INVENTORY MANAGEMENT FOR PENDA HEALTH

Prepared for: Penda Health
Spring 2015

Submitted by: Aylin Elmaagacli, Zueber Juma, Theo Kelly, Tracy Koczela, Bryan Murphy, Pablo Garza Villarreal
Professors Werwath and Klabjan
Industrial Engineering Design Project
McCormick School of Engineering and Applied Science
Northwestern University
Evanston, IL 60208

McCormick Northwestern Engineering

PENDA HEALTH
Your Caring Health Care Provider

ACKNOWLEDGEMENTS

Prof. Werwath - For his assistance in project management.

Prof. Klabjan - For his guidance on a variety of project matters.

Prof. Iravani - For his input on technical issues.

Prof. Ankenman– For his advice on estimating demand data with sales data.

Prof. Staum– For his insight into fitting distributions to data.

Prof. Perry– For his instruction on how to properly use the exponential distribution.

Prof. Wilson– For her willingness to connect our team with several helpful sources.

Prof. Nelson – For his assistance on handling limited sales data.

Matt Hawkins- For his partnership and support throughout the quarter.

John Njoroge– For his collaboration and descriptions of current inventory procedures.

Jack Koczela - For his business advice and general guidance.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Penda Health requested support with their inventory management policy, with the ultimate goal of achieving a greater service level for approximately 200 pharmaceutical drugs. By incorporating a periodic review system with a base order up-to level for all drugs, a dynamic and easy-to-use user interface was provided, where Penda Health employees will get the optimal order recommendation for each drug in order to achieve a 95%, 90%, or 85% cycle service level (terminology definitions can be found in Appendix A). The model provides a dramatic improvement of cycle service level for all of their drugs when applied to historical data, as can be seen below in Figure 1.

Target CSL	CSL in Current Policy	Historical CSL Achieved in Proposed Policy	CSL Improvement
95%	58%	96%	38%
90%	53%	90%	37%
85%	63%	89%	26%
Avg. for All Drugs	57%	92%	35%

Figure 1: CSL improvements in Proposed Policy vs Current Policy

PROBLEM STATEMENT AND BACKGROUND

COMPANY OVERVIEW

Penda Health is a for-profit social enterprise that aims to provide high quality primary healthcare services in Nairobi, Kenya. Penda Health was founded in 2012, and currently has two operational health centers and one university sponsored health clinic, where they provide patient centric approach to health care requests. The scope of the services Penda Health provides includes the full range of primary care. Their mission is to make primary health care more accessible and affordable, while providing higher quality medical care and services for the whole family.

PROBLEM

Currently, Penda Health uses a general trial-and-error approach to ordering supplies; they have limited insight into their drug ordering policy and often stock out or hold too much stock in their clinics. These are immediate problems the company wants to address because under stocking results in lost revenue and a detriment to company image and patient experience, while overstocking may result in not enough space to keep the extra supplies. Penda Health keeps all their drugs in-shop, and there is an added risk of theft or natural disasters damaging excess inventory. Penda Health would like us to develop a more sophisticated inventory management system that aims to prevent stock outs while simultaneously minimizing the amount of excess inventory.

OBJECTIVE

The objective for the inventory management policy design is to create a dynamic and user-friendly model that incorporates all previous and future sales information of Penda Health's clinics and allows the user to choose a desired cycle service level (CSL) for every drug, although the model will already have included a predetermined recommendation of what each drug's CSL should be based on their profit margins, importance, and historical data. This CSL is used in combination with that drug's historical demand data to form an order up-to level. The model then tells the user how much to order based on the difference between this order up-to level and the drug's current inventory. The model also provides visual aids so that the user can see in a more detailed manner how much cost increases by each percentage of cycle service level improvement, and whether the user should choose the larger quantity or the smaller when comparing two viable order quantities based upon cost and CSL. All calculations are done "behind the scenes" as the model takes into account all the variables that can affect these outputs.

SCOPE

The order policy will apply to all of the approximately 200 drugs used in Penda Health's two clinics, Umoja and Kitengela, using the earliest available sales data, which extends back to December 2014. The model advises the client on the order up-to level of each drug and will also allow the client to easily update the model by adding future demands into its calculation, so that the recommended inventory policy remains operational even after the policy is handed over to the client. Furthermore, the inventory management model can be easily expanded to accommodate new drugs and clinics as Penda continues to grow.

BACKGROUND RESEARCH

VEND POS SYSTEM

Vend POS is a cloud based point of sale (POS) and retail management software targeted to small to medium businesses, and is used by Penda Health. Vend POS can manage inventory levels across multiple locations, and includes back-office features and customer management tools.

Vend provides POS reports for product and sales management, and records the amount of sales for each product over time. The historical sales data of the drugs for the two clinics of Penda Health are able to be obtained from Vend POS. Vend POS offers bulk import and export options for product management by using a

CSV file that contains information on each product such as product name, product price, and user specified reorder point (ROP).

The current inventory stock levels for the products of the different clinic locations are recorded on Vend POS, and product levels are displayed only for the time of viewing; historical inventory levels over time are not displayed. Weekly inventory data or data of other granularity can be obtained by viewing stock levels and recording periodically.

MULTICRITERIA INVENTORY CLASSIFICATION

Managing inventory levels for more than 200 different drugs necessitated a method to classify the inventory of the different products based on the importance of the drug determined by both cost and non-cost criteria. The two criterion we used for the classification of the products were product urgency and weighted average contribution margin. These were decided on by considering two driving factors behind Penda's business model: they wish to provide the best healthcare possible to their users, and they wish to make a profit so they can continue expanding the business.

DATA COLLECTION

SALES AND INVENTORY DATA EXPLORATION

Penda Health works with multiple suppliers, and places orders to the different suppliers for each product group. The six product groups are; medicinal drugs, lab supplies, dental supplies, stationery and printing, sanitation and cleaning supplies, and patient experience. By Penda's request, the inventory management system focuses on managing the inventory for the "drugs" product group and the ~200 items within this group. The inventory management procedure incorporates all previous and upcoming future sales information of Penda Health's clinics for every drug category.

