. Balanced cycles are ideal. Some results could be attributed to holidays or bad weather.

Project Scope

Through a demand prediction model, our goal was to **better align** the Community Food Pantry's orders from the Regional Food Bank of Northeastern New York with **actual client needs.**

In Rockland County, the Community Food Pantry's efforts are becoming increasingly crucial. In 2023, out of the 104,325 households in the county, 12,713 lived below the federal poverty line, and 9,469 households lived below the ALICE—Asset Limited, Income

What is the ALICE threshold?

Developed by United Way, the ALICE threshold represents the income needed for a household to afford basic necessities in their specific county.

Constrained, Employed—threshold. 50.02% of the county was experiencing poverty. In the Suffern village, where the pantry is located, 254 households were below the federal poverty line, and 2,046 were under the ALICE threshold out of a total 4,551, meaning 50.54% of households were experiencing significant financial struggles. Local communities that the Community Food Pantry also serves report even higher rates of struggle. In the Kaser village and Spring Valley, 92.37% and 77.41% of households experience financial hardship, respectively. In Hillburn, 67.62% of households did.

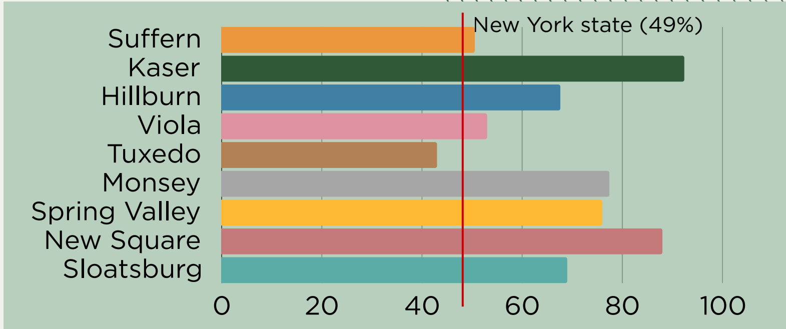

Percentage of Households in Nearby Communities under the Federal Poverty Line or ALICE

In Rockland County, the percentage of households under the Federal Poverty Line or ALICE increased by 38.89%, from 36% to 50%, between 2019 and 2023. As economic pressures such as persistent inflation, rising housing costs, and slower wage growth intensify, this trend is unlikely to reverse and likely to worsen, meaning food pantries must prepare for sustained increases in demand.

With this continuing increase in financial struggles in the area, the Community Food Pantry's work is not only increasingly crucial but also increasingly challenging. We aimed to help the pantry adjust to not only this surge in clients, but to prepare adequately for inevitable future increases as well.

Our **demand prediction model** aims to enable the Community Food Pantry to operate effectively and efficiently by understanding how different factors affect their demand. **Trained upon past data** from the Community Food Pantry itself, the model will factor in weather and time of year to estimate the number of households and individuals expected to visit within the two-week distribution cycle. It will provide the total units/lbs of food required, specifying the amount for different food groups and categories. This will streamline and ease the ordering process.

Moreover, volunteers are a vital part of the Community Food Pantry, but availability can be inconsistent, with some days being oversaturated and others understaffed. To counter this, our model **estimates the number of volunteers needed** to serve the estimated number of clients. After inputting the conditions of the coming cycle, all of these suggestions will become available to pantry staff on a **user-friendly, comprehensive dashboard.** A weekly "Pantry Prep Plan" will also be available to briefly summarize the model's findings for the following week.

After the model is implemented, Community Food Pantry staff will **continue to collect data** on the actual client attendance, and **new data will be added to the model's training set**, encouraging continuous learning. Thus, the model will grow with the pantry and its surrounding community, with the rationale that as economic and social trends change, the model and its predictions will adapt alongside these changes.

III. PLANNING AND ORGANIZING

Project Goals

MAIN GOAL

Ensure that at least **80%** of distribution cycles maintain adequate stock through at least **75%** of the cycle, while keeping food waste under **30 lbs**

FORECAST ACCURACY

With a model trained on historical data, we aimed to accurately estimate client turnout and food demands to ensure smooth and consistent service.

RESOURCE EFFICIENCY

Through a weekly "Pantry Prep Plan," we aim to identify potential high-demand weeks and allocate volunteers more effectively, as well as present predictions simply.

Human Resource Management Plan

To execute our project, we worked closely with the HR and Operations Coordinator, Director, and Food Pantry and Recovery Coordinator.

HR and Operations Coordinator

Elizabeth Takacs, HR and Operations Coordinator of the Community Food Pantry, played a crucial role in the operational side of our project. She provided the historical data needed to train our ML models, including household turnout, food distributed, common categories, and inventory logs. Ms. Takacs also reviewed and approved the model's staffing proposals for each biweekly cycle, making recommendations for potential aspects of our model to improve pantry operations. She ensured our plan integrated smoothly into existing HR processes, and her help enabled us to propose a system that improved volunteer efficiency while maintaining a high quality of service for clients.

HR and Operations Coordinator

Elizabeth Takacs

Director

We worked closely with Marisol Roman, Director of the Community Food Pantry, who served as our primary point of contact throughout the project. When initially identifying potential projects, we first met with Ms. Roman over Zoom to discuss the goals and unique operational challenges the pantry was facing, including demand volatility, and volunteer scheduling limitations. She aided us in

Marisol Roman
Director

developing our initial project proposal and presenting it to the Community Food Pantry Board of Directors and their chairperson, Susan Meyer. Throughout the project, Ms. Roman offered critical insight into the inner workings of the pantry and detailed feedback that helped us adjust our system to meet the pantry's specific needs. We met with her biweekly to discuss our progress. Her support was crucial in coordinating between staff and volunteers.