The data for the inventory management policy was approached by first collecting all of the data available for each drug to approximate their demands. All of the daily sales data for each drug in both of Penda's Health clinic locations, Kitengela and Umoja was exported. However, noticing that the sales data might not reflect true demand since there would be no way to know whether there was more demand for a product when it stocked out, inventory data was utilized.

The inventory data was cross-referenced for each drug to the sales data, for the purpose of properly identifying the drugs with true demand information and the drugs without. Unfortunately, the inventory data was found to be very inaccurate. Since the inventory is counted manually once a week and entered manually, it is prone to error. This error manifested in a large number of negative inventory counts, and missing inventory counts in some dates for several drugs. Furthermore, since the inventory data only went back around four months, it was insufficient for further analysis.

Penda Health utilizes Vend as a Point of Sales system, and for back office computing to handle and track inventory levels. The historic data for drug sales was available starting from December 2014, and the available sales data was acquired from Vend. Some of the sales observations recorded are a censored representation of the underlying true demand due to frequent stock outs. In order to be able to utilize this data provided by Vend POS, further investigation was done into estimation approaches that accounts for this censoring. The sales data for all product categories is used to approximate the underlying demand process, and is used to develop the inventory management procedure.

Vend POS tracks the items sold, and updates the stock levels for the products sold. The inventory levels of the products are available for only the time checked, and are not reported in a historical and iterative level. After modifying the inventory data available to obtain historical weekly inventory levels, it was observed that Vend POS electronic inventory records occasionally did not match the quantity of customer available stock, which would create difficulties due to recording inaccuracy. In order to be able to implement the inventory management policy for each product, further research was conducted in order to eliminate any unnecessary discrepancy.

DISCUSSION OF ASSUMPTIONS

There were several assumptions made when creating the model. Perhaps the biggest assumption is that the current inventory count for each drug is accurate. Since the model is recommending the optimal order up-to level, the amount of drugs it will suggest ordering is contingent on the current inventory, since the user would then order up to the reorder point. If the inventory numbers are wrong for a drug or not up to date, then the recommendations for how much to order would also be wrong. Penda Health has been informed on the importance of maintaining proper inventory, and has been implored to perform a full inventory check before implementing our model to set all inventory levels to their true values. Unfortunately, there is no other way to bypass this assumption without having inventory counts be double checked and/or switching to an electronic method for counting inventory (which is unfeasible for a fledgling startup such as Penda due to the high cost of entry).

Another major assumption, also related to the data provided, is that the sales data is accurate. Since all sales are entered manually into their point of sale database system, there is room for error. For example, a user might forget to enter the sale of a drug or might enter a sale for a drug under a different drug's name. There is evidence for errors of this type in Penda's Vend database; for example, some drugs are listed multiple times under similar but unique names. These errors would make for inaccurate sales data that would in turn skew the predicting model to provide an inaccurate reorder point.

Furthermore, the sales data recorded are a censored report of the underlying demand because it is not possible to sell more units than the amount available in stock, and backorders/lost sales are not tracked by Penda. The inventory management models used in our calculations require demand data, and demand distribution parameters as an input. Demand is estimated from available past sales data in our model because upon analysis into the sales demand and after interviews with Penda, we have determined that stock outs are not observed frequently, and the observed few occurrences have been noted. Nonetheless, it is an assumption to assert that sales data is a proper representation of demand

Another assumption critical to the model is that weekly demand is independent and identically distributed from week to week. This is a common and almost ubiquitous assumption used in inventory modeling. However, it is an assumption nonetheless. In the case that demand in a given week is influenced by demand occurring in other weeks, this assumption could hinder our model's accuracy.

Finally, the model operates under the assumption that the exponential distribution provides the best fit for each and every drug. While this assumption was supported by both research and data analysis (as shown in the **Estimating Demand** section,) it could not be confirmed for each drug individually due to analysis restrictions. Conversely, given the highly variable nature of the exponential distribution, this could be seen as a conservative assumption as it results in our model accounting for a large amount of variability.

PROBLEM FORMULATION

APPROACH

Penda Health has determinate insight into their drug ordering policy and has experienced stock outs or over stocking in their clinics. The first step of designing the inventory management policy consisted of understanding the current inventory procedure of Penda Health, and their unmet needs from the current policy. Sales and inventory data was collected from Vend POS to gain an understanding into the data and to fit distributions to the sales data for the purpose of determining the model parameters. Using the previously determined parameters, the inventory model outputs an optimal order quantity and an order up-to level for each drug in each clinic based on the drugs cycle service level. Sensitivity analysis on the order quantity based on the cycle service level, and the resulting cost was conducted.

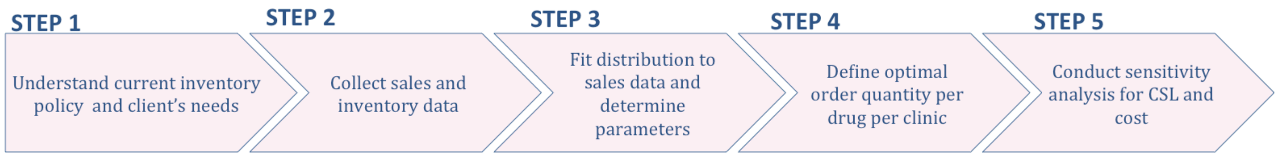

Figure 2: Outline of Approach

BUILDING THE MODEL

Estimating Demand

The demand for each drug was estimated from the historical sales data obtained from Penda Health Vend POS. @RISK was utilized to perform statistical analysis and to fit distributions, with AIC as the ranking statistic. The ten drugs with the highest total units sold were chosen for analysis due to the fact that they had the most data with the fewest periods with no demand. Demand for each drug was found to be best modeled by an exponential distribution when compared to other discrete and continuous distributions. Furthermore, research determined the exponential distribution best models demand data that can be classified as "lumpy demand," which is "extremely irregular with high level of volatility coupled with extensive periods of zero demand."[1] This description accurately applies to the context of Penda's drug ordering policy as numerous drugs experience zero demand for over half of the time periods recorded, and the periods that exhibit demand fluctuate largely in value.