Food Pantry and Recovery Coordinator

To ensure that all operational and inventory data were accurate and reliable, we worked alongside Community Food Pantry's Food Pantry and Recovery Coordinator, Terry Deasy. Terry supervised

Food Pantry and Recovery Coordinator
Terry Deasy

the detailed records collected during every distribution cycle, including the total pounds of food distributed, the amount of waste, and household counts. Her close oversight of day-to-day operations and inventory allowed us to validate our predictive model's outputs against real outcomes. Moreover, her high level of experience in the pantry enabled her to analyze the plausibility of our model's early predictions. She provided guidance on workflow adjustments and organizational categories for the groceries, which improved the model's efficacy and efficiency.

Schedule

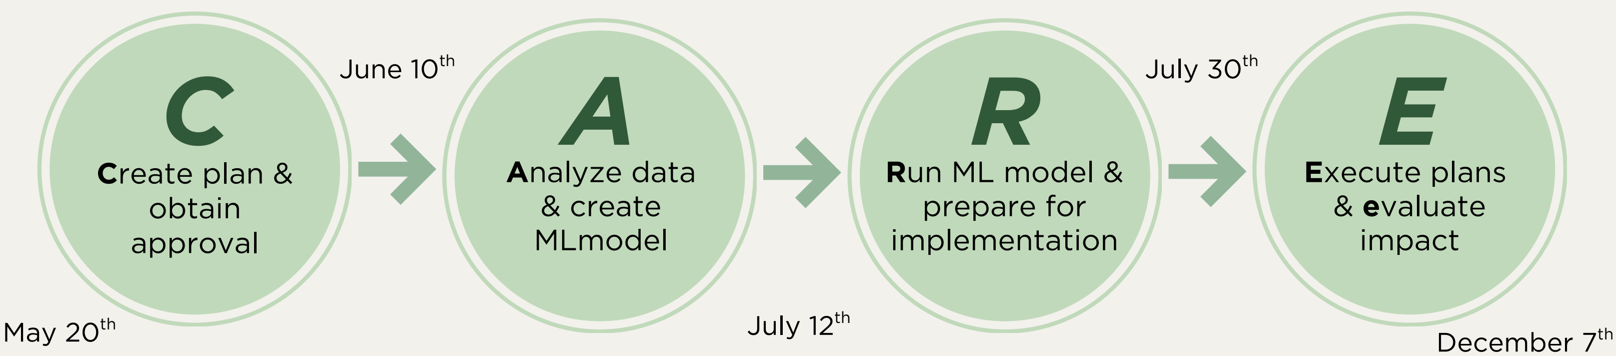

Milestone 1: Creating Project Plan and Obtaining Approval

Our milestone was to obtain approval for our project proposal from Ms. Takacs, Ms. Roman, and the board of directors of the Community Food Pantry. Initially, we met with Ms. Roman online on May 20th to discuss the possibility of working with them, and were introduced to the problems they were experiencing with unpredictable demand. Next, we visited the pantry in person, learning about their pantry inventory, operational systems, and supply sources. We also spoke to more of the staff, including Ms. Deasy and some board members. As we continued to keep in contact with the pantry, we began developing an initial proposal for a machine learning model to predict the demand of customers, adding features according to the needs of the organization that volunteers and staff informed us they needed. On June 10th, we presented our proposal at a meeting of the Board of Directors after being introduced by Ms.

Roman. Our project was approved, so we began the process of data analysis, and model creating and training.

Milestone 2: Analyzing Data and Creating Machine Learning Model

The second milestone focused on collecting historical data and building the predictive model's foundation. With the help of Ms. Deasy and Ms. Takacs, we gathered comprehensive pantry records from previous distribution cycles, including household turnout, food distributed by category, volunteer logs, and amount of perishable waste. We also incorporated external factors such as holidays and weather patterns that influenced demand. Once data was collected, we standardized and aggregated it to ensure accuracy while keeping client privacy. We began constructing a machine learning model to forecast individual and household turnout, food requirements, and needed volunteer staffing during a specific distribution cycle. To ensure that the model was accurate and reliable, we tested it with the parameters of a past distribution cycle and compared it to the actual data.

Milestone 3: Running Machine Model and Preparing for Implementation

The third milestone focused on moving from model development to real-world application. After finalizing our training pipeline, we ran the machine learning model using inputs for an upcoming distribution cycle (including seasonal timing, weather conditions, and holidays) to generate actionable forecasts. The model produced estimates for expected household and individual turnout, total pounds/units of food needed, a category-level breakdown to better align orders with client preferences, and a recommended volunteer count to avoid over- or understaffing. To make these outputs usable for day-to-day operations, we translated the model's results into a **simple, staff-friendly dashboard and a one-page weekly "Pantry Prep Plan"** that summarizes predicted demand, high-demand days, suggested staffing, and ordering considerations. We met with Ms. Takacs and Ms. Deasy to review whether the predictions were realistic and operationally helpful, then refined the format and categories to match the pantry's existing workflow. By the end of this milestone, the model was producing accurate forecasts and was packaged to ensure the pantry could implement consistently and confidently.