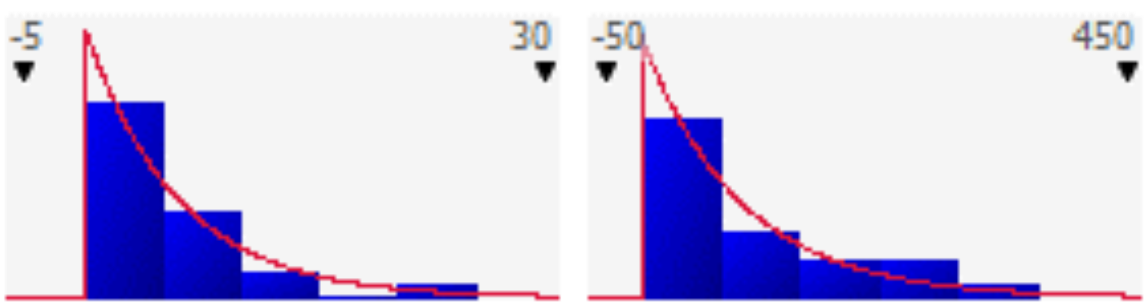

Figure 3: Exponentially Distributed Demand Fitting for "Cough Suppressant Syrup" At Umoja, Kitengela

Given the small quantity of existing data, the distributions inherently have a degree of uncertainty that will lessen over time. To benefit from this, the parameters for the distribution are fit periodically at the time an order is placed to provide the most accurate and up-to-date forecasts. The fitting procedure is carried out automatically by an R script written specifically for Penda Health (as described in the next

[1] Choy, Murphy, and Michelle Cheong. "Identification of Demand through Statistical Distribution Modeling for Improved Demand Forecasting." (n.d.): n. pag. School of Information Systems, Singapore Management University. Web. <http://arxiv.org/ftp/arxiv/papers/1110/1110.0062.pdf>

section). Each drug is fitted to an exponential distribution with a unique rate parameter λ so as to get the most accurate prediction for each drug individually.

When applying these exponential distributions for weekly demand to a monthly ordering policy, it is incorrect to carry out calculations with the standard exponential distribution. This is because a month consists of four weeks; therefore, a monthly order actually takes into account four weeks of demand. Another way of viewing the expected demand for a month is by considering it the sum of four random variables representing weekly demand, each of which follow an exponential distribution. When summing exponential random variables such as this, it is customary to use an Erlang distribution with parameters k equal to the number of random variables to be summed, and λ equal to the rate parameter of the underlying exponential random variables. Thus, our model implements an Erlang distribution for each drug with parameters $k=4$ and $\lambda = \lambda$ where λ is taken from the exponential distribution fitted to that drug.

The R script effectively does all of the necessary calculations for fitting demand distributions for each drug without the client needing to know how it works behind the scenes. The only responsibility of the client is to download 3 separate CSV files from their current POS system, VEND. The first 2 files are the weekly sales data over a time interval of their choosing for their Umoja and Kitengela clinics respectively, and the third file is the current inventory level for each drug. They then need to save them as regular excel files onto their Desktop, and run the script in an R console, which is free to download on the web. The script uses the data on the three files as inputs, and first calculates the number of data points that have zero weekly demand for each drug. If the proportion of zero weekly sales to total data points is too high, the fitted exponential distribution will be over fitting due to the lack of any real data. Thus, the code only fits the exponential distribution to a drug if its proportion of zero weekly sales is less than a certain threshold, which in this case is assigned to be 70%. When fitting the exponential distribution, the script calculates the best-fit lambda parameter using maximum likelihood estimation. For the rest of the drugs, a worst-case analysis approach is implemented. The script finds the top 4 weeks with the highest demand and sums them to get an inferred "worst-case-scenario" monthly sales number. This number is automatically assigned as the order up-to level for these drugs, and no further analysis is needed. Once all the drugs have been analyzed, the information is exported to the blank excel sheet that has 4 sheets. The first two sheets are for the Umoja and Kitengela clinics respectively. It is a list of all the drugs and either has the rate parameter in the column next to it that was fitted using the exponential distribution, or 0 if the drug's OUL was directly computed using the worst-case analysis. For those drugs, that number is located in the 3rd column. The final 2 sheets are the current inventory levels for each drug in the Kitengela and Umoja clinics respectively.

Inventory Management

This model implements a base-stock, or S*, version of the (s,S) inventory model for periodic review. This model was selected for a number of reasons as follows. Firstly, a periodic review model is necessary given Penda's current methods of inventory management. They do not have any form of electronic inventory control such as RFID tracking or barcodes, and therefore rely on manually computing order amounts at a regular interval (typically, every month). This model follows the same format, allowing Penda to only concern itself with inventory-related matters once a month for low-maintenance inventory management.

Typically, a periodic review system comes in the form of an (s,S) policy, where s represents the minimum value that inventory should be observed at before an order is placed, and S is the maximum inventory value that should be achieved. Thus, once inventory is observed at a level below s, and order is placed for S-current inventory to bring the inventory level back up to S. S and s are calculated by considering cost of under stocking and over stocking; however, in this context these costs are intangible (e.g. cost of medical harm induced by not having a drug, cost of having excess storage for inventory that goes unused.) Thus, the model is instead driven by the ideal CSL to be achieved by the inventory policy. When this is the case for a periodic review policy, the s and S values converge to a single value, S*, that is referred to as the order up-to level. At the time of each periodic review, if the inventory level falls below S* then an order is placed for the amount S*-current inventory to bring the level back up to S*. This policy provides optimal catering to service levels while only requiring periodic inspection of inventory. Thus, it is perfectly suited for Penda Health.

The model determines the proper S* level using the cumulative distribution function of the exponential distribution. Considering that a CSL value is equivalent to P(no stock outs in the order period), the model finds the value of the c.d.f. at a given CSL level determined individually for each drug. This provides an ideal order value; however, given orders must be placed in discrete values and often in terms of large pack sizes rather than individual drugs, this ideal order value must be approximated to a multiple of the pack size. This approximation is always either of the nearest multiple of the pack size on either side of the ideal value, and the choice between the two is made by weighing the two pairs of CSL's and costs that are associated with orders of that size. This comparison is made visible to the user through graphical readouts in the interface (see attached Excel file).

The formula for the ideal ordering amount was calculated as follows and the parameters are discussed more in depth in the next section:

GammaInv(CSL, k=4, 1/ λ)

GammaInv: CDF of Gamma distribution

CSL: cycle service

K: shape parameter, chosen as 4 to scale weekly data to monthly periodic ordering

1/ λ: Inverse of rate parameter

Model Parameters

As stated above, each drug is fitted to an exponential distribution and represented by a rate parameter λ. This parameter denotes the specific shape and scale of the exponential distribution that describes each drug individually. In terms of calculations, λ is equal to the inverse mean of the distribution and is thereby fairly simple to calculate. In practice, λ is determined for each drug automatically through the R script that fits the demand distributions. This same value for λ is then passed on to the Erlang distribution to determine the proper order up-to level for a given drug over a four-week period (as described earlier). This is the only statistical parameter that describes the demand distributions; all other parameters in our model apply only to the inventory management section of the model and not to the statistical distribution fitting.

CSL is a parameter in our model, and is determined on an individual basis for each drug. Traditionally, CSL levels are estimated from the ratio of under stocking to cost of under stocking plus the cost of overstocking. In this case however, the cost of under stocking cannot be measured because under stocking one unit of an urgent drug may be a life or death situation. Since there was no appropriate measurement of under stocking, high CSL levels were used in the creation of the model to provide a conservative order policy for Penda. As their mission is to provide high quality care to all patients, the additional cost of extra inventory may be necessary to ensure limited number of stock outs. CSL designates how conservative each drug's policy is, and a higher CSL corresponds to a more conservative policy. Specifically, CSL describes the probability that there will not be a stock out for the drug in question in a given month. Our model allows for CSL to take on three unique values (0.95, 0.9, and 0.85), and the CSL value for each drug is determined by two parameters: weighted average contribution margin and urgency ranking.

Weighted average contribution margin, found by taking the product of contribution margin (sales price-cost) and relative sales revenue, equates to the amount of money each drug contributes to paying down the fixed and variable costs of the business. This was determined as a good measure of how profitable each drug is to the company. Accordingly, our model utilizes this metric in deciding what cycle service level each drug should maintain. The reasoning behind this decision is that highly profitable drugs should be kept in stock and held to a higher service level so that Penda is never turning away customers seeking to buy them. Once weighted average contribution margin is calculated, drugs are designated with a ranking 1-3, where 3 represents highest weighted average contribution margin and 1 represents lowest.

Urgency ranking is a metric describes how essential it is to have the drug on-hand when needed from a medical perspective. It serves as the second factor in deciding what cycle service level each drug should maintain. Urgency is assigned to each drug by the team at Penda Health (although initial rankings are provided in the model). Here again, rankings range from 1-3 with 3 representing the most urgent of drugs to have on hand, and 1 representing the least. The reasoning behind this decision is that as a healthcare provider, Penda would like to have the drugs that it's patients' need on hand. However, if a drug could result in a death if not immediately administered, that drug should obviously be held in higher quantities than a drug that could be administered a week later.

To determine the final CSL for each drug, both the weighted average contribution margin ranking and the urgency ranking are considered. The model takes the maximum of the two rankings, and assigns CSLs according to this value as follows: 3=0.95, 2=0.9, 1=0.85.

RESULTS AND ORDER RECOMMENDATIONS

Using the above formulation, an ideal order amount was able to be determined based off the distribution parameters and CSL levels determined from the WACM and Urgency ranking. However, the supplier Penda uses requires some of the drugs to be ordered in pack sizes, or quantities greater than 1, leading to an extension of the model to incorporate lower and upper bounds to the ideal ordering amount.

For each drug, an algorithm was written to round the ideal order amount to the nearest lower and upper bound pack sizes. From here, the CSL at these bounds were recalculated and compared with the CSL level of the drug when ordering the ideal amount (see Figure 4 and Figure 5).

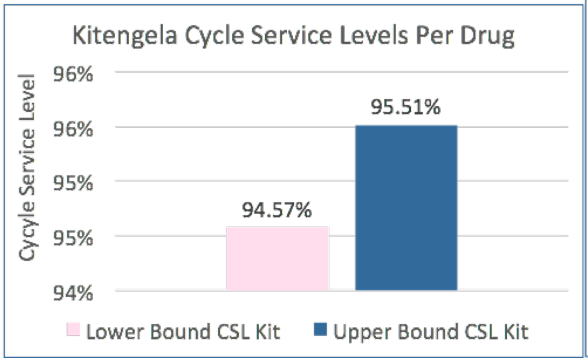

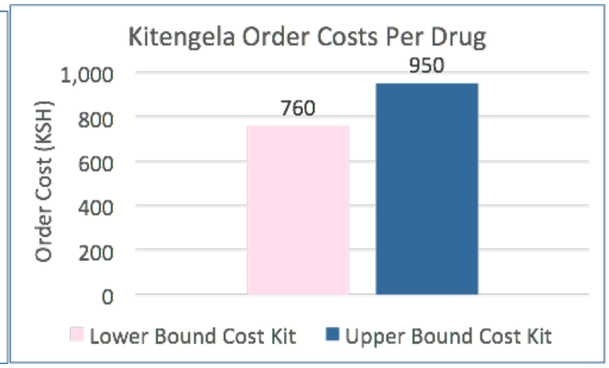

Figure 4 and 5: Amoxiclav - 625 mg Lower and Upper Bounds

A final order recommendation was determined by analyzing the deviation of CSL levels of the lower and upper bounds in comparison to the ideal CSL level. The order amount with the CSL level closest to the ideal CSL level was chosen. It is believed this would be an appropriate measure for comparison as Penda is concerned greatly with patient satisfaction and care. As seen in the figures above, ordering at the lower bound would result in a cost savings of 190 shillings and the CSL level is only slightly reduced.

MODEL VALIDATION

COMPARISON TO CURRENT POLICY

In order to validate the functionality of the model, the recommended order up-to levels for all drugs were compared to their respective historical sales data, which extended for five months-December 1st to April 27th-, with a weekly granularity. The weekly sales for each drug were summed up in each of the five months, and that summation was then compared to the model's recommended order up-to level. If any of the monthly sales for a drug were larger than the recommended order up-to level for said drug, then it would be assumed that a stock out, or a shortage for said drug, would have taken place in that particular cycle.