Milestone 4: Executing Plans and Evaluating Impact

The fourth milestone focused on implementing our system during a real distribution cycle and measuring its impact on pantry operations. Using the model's predictions, we used the weekly "Pantry Prep Plan" to guide ordering and volunteer scheduling ahead of high-demand days. Staff used forecasted turnout and category-level estimates to better align inventory with client needs, while the recommended volunteer count reduced congestion at check-in, restocking, and client assistance. After the distribution cycle, we evaluated impact by comparing the model's predictions to actual outcomes recorded by the pantry (household counts, pounds distributed, and perishable waste). We also assessed operational success using our project goal benchmarks (whether adequate stock was maintained through at least 75% of the cycle and whether food waste stayed under 30 lbs) while collecting feedback from staff on ease of use and workflow fit. Using these findings, we debugged, iterated, and refined the system by fixing issues with the code, correcting calculation or formatting errors, and tuning the model and dashboard outputs. This would allow the pantry to run the tool reliably and add new data to the training set for continuous improvement.

Quality Management Plan

Key Performance Indicators
Longevity of Orders: • How long groceries stay in stock in the distribution cycle
Food Waste: • Pounds of food and perishables discarded per cycle
Forecast Accuracy: • Average percent error between predicted and actual turnout
Volunteer Efficiency and Utilization: • Average households served per volunteer-hour

To evaluate the efficacy of the demand prediction model in not only accurately estimating the number of households visiting but also in reducing the number of distribution cycles resulting in either stockouts or wasted food, we established key performance indicators (KPIs). These KPIs were chosen to **directly reflect the pantry's core goals and values:** reliably serving families, minimizing food waste, and allocating resources efficiently.

The primary KPI was the longevity of orders from the Regional Food Bank of Northeastern New York. This is measured by when, during the distribution cycle, if at all, groceries run out, and families can no longer receive goods. This metric measures the pantry's ability to consistently meet client needs without running out of food early in a cycle. This KPI will be assessed weekly, and extended levels of low longevity indicate errors or bugs in code, leading to the model being examined and debugged immediately.

Another critical KPI was food waste, measured by the pounds of food discarded per distribution cycle. Although having enough food to feed all families is crucial, so is ensuring there is not an excess of perishables. This indicator assesses how well the pantry **balances ordering enough food to meet demand while avoiding expiring inventory.** It reflects the ML model's ability to predict an accurate and precise estimate. Reductions in food waste indicate high alignment between demand predicted by the model and actual demand. Like longevity of orders, issues with food waste would indicate issues with the model, prompting investigation and debugging.

Forecast accuracy is a KPI for the ML predictive model itself, tracking the average percent error between the model's predicted turnout and the actual number of households and individuals served. Improving forecast accuracy indicates a high-functioning model and is necessary to achieve the other three KPIs. It's essential to build trust in the model and ensure that its recommendations are accurate and actionable for staff.

We also measured volunteer efficiency and utilization as a KPI. Before our project, the Community Food Pantry's volunteer force was often inconsistent, with most volunteers working on similar days and times of the year. By considering the number of households served per volunteer hour, this measure can analyze how efficiently volunteer labor is matched to client demand by the ML model. Improvement in volunteer utilization leads to shorter wait times for clients and more manageable workloads for both staff and volunteers.

Together, these KPIs form a framework to evaluate the project's impact. In the case of disparities or problems with certain KPIs, we would meet and analyze the model to find the problem in the code or data, and remove it as soon as possible.

Risk Management Plan

We identified data quality gaps, forecast risk, inventory/supply variability, staff adoption challenges, and client privacy as the primary risks to the success of our demand prediction model and "Pantry Prep Plan."

Data Quality & Consistency: Since our model relies on historical pantry records, incomplete or inconsistent logging could reduce forecast accuracy and lead to poor ordering or staffing decisions. To reduce this risk, we worked with pantry staff to standardize the fields collected each cycle (household turnout, pounds distributed, category totals, and perishable waste). We also built basic checks into our workflow to flag missing entries or unusual values. For example, unexpectedly low turnout or category totals that don't match the overall pounds.

This risk matrix summarizes the main risks to our Demand Prediction Model and Pantry Prep Plan by rating each risk's likelihood and operational impact on a 1-5 scale. Risks in the upper-right quadrant are the highest priority because they are more likely to occur and more disruptive to pantry operations.

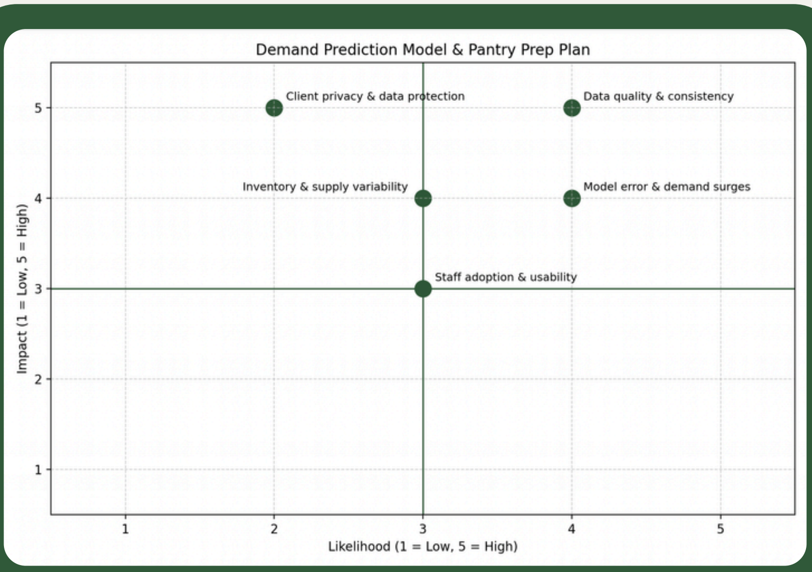

Model Error & Demand Surges: A major risk was that the actual demand could exceed predictions, especially during busy periods such as holidays, severe weather weeks, or periods of economic change, causing stockouts earlier in the cycle. To manage this, we created a two-tier response plan:

- **Standard Plan:** Use the model's predicted turnout and category estimates to guide ordering and volunteer scheduling for typical weeks.
- **Surge Plan:** If a holiday, severe weather, or recent turnout trends suggest higher-than-normal demand, the model outputs a prediction plus an uncertainty range. Staff would use the upper end of the range to add a buffer for ordering and staffing, depending on which risk

signals were present (Severe weather weeks may need a higher buffer than holidays, for example).