To begin, all of the drugs were separated into the three target cycle service levels—95%, 90%, and 85%. Then, for each particular drug, the cycles that faced a stock out were summed up and divided by five (the total number of months in the historical data). This number was then subtracted from the number one, to obtain the theoretical achieved CSL for each drug under the recommended order up-to levels. Finally, all of the theoretical achieved CSLs for each of the three target CSL categories were averaged up, and that final number was compared to their target CSL. Ultimately, providing insight into the performance of the proposed model.

In order to compare the performance of the recommended inventory management policy to the current inventory management policy, a method similar to the one explained above was implemented. All of the ~200 pharmaceutical drugs were separated into the same three target CSLs of 95%, 90%, and 85%. Then the inventory levels for each drug in each of the five months was analyzed, and it was assumed that if the inventory level for a drug ever reached the number zero, it meant there was a stock out for said drug. Afterwards the same previously described process to obtain the CSL for each drug was implemented--adding up the number of stocked up cycles, dividing by five, and subtracting from the number one. Finally, the CSL's of each drug for each category were again averaged up. Figure 6 below describes the comparison of the proposed inventory policy to the current inventory policy.

Target CSL	CSL in Current Policy	Historical CSL Achieved in Proposed Policy	CSL Improvement
95%	58%	96%	38%
90%	53%	90%	37%
85%	63%	89%	26%
Avg. for All Drugs	57%	92%	35%

Figure 6: Visual Representation of CSL improvements in Proposed Policy vs. Current Policy

From figure 6, once can appreciate the CSL improvement the proposed inventory policy would provide Penda Health**.** An interpretation of why the CSL improvement is so dramatic in the proposed policy versus the current policy can be further explained by Figure 7.

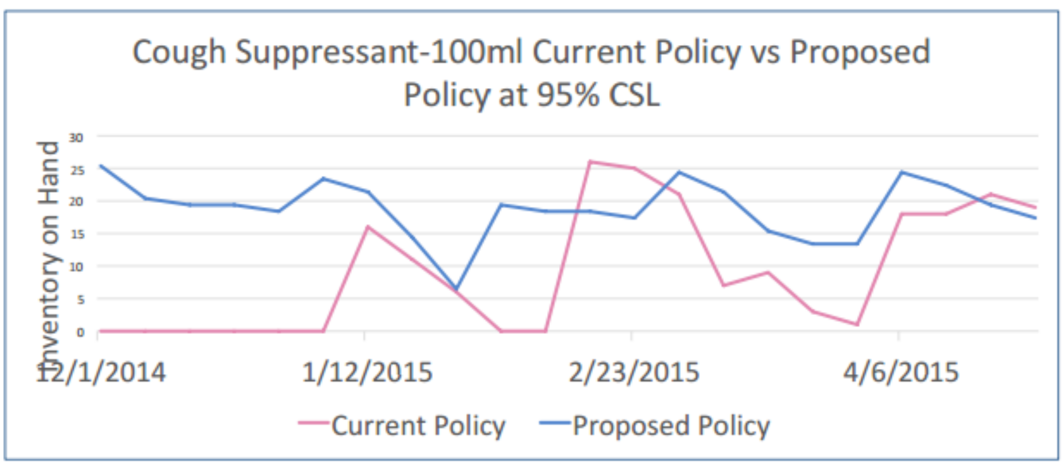

Figure 7: Visualization of a Typical Drug's Inventory Level

From Figure 7, one can deduce that under the current inventory policy a lot of the drugs are ordered once they have stocked out or just before they stock out. Furthermore, the order quantities are also changing, leading to a highly variable--and difficult to manage--inventory policy for Penda Health's pharmaceuticals. Under the proposed policy, the inventory levels for the ~200 drugs would be significantly more stable, and simpler to manage, as one would only have to order up to the recommended level.

SENSITIVITY ANALYSIS

In order to analyze how well the proposed model would fare under increasing demand, the ~200 pharmaceutical drugs were again divided into the same three target CSLs and the total demand of each drug for each month was summed up. This sum was then increased by 5%, 10%, and 15%, and the same process implemented during validation of comparing the monthly demand to the recommended order up-to level to obtain a drugs CSL was implemented. As can be seen in figure 8, the proposed inventory policy is very robust as the achieved CSLs remain very close to the target CSLs even under a 15% increase in demand. This analysis also further corroborates that the proposed inventory model can continue to be used in the future as Penda Health's client base and demand grows.

Target CSL	Historical CSL Achieved with Demand Increase of:			
	0%	5%	10%	15%
95%	96%	96%	94%	94%
90%	90%	90%	88%	86%
85%	89%	87%	87%	87%
Avg. for All Drugs	92%	92%	91%	90%

Figure 8: Visualization of Achieved CSL Under Different Increases in Historical Demand

RECOMMENDATIONS AND FUTURE CONSIDERATIONS

The biggest recommendations stem from more accurate data in terms of current inventory numbers and overall demand. Utilizing the periodic review system is a much easier model to implement for a small business since it does not require continuous monitoring of each drug's inventory level. However, it does necessitate having accurate records of inventory levels since the ordering policy is calculated from: order up-to level - current inventory. Thus, not having accurate inventory levels means ordering an inappropriate amount in any given ordering period, which diminishes the value that this inventory modeling system aims to provide. Looking through VEND and the current inventory numbers Penda has on file, it is clear that there is inaccuracy due to the fact that some drugs are currently listed with negative inventory. In order to derive the maximum benefit, Penda needs to correct their current inventory numbers.

Another key recommendation for the future is to keep track of true demand data. Currently, Penda only records sales data from its customers; however, when modeling inventory management, sales data in only a proxy for demand data, and demand is the true indicator of how much inventory you want to have supplied at all times. Thus, Penda needs to be able to also record customers who inquire about a certain drug, but are not able to purchase it due to stock outs or other complications.
Finally, a key reliance of the model is that order quantity is driven only by a target CSL since the holding and order costs for Penda are intangible. However, if their suppliers change their policy and implement a fixed ordering cost, or if Penda realizes there exist costs that they had not previously considered, this model may not adequately account for these costs. Thus, Penda should remain aware of their suppliers' policy, and enlist the help of their new Inventory Management Specialist to alter this model if necessary.