Inventory & Supply Variability: Even with accurate forecasts, the pantry could face inventory risk due to changes in food bank availability, donations, or delivery timing. Therefore, the "Pantry Prep Plan" emphasizes flexible category substitutions and focuses on ordering priorities rather than rigid item lists. We also recommend tracking common "low-supply" categories so the pantry can anticipate where substitution planning is most needed.

Staff Adoption & Usability: If the dashboard or Pantry Prep Plan were too complex, pantry staff may not use it consistently. To address this, we designed outputs to be simple, visual, and action-oriented, and we met with pantry leadership to align the plan with their existing workflow. During implementation, we collected feedback and debugged issues in the code and reporting format, so the tool could be run reliably with minimal extra work.

Client Privacy & Data Protection: Because pantry data includes sensitive information, client privacy was a major risk factor. To reduce this risk, we aggregated data to the cycle level (counts and totals) and avoided storing personally identifiable information in the model pipeline. We prioritized secure handling and limited the model's inputs to operational variables necessary for forecasting.

Proposed Project Budget

Our solution is software-based, so costs are low and primarily cover deployment, printing, and training. In fact, many tasks, including training our model, have no cost. Using open-source tools and existing pantry systems to improve efficiency and reduce costs, we estimate a **$290** out-of-pocket budget for a 6-month pilot.

Category	Details	Cost
Data & Workflow Development	Cleaning past records, building standardized intake template, validation checks	**$0**
Model & Dashboard Development	Model training pipeline, dashboard build, automation script coding	**$0**
Hosting & Deployment (6 months)	Lightweight cloud hosting (Streamlit/Render/AWS small instance) and storage	**$120 ($20/mo)**
Backups & Security (6 months)	Secure drive storage. access controls, backup snapshots	**$30**
Staff Training & Handoff	One to two training sessions, documentation and how-to guide	**$40**
Contingency	Unexpected fixes and maintenance (data gaps, format changes, extra hosting)	**$100**
TOTAL COST		**$290**

IV. EXECUTION

Creating Project Plan and Obtaining Approval

Our first contact was with Ms. Roman, the Director of the Community Food Pantry, in a Google Meet. After learning about some of the logistical struggles that the pantry had been facing in recent

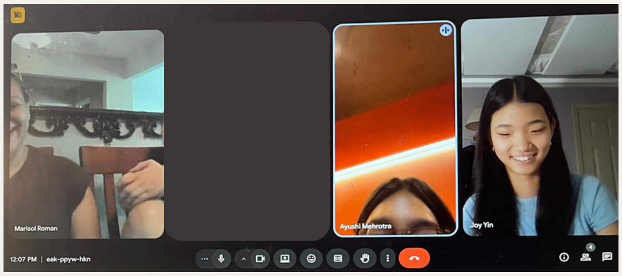

years due to a surge in clientele, we decided that we wanted to create a solution to help the pantry staff serve the community with more ease and confidence. When we visited the Community Food Pantry in person the next week, Ms. Roman and Ms. Deasy explained the ordering process, which they often found frustrating because of the unpredictability of the attendance and amount of clients they would see during that two-week cycle.

We decided to create a machine learning model to predict the number of clients visiting and the amount of groceries needed in a specific two-week period based on different factors. After discussing our ideas with Ms. Takacs, the HR and Operations Coordinator, she also suggested including a volunteer prediction because she noticed an inconsistency, with certain days having significantly more volunteers than others, leading to the pantry being overstaffed on some days and understaffed on others. In both cases, this led to inefficiency. Thus, we decided to add a prediction of the number of needed volunteers as well.

When presenting our proposal to the Community Food Pantry Board of Directors, the main concern was integration into existing pantry systems, especially since a majority of the staff had little if any machine learning experience. With this suggestion, **we modified our strategy to include a comprehensive dashboard and a weekly "Pantry Prep Plan"** to simplify the predictions for the staff. Once our project proposal was approved, we started working on our model.

Analyzing Data and Creating Machine Learning Model

Step 1

Inputs
- 5 years of pantry data
- Seasonality/holiday tags
- Weather risk flags

Step 2

Processing
- Clean & merge spreadsheets
- Standardize units & categories
- Missing-value checks

Step 3

Outputs
- Turnout forecast
- Category demand estimate
- Volunteer recommendation

We began gathering and analyzing historical pantry data from the last five years. To standardize the data, we imported past spreadsheets into **Python (Pandas library)**, then cleaned and merged them into **one consistent dataset**. We renamed columns to match one format, converted all weights to pounds, mapped old category names into a single category list, excluded rows with missing values, and removed duplicates/outliers so each distribution cycle could be compared reliably.

As the pantry's clientele grew significantly over the last two years, we continuously added new biweekly distribution cycle into the training set so the model captured the newer, higher-demand pattern instead of only the older baseline. To improve prediction quality, we also **tagged external factors** such as seasonality, holidays, and severe weather weeks. Additionally, to protect client privacy, we aggregated the data into weekly totals (households served, pounds distributed by category, and waste) and **removed any personally identifiable information**, ensuring the model only used data trends rather than individual client records.