LIMITATIONS AND CONCERNS

Since Penda is a startup company, recorded sales data only extends five months, preventing our model from incorporating seasonality, a customary practice in inventory modeling. As more data becomes available, a seasonality component may need to be added to the model. For the short term, Penda can incorporate seasonality via a moving average and only utilize the most recent sales data when fitting the model. The model was built utilizing all 5 months of data, but has the capabilities to incorporate less data. For example, if it is believed that sales in July will be very similar to June, Penda can only download June data from Vend and use it to estimate demand in July. As Penda continues to record new historical sales data, the sales data might no longer approximate an exponential distribution, and the proposed model will have to be modified to account for this.

Orders and operations in Kenya are not always reliable; some drug shipments can take from weeks to months to get to the clinics. This large variability in re-order lead times may affect the accuracy and reliability of our model. Another difficulty to consider is that Penda Health is still a startup in only its third year, so there are concerns about their limited supply of data they keep on file, its quality, and how extensive it, creating a challenge for our analysis purposes.

Lastly, in order for the model to work effectively, the inventory data must be true to reality since the recommended order is based off of the current inventory on hand. Accordingly, Penda must ensure they have accurate records of the current inventory levels along with the supply and sales prices for each drug. Inventory records are not easily accessible from Vend, so a data scraping code was utilized to "collect" inventory data. However, after comparison of inventory data to sales data, many weeks contained negative values for inventory. The inaccuracy of inventory data was a significant limitation as it was assumed that sales data is equivalent to true demand. This assumption could not be confirmed by cross-referencing sales with inventory data because of the frequency of negative inventory data points.

CONCLUSION

The dynamic inventory model can be utilized to determine the appropriate order quantity for each drug. After the sales data is fit to the appropriate distribution, the excel model automatically updates to provide the user with the specific order up-to level for all 200 drugs. The current inventory is then used to calculate the order quantity. It is of the utmost importance that current inventory data in vend is accurate since the order quantity is determined from the inventory. It is recommended that Penda conduct an inventory audit prior to placing large orders if they believe their inventory records are not accurate. The inventory model was validated from historical data to ensure accuracy and a sensitivity analysis was conducted to prove robustness of the model. Penda should start tracking unfulfilled demand, providing them with the most accurate demand data, rather than approximating demand with sales data.

APPENDIX A: INVENTORY TERMINOLOGY

This appendix is intended to help the reader of this report better understand the terminology being used within it. Note: All of these definitions and images are gathered from Dr. Seyed Iravani's notes for IEMS 382: Production Planning and Scheduling

Backorder: Unsatisfied demand due to stock-out will be satisfied when inventory become available (i.e., customers wait for the product).

Cycle Service Level: Probability of not having a shortage in a cycle (or during lead time). Also thought of as probability of not having a stock-out in a cycle.

Example:

Order Cycle	Demand	Shortage
1	100	0
2	110	5
3	80	0
4	150	25
5	120	0
6	130	10
7	80	0
8	105	0
9	95	0
10	140	5
Total	**1100**	**45**

In the image above, although the demand shortage is of only around 4%, the Cycle Service Level is of 60% given 4 of the 10 cycles had a shortage.

Continuous Review System: Exact inventory levels are known at any time (using a computerized data base) and appropriate inventory management policy is selected through this information. (See Periodic Review System for further information).

Holding Cost: Cost of holding your inventory product in storage for one year. This number includes the physical costs (insurance, storage space, air conditioning, security, etc.) and it also includes the opportunity cost of holding said product inventory.

Lead Time: Time taken for order to arrive at its destination from the moment it was ordered.

Lost Sales: Unsatisfied demand due to stock-out is lost (i.e., in case of stock-out, customers do not wait, they buy it from the competition).

Order Setup Cost: The cost of setting up an order, including man-hours and shipping costs.

Order Up-To Level: Utilized within the Periodic Review System that uses CSL. Here, the inventory policy proposes the user orders up to a certain level of inventory at the start of every month. Having a predetermined CSL, the inventory should not stock out in more cycles than the predefined CSL.

Periodic Review System: Inventory levels are only known at the end of periods (e.g., end of a day, end of a week, end of a month) In a periodic review system, the inventory policy consists of ordering enough inventory to reach the assigned order up-to level for each item at the end of each period.

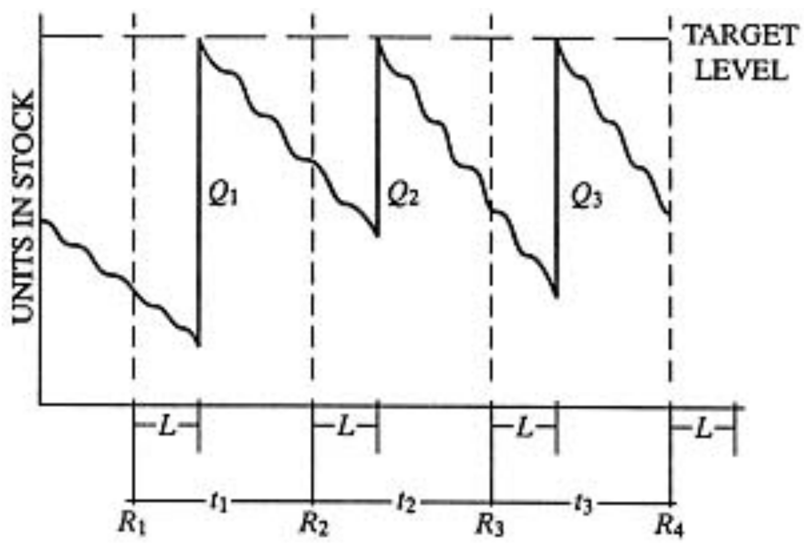

The image above demonstrates how the inventory levels under a periodic review system with order up-to levels would look: Orders are made at the end of each period, and the inventory is always increased to the same level.