Using this dataset, we trained and tested a demand forecasting model with scikit-learn to predict:

- **Expected household + individual turnout**
- **Category-level food demand (estimated pounds/units)**
- **Recommended volunteer coverage based on expected traffic**

We validated performance by comparing predicted vs. actual outcomes from past cycles and tuning the model and preprocessing rules to reduce error and improve reliability before implementation.

Running Machine Learning Model and Preparing for Implementation

After training our model, we **tested the performance** by running it on past distribution cycles and comparing predictions to the pantry's recorded outcomes. We used a **regularized regression model** (Ridge) because it handles correlated features (such as different food categories and seasonal patterns) without overfitting. We used historical distribution cycles as our training data and included features such as season, holiday indicators, school breaks, severe weather, and donations. We randomly split our dataset of past distribution cycles into 80% training and 20% testing data and evaluated the model using Mean Absolute Percent Error (MAPE).

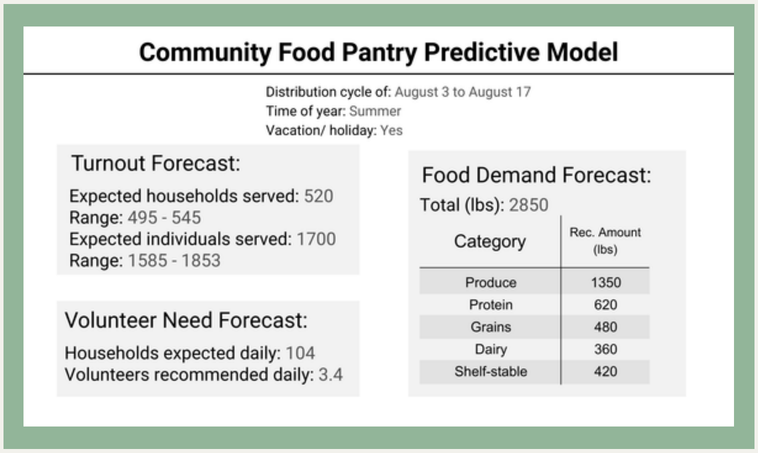

To make sure our results were not due to a lucky split, we repeated this process 30 times with different random seeds. For each candidate regularization strength (alpha) in our Ridge model, we averaged the MAPE across the 30 runs and selected the value of alpha=1.0, which produced the lowest average error. Across historical test cycles with this tuned model, the Ridge regression achieved an average MAPE of 4.2% for turnout and 4.9% for total food demand.

We found this testing to be successful, demonstrating our model was reliable and accurate enough to be implemented. To prepare for implementation, we created a simple and comprehensive dashboard to summarize all the model's predictions for the coming distribution cycle, consolidating key forecasts. Moreover, a weekly "Pantry Prep Plan" was also designed to signify benchmarks and expected traffic levels to the staff for a shorter time frame. This enabled staff to make **informed adjustments** in advance rather than reacting during the distribution cycle. These tools ensured the model's outputs were accessible and actionable for pantry staff with varying levels of technical experience.

Executing Plans and Evaluating Impact

After testing and refining the demand prediction model, we moved into the execution phase by implementing the model during live distribution cycles. Our first live distribution cycle was from August 3rd to August 16th. For each cycle, we input variables such as seasonal timing, expected weather, and whether it was during vacation or a holiday. The model then generated forecasts for household turnout, recommended food quantities, and suggested volunteer staffing levels. These outputs were shared with pantry staff and used to order groceries from the Regional Food Bank of Northeastern New York and monitor volunteer sign-up levels. They received the predictions through a dashboard and a weekly "Pantry Prep Plan," allowing staff to prepare orders in advance. These elements were especially popular with the staff, who found it comprehensive and easy to understand. To ensure the staff understood how to employ the model, we suggested a **brief online training session** for a few staff, such as Ms. Roman and Ms. Takacs. After getting approval from the Board of Directors and Ms. Takacs for this training, three staff members were trained in preparation for future complete handoff of the prediction model.

> "The **Pantry Prep Plan** was **extremely helpful** to plan this next week out. I think it just saved me hours of work!"
> -Staff member

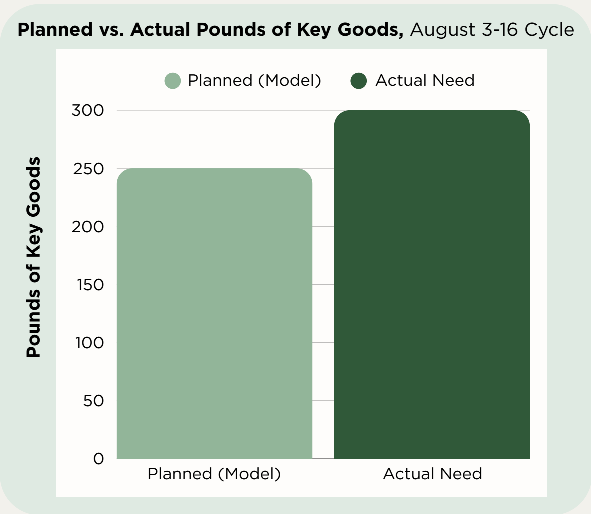

To evaluate the impact of the model, we compared forecasted results to actual outcomes after each distribution cycle. Our first cycle did not go as planned: the model underestimated turnout and several high-demand items ran out sooner than planned. Based on the model's recommendations, we ordered about 52 fewer pounds of high-importance goods than were actually needed for that pantry cycle, leading to a pantry stockout. At the same time, the August 3–16 cycle ended with no perishable food wasted as staff used every pound we ordered. Therefore, households arriving later in the week faced limited choices once high-demand items ran out. So, we investigated the training data immediately and found that in the earliest live cycles, the model tended to slightly underestimate demand because it was trained on historical data from before the most recent client surge (approx. 3 years ago). By reviewing these discrepancies, we were able to adjust our feature weights, incorporate additional recent distribution cycles into the training set, and tightened weather and seasonality tags. As more up-to-date data was added, forecast error decreased and the model's recommendations aligned more closely with actual turnout. Over time and more runs, the Demand Prediction Model evolved into a reliable tool that staff could use to plan orders and volunteer shifts confidently. In its most-recent configuration, the model was able to achieve a MAPE of 1.32% for household turnout.

V. MONITORING AND CONTROLLING

Monitoring: Data Collection, Model Performance, and Pantry Operations

After launching the Demand Prediction Model and weekly "Pantry Prep Plan," we monitored each distribution cycle to confirm the system was accurate, easy to use, and operationally helpful. Staff logged cycle outcomes (household/individual turnout, pounds distributed by category, and perishable waste) using a standardized format so we could compare results consistently over time. Across the cycle, we compared predicted turnout and category estimates to actual demand and tracked goal-aligned indicators: whether stock lasted through at least 75% of the cycle, whether perishable waste stayed under 30 lbs, and whether volunteer coverage was sufficient or lacking.

We also gathered staff feedback on the dashboard and Prep Plan, ensuring it was clear and accessible during high-traffic moments (check-in, restocking, and client assistance). When unusual conditions occurred, such as holidays or severe weather, we tagged those cycles so the model could account for them in future updates.

KPI Scorecard (Example Week/Cycle)	
Stock lasted more than 75% of cycle	**Met**
Perishable waste < 30 lbs	**Met**
Forecast error (MAPE)	**3%**
Volunteer coverage	**Adequate**

Controlling: Plan Adjustments and Continuous Improvement

To manage forecast error and sudden demand spikes, we used a two-tier approach. First, for typical weeks, staff followed the Standard Plan using the model's predicted turnout, category estimates, and volunteer recommendations to guide ordering and scheduling. When risk signals appeared, the Surge Plan shifted decisions toward the upper end of the model's prediction range, reducing stockouts and long lines since the pantry was adequately prepared.

During implementation, we also made refinements based on what we observed, fixing calculation or formatting issues, clarifying category labels, and simplifying visuals so staff could act quickly. We ran basic data checks to flag missing or inconsistent entries and correct or exclude them during preprocessing to protect forecast quality.