APPENDIX B: @RISK DISTRIBUTION OUTPUT

Write here

APPENDIX C: R-CODE AND DOCUMENTATION

```
install.packages("fitdistrplus")
library(fitdistrplus)

install.packages("XLConnect")
library(XLConnect)

install.packages("sqldf")
library(sqldf)

install.packages("rJava")
library(rJava)


#reading in weekly sales numbers from a file that was exported from VEND.
Kit_sales <- readWorksheetFromFile('/Users/zueberjuma/Desktop/vend-item_count-for-type-by-week-kitengela.xlsx',
sheet=1, header=T)
Umoja_sales <- readWorksheetFromFile('/Users/zueberjuma/Desktop/vend-item_count-for-type-by-week-umoja.xlsx',
sheet=1, header=T)

#cleaning up the data a bit. removing columns that are not needed.
Kit_sales <- subset(Kit_sales, select = -c(SKU, Supplier.Code, Brand, Supplier, Revenue, Cost.of.Goods, Gross.Profit, Margin,
Tax))
Umoja_sales <- subset(Umoja_sales, select = -c(SKU, Supplier.Code, Brand, Supplier, Revenue, Cost.of.Goods, Gross.Profit,
Margin, Tax))

#Filtering out the rows that not actual drugs
Kit_sales <- sqldf("select * from Kit_sales where Type='Drugs'")
Umoja_sales <- sqldf("select * from Umoja_sales where Type='Drugs'")


#getting rid of column now that have filtered it
Kit_sales <- subset(Kit_sales, select = -c(Type))
Umoja_sales <- subset(Umoja_sales, select = -c(Type))

#Deleting specific drugs that are either repeated in the dataset, or are not actual true drugs
Kit_sales <- sqldf(c("DELETE FROM Kit_sales WHERE Product='Drugs'", "select * from Kit_sales"))
Kit_sales <- sqldf(c("DELETE FROM Kit_sales WHERE Product='Drug Unlisted'", "select * from Kit_sales"))
Kit_sales <- sqldf(c("DELETE FROM Kit_sales WHERE Product='Tinidazole -'", "select * from Kit_sales"))
Umoja_sales <- sqldf(c("DELETE FROM Umoja_sales WHERE Product='Drugs'", "select * from Umoja_sales"))
Umoja_sales <- sqldf(c("DELETE FROM Umoja_sales WHERE Product='Drug Unlisted'", "select * from Umoja_sales"))
Umoja_sales <- sqldf(c("DELETE FROM Umoja_sales WHERE Product='Tinidazole -'", "select * from Umoja_sales"))

#vector that will hold the names of every drug in our datasets.
#Doing this because the indexing will be easier if the drug names are row headers and not the first column of the dataset.

names_Kit <- c()
names_Umoja <- c()
```

```
#filling the vectors with the names in these loops
for (i in 1:nrow(Kit_sales))
{
names_Kit <- append(names_Kit, Kit_sales[i,1])
}

for (i in 1:nrow(Umoja_sales))
{
names_Umoja <- append(names_Umoja, Umoja_sales[i,1])
}

#getting rid of 1st column that contains the drug name

Kit_sales <- Kit_sales[,-1]

Umoja_sales <- Umoja_sales[,-1]

#renaming row names as the drug names

rownames(Kit_sales) <- names_Kit

rownames(Umoja_sales) <- names_Umoja



#rounding all entries in the dataframe to make sure everything sold is a whole number
for (i in 1:nrow(Kit_sales))
{
for (j in 1:ncol(Kit_sales))
{
Kit_sales[i,j] <- round(Kit_sales[i,j])
}
}

for (i in 1:nrow(Umoja_sales))
{
for (j in 1:ncol(Umoja_sales))
{
Umoja_sales[i,j] <- round(Umoja_sales[i,j])
}
}


Kit_sales$count <- 0                    #count is the number of weeks that had zero sales for each drug in timeframe
Umoja_sales$count <- 0
Kit_sales$delete <- 0                   #delete is a binary indicator that is 1 if said drug has too many zeros to fit an
exponential distribution
Umoja_sales$delete <- 0



#this loop goes through every column for each drug and adds up the number of zero demand sales weeks for each drug.
#if the percentage of zero sales weeks is greater than the threshold (70%), the delete indicator gets a 1 for said drug

counter=0
for (i in 1:nrow(Kit_sales))
{
for (j in 1:(ncol(Kit_sales)-3))
{
if (Kit_sales[i,j] == 0)
{
```

```
counter = counter+1
}
}
Kit_sales[i,(ncol(Kit_sales)-1)] = counter
counter=0
}

threshold_zero_Kit <- (.7*(ncol(Kit_sales)-2))
name_kill_Kit <- c()

for (i in 1:nrow(Kit_sales))
{
if (Kit_sales[i,(ncol(Kit_sales)-1)] > threshold_zero_Kit)
{
name_kill_Kit <- append(name_kill_Kit, names_Kit[i])
Kit_sales[i,ncol(Kit_sales)] <- 1
}
}

counter=0
for (i in 1:nrow(Umoja_sales))
{
for (j in 1:(ncol(Umoja_sales)-3))
{
if (Umoja_sales[i,j] == 0)
{
counter = counter+1
}
}
Umoja_sales[i,(ncol(Umoja_sales)-1)] = counter
counter=0
}

threshold_zero_Umoja <- (.7*(ncol(Umoja_sales)-2))
name_kill_Umoja <- c()

for (i in 1:nrow(Umoja_sales))
{
if (Umoja_sales[i,(ncol(Umoja_sales)-1)] > threshold_zero_Umoja)
{
name_kill_Umoja <- append(name_kill_Umoja, names_Umoja[i])
Umoja_sales[i,ncol(Umoja_sales)] <- 1
}
}