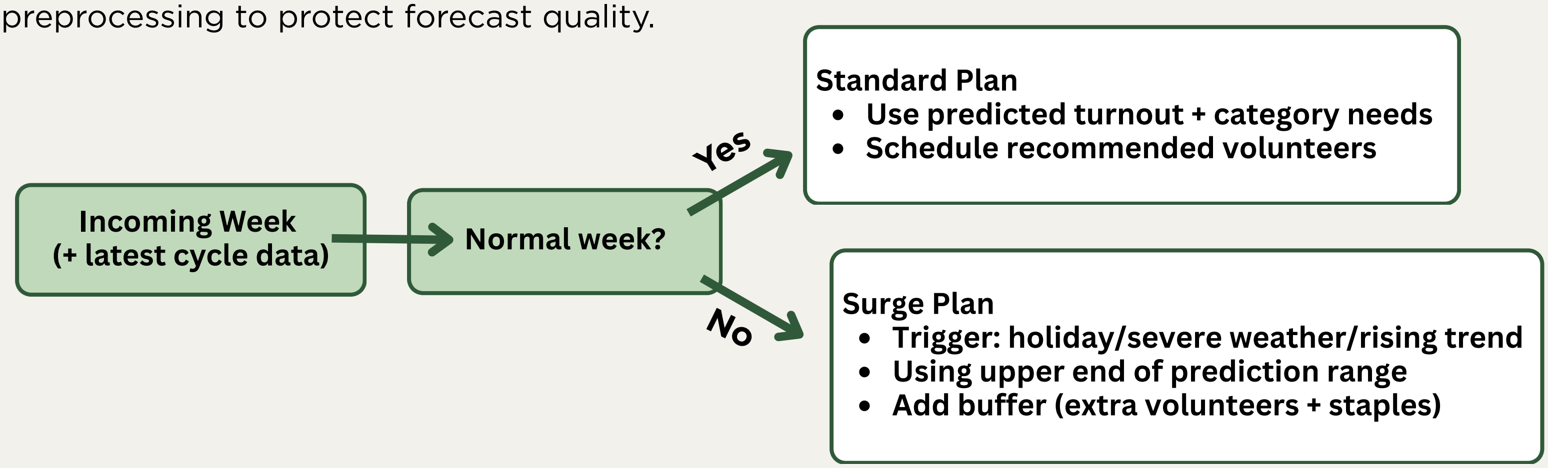

VI. CLOSING THE PROJECT

Evaluation of Key Metrics

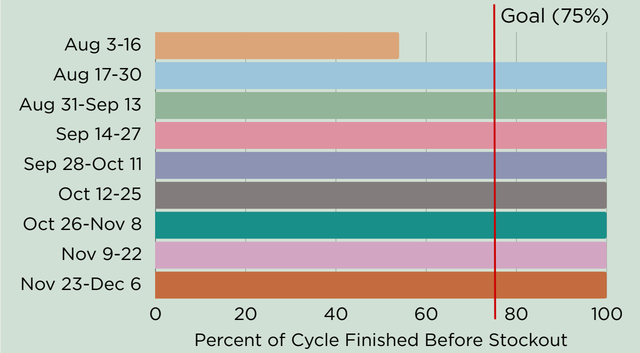

Percentage of Distribution Cycle Completed Before Stockout for Each Cycle

Longevity of Orders

In early 2025, before our project was implemented, five out of twelve distribution cycles resulted in stockouts before 75% of the cycle was complete, which is 41.67%. After we deployed the demand prediction model, inventory lasted much longer. In the nine cycles where the model was used, groceries stayed in stock past the 75% mark in eight cycles, and only one cycle (August 3–16) experienced any stockout at all. This is a significant improvement and demonstrates a more informed ordering process as a result of our predictive model. Our first live cycle, August 3–16, ended in a full stockout of a high-priority staple item, and inventory across several key categories ran out before 75% of the distribution cycle was complete, so this first run did not meet our longevity goal. Because the model had under-estimated turnout, families who arrived later in the week had limited choice, but there was essentially no perishable waste as almost everything we ordered was used. Factors such as pre-holiday demand or unexpected community needs may have contributed to this surge. To account for this, we used that outlier to refine how the model handles seasonal demand surges so that upcoming holiday cycles are treated more conservatively. Additionally, we implemented the two-tier approach and added a range to each prediction, allowing staff to use the upper end of the range as a buffer for ordering food and scheduling volunteers during high-risk weeks. After these refinements, the final version of the model maintained a Mean Absolute Percentage Error of 1.32% for household turnout across 9 live cycles, helping staff keep groceries in stock through at least 75% of each cycle in eight out of nine cases (no other cycles had a stockout).

Waste Reduction

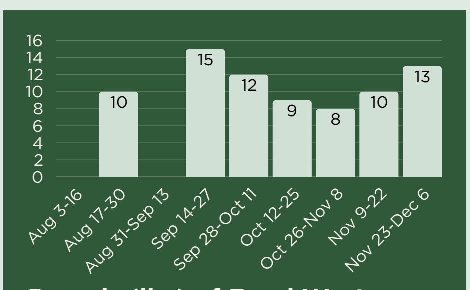

Pounds (lbs) of Food Waste per Distribution Cycle

The average amount of perishable waste after distribution decreased significantly after our model was implemented. All distribution cycles had waste under 30 lbs, exceeding our goal of at least 80% having under 30 lbs of waste. In the first live cycle (August 3–16), turnout was higher than expected, so waste was close to zero. After we retrained the model and added safety buffers, later cycles had smaller amounts of unavoidable waste (typically 10–20 lbs per cycle) but no shelves full of unsent produce. This pattern shows that the model helped staff order enough food to meet demand while still minimizing what had to be discarded.

Forecast Accuracy

Forecast accuracy is essential to effectively predict the turnout. Across the nine distribution cycles, the MAPE for turnout was 1.32%, indicating that predictions were very close to actual households served. For 7 out of 9 cycles, the percent error was under 2%. For 5 of those 9 cycles, it was under 1 percent. The clear outlier was the first live cycle, August 3–16 (3.7% error), which ended in a full stockout of a high-priority staple item because the model under-estimated turnout. August 17–30 (2.3% error) still slightly under-estimated demand, but with added buffer stock, it did not result in a stockout. We went back and investigated the model's data, added the most recent years of data (approx. 3 years ago, when turnout surged) and strengthened indicators for holidays, breaks, and weather. We also added the range to provide a buffer volunteers could use during high-demand cycles. After these refinements, the percent error steadily declined over time: the last five cycles all had error under 1%, and the final two cycles were below 0.5%. As shown in the Predictive Model Percent Error chart, this pattern shows that the model became more accurate with each additional cycle of data. The final version of the model maintained an average MAPE of 1.32% across nine live distribution cycles, accurately capturing both typical weeks and high-demand surges.

Predictive Model Percent Error (Turnout)

Distribution Cycle	% Error
Aug 3-16	3.7
Aug 17-30	2.3
Aug 31-Sep 13	1.5
Sep 14-27	1.0
Sep 28-Oct 11	0.8
Oct 12-25	0.9
Oct 26-Nov 8	0.7
Nov 9-22	0.6
Nov 23-Dec 6	0.4

1.32%
average MAPE

Volunteer Efficiency and Utilization

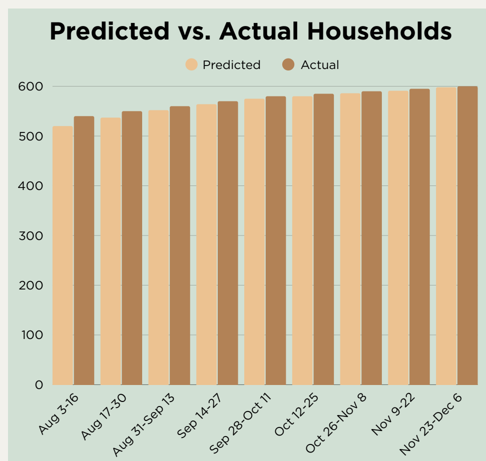

Before our project, volunteer turnout was highly uneven, with some distribution days being overstaffed and others understaffed. By predicting household turnout and recommending a target number of volunteers, the model helped align staffing levels with expected demand. Staff reported that queues at check-in and restocking became more manageable, and busy periods were easier to cover because they could identify high-traffic days in advance. Across the nine evaluated cycles, average households served per volunteer-hour increased from about 9 to roughly 12, a gain of more than 30%, and the number of days classified by staff as significantly overstaffed or understaffed dropped from 7 out of 12 pre-project cycles to just 2 out of 9. This demonstrates that the model supports not only better ordering decisions but also more efficient use of the pantry's volunteer team.

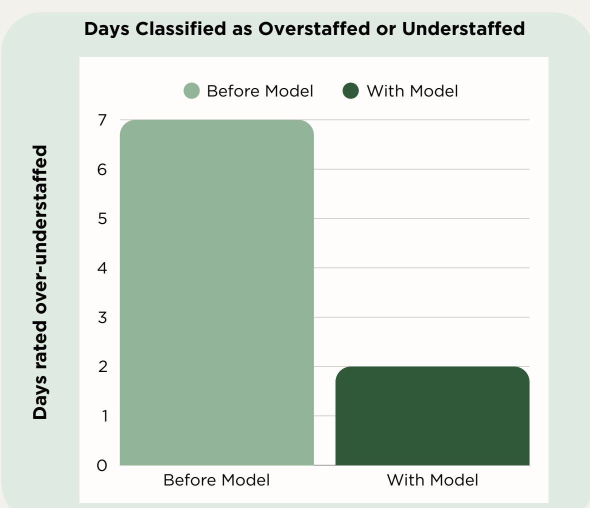

Lessons Learned

Understanding Staff Needs

Initially, we had focused our model on predicting household turnout and the amount of groceries needed. However, we learned during a meeting with Ms. Takacs that the Community Food Pantry struggled not only with the unpredictability of client attendance but also of volunteers. To address this, we updated our plan to include predictions for the recommended number of volunteers for each distribution day.

Moreover, the Board of Directors of the Community Food Pantry expressed concern during our initial presentation of our project proposal over the accessibility of our model to staff members unfamiliar with technology and machine learning. This insight led us to add the comprehensive dashboard and weekly “Pantry Prep Plan” so that all staff, regardless of their background, could make use of our project. This made our data accessible and allowed staff to better coordinate their efforts. We learned that aligning technical decisions with staff realities is essential for operational success. It is necessary to inquire about and consider various factors and needs of the employees and the organization in order to create an accessible and effective solution.

Continuous Improvement

Despite training going well, our first distribution cycle did not go as planned, with our initial prediction heavily underestimating demand because the model was trained on older data and didn’t completely account for the recent surge in the population in need in the area the pantry serves. Through monitoring the earliest cycle and comparing it to predictions, we were able to adjust parameters, leading to higher accuracy in the following cycles.

Moreover, in updating the model with fresh data after every distribution cycle, we realized the value of **continuous feedback and recalibration,** enabling the model to develop and change along with real-world circumstances and live trends. Because of this, the model’s predictions of latter distribution cycles were significantly more accurate because they had more recent data that continued to reflect the pantry’s evolving demand.

Recommendations for Future Projects

Expand Scope

Our predictive model was applicable only to the Community Food Pantry. However, the underlying framework has the potential to be scaled and adapted to serve more than one pantry or even entire regions. This larger, more complex model would need additional data sources, such as more detailed seasonal demand fluctuations and population changes. It would be significantly helpful for many organizations to have a widely available, more general model to adapt to their own needs. Moreover, not every food pantry struggles with demand prediction. A machine learning model expanded to predict other elements of community service, such as donations and fundraising, would be helpful for many food pantries.

The scope could also be expanded by creating predictive models for other charities and 501(c)(3) organizations that serve the community in different ways.

Community Engagement and Feedback

To plan and execute our project, we worked almost exclusively with the management and staff of the Community Food Pantry. Although this provided extremely useful and significant insights into the inner workings and operational challenges of the pantry, we lacked a community or client viewpoint. Thus, future projects could incorporate direct feedback from pantry clients and volunteers to better understand the community's needs and preferences. Using surveys or focus groups could provide insights that visitor statistics and data alone could miss, such as accessibility challenges or workload concerns for volunteers. Since clients may have concerns missed by management, this would allow for a broader and more complete perspective on the problems facing the organization being studied, and potential solutions.

By integrating this feedback, the predictive model could allow for more community-and-client-centered decision-making. This would ensure that the project serves not only the operations and efficiency of the pantry, but also the needs and wants of the community and those serving it.

VII. BIBLIOGRAPHY

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., … Oliphant, T. E. (2020). Array programming with NumPy. Nature, 585(7825), 357–362.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. Technometrics, 12(1), 55–67.

Kramer, P. D. (2025, July 26). Rockland's poorest: Report shows where county's poor and working poor call home. The Journal News. https://www.lohud.com/story/news/local/rockland/2025/07/25/rockland-ny-communities-with-most-poor-working-poor-households/853536480007/

McKinney, W. (2010). Data structures for statistical computing in Python. In S. van der Walt & J. Millman (Eds.), Proceedings of the 9th Python in Science Conference (pp. 51–56). Austin, TX.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

Python Software Foundation. (2023). Python [Computer software]. https://www.python.org

Sloatsburg Food Pantry. (2026). Community food pantry [Website]. Retrieved January 2026, from https://www.sloatsburgfoodpantry.org

United For ALICE. (n.d.). ALICE county profile: Rockland County, New York. United Way of Northern New Jersey. https://www.unitedforalice.org/county-reports/new-york

Wetherill, M. S., White, K. C., Rivera, C., Seligman, H. K., & Cannuscio, C. C. (2018). Food pantries as partners in population health: Assessing organizational and personnel readiness for delivering nutrition-focused charitable food assistance. Journal of Hunger & Environmental Nutrition, 14(1–2), 50–69.