#this code is the meat of the code. it goes through each drug and if the delete indicator had a 1, it will define an OUL in a different way than fitting the exponential distribution
#instead, it adds up the 4 biggest sales points the drug ever had in the timeframe, and adds them up to replicate a "worst case scenario for the next months demand projections.
#if the drug has a 0 for the indicator variable, then the exponential distribution is fit, and the optimal lambda parameter is noted. The results are then outputted to the
#blank excel sheet for reference.


OUL_Kit <- list()

OUL_Umoja <- list()


vec <- c()
```

```
max1 <- 0

max2 <- 0

max3 <- 0

max4 <- 0

index <- c(0,0,0)


for (i in 1:nrow(Kit_sales))
{
if (Kit_sales[i, ncol(Kit_sales)] == 1)
{
for (j in 1:(ncol(Kit_sales)-3))
{
if (Kit_sales[i,j] > max1)
{
max1 = Kit_sales[i,j]
index[1] = j
}
}
for (j in 1:(ncol(Kit_sales)-3))
{
if ((Kit_sales[i,j] > max2) && (Kit_sales[i,j] <= max1) && (j != index[1]))
{
max2 = Kit_sales[i,j]
index[2] = j
}
}
for (j in 1:(ncol(Kit_sales)-3))
{
if ((Kit_sales[i,j] > max3) && (Kit_sales[i,j] <= max2) && (j != index[1]) && (j != index[2]))
{
max3 = Kit_sales[i,j]
index[3] = j
}
}
for (j in 1:(ncol(Kit_sales)-3))
{
if ((Kit_sales[i,j] > max4) && (Kit_sales[i,j] <= max3) && (j != index[1]) && (j != index[2]) && (j != index[3]))
{
max4 = Kit_sales[i,j]
}
}

vec <- append(vec,(max1+max2+max3+max4))
vec <- append(vec,i)
OUL_Kit <- append(OUL_Kit, list(vec))

vec <- c()
max1 <- 0
max2 <- 0
max3 <- 0
max4 <- 0
index <- c(0,0,0)
}
}
```

```
vec <- c()
max1 <- 0
max2 <- 0
max3 <- 0
max4 <- 0
index <- c(0,0,0)


for (i in 1:nrow(Umoja_sales))
{
if (Umoja_sales[i, ncol(Umoja_sales)] == 1)
{
for (j in 1:(ncol(Umoja_sales)-3))
{
if (Umoja_sales[i,j] > max1)
{
max1 = Umoja_sales[i,j]
index[1] = j
}
}
for (j in 1:(ncol(Umoja_sales)-3))
{
if ((Umoja_sales[i,j] > max2) && (Umoja_sales[i,j] <= max1) && (j != index[1]))
{
max2 = Umoja_sales[i,j]
index[2] = j
}
}
for (j in 1:(ncol(Umoja_sales)-3))
{
if ((Umoja_sales[i,j] > max3) && (Umoja_sales[i,j] <= max2) && (j != index[1]) && (j != index[2]))
{
max3 = Umoja_sales[i,j]
index[3] = j
}
}
for (j in 1:(ncol(Umoja_sales)-3))
{
if ((Umoja_sales[i,j] > max4) && (Umoja_sales[i,j] <= max3) && (j != index[1]) && (j != index[2]) && (j != index[3]))
{
max4 = Umoja_sales[i,j]
}
}

vec <- append(vec,(max1+max2+max3+max4))
vec <- append(vec,i)
OUL_Umoja <- append(OUL_Umoja, list(vec))

vec <- c()
max1 <- 0
max2 <- 0
max3 <- 0
max4 <- 0
index <- c(0,0,0)
}
}


Kit_sales <- subset(Kit_sales, select = -c(count))
Umoja_sales <- subset(Umoja_sales, select = -c(count))
```

```
Kit_sales$rate <- 0
Kit_sales$OUL <- 0
Umoja_sales$rate <- 0
Umoja_sales$OUL <- 0



for (i in 1:nrow(Kit_sales))
{

if (Kit_sales[i,(ncol(Kit_sales)-2)] != 1)                    #testing whether “delete” is 1 or 0
{
a <- fitdist(as.numeric(Kit_sales[i,-((ncol(Kit_sales)-3):ncol(Kit_sales))]), "exp", method="mle", discrete=F)
Kit_sales[i,(ncol(Kit_sales)-1)] <- coef(a)
}
else
{
for (j in 1:length(OUL_Kit))
{
if (i == (OUL_Kit[[j]][2]))
{
Kit_sales[i,ncol(Kit_sales)] = OUL_Kit[[j]][1]
}
}
}
}

for (i in 1:nrow(Umoja_sales))
{

if (Umoja_sales[i,(ncol(Umoja_sales)-2)] != 1)                    #testing whether “delete” is 1 or 0
{
a <- fitdist(as.numeric(Umoja_sales[i,-((ncol(Umoja_sales)-3):ncol(Umoja_sales))]), "exp", method="mle", discrete=F)
Umoja_sales[i,(ncol(Umoja_sales)-1)] <- coef(a)
}
else
{
for (j in 1:length(OUL_Umoja))
{
if (i == (OUL_Umoja[[j]][2]))
{
Umoja_sales[i,ncol(Umoja_sales)] = OUL_Umoja[[j]][1]
}
}
}
}

Kit_sales <- subset(Kit_sales, select = c(rate,Item_Count,OUL))
Umoja_sales <- subset(Umoja_sales, select = c(rate,Item_Count,OUL))

writeWorksheetToFile('/Users/zueberjuma/Desktop/CodeOutput.xlsx', Kit_sales, sheet="Kitengela", clearSheets=F,
rownames="rownames(Kit_sales)")
writeWorksheetToFile('/Users/zueberjuma/Desktop/CodeOutput.xlsx', Umoja_sales, sheet="Umoja", clearSheets=F,
rownames="rownames(Umoja_sales)")


inv <- readWorksheetFromFile('/Users/zueberjuma/Desktop/vend-inventory_count-for-type-by-month.xlsx', sheet=1,
header=T)
```

```
trim.trailing <- function (x) sub("\\s+$", "", x)
inv$Product <- trim.trailing(inv$Product)

inv_Kit <- sqldf("select Product,Count from inv where Type='Drugs' and Outlet='Kitengela'")
inv_Umoja <- sqldf("select Product,Count from inv where Type='Drugs' and Outlet='Umoja'")

writeWorksheetToFile('/Users/zueberjuma/Desktop/CodeOutput.xlsx', inv_Kit, sheet="Kitengela Current Inventory",
clearSheets=F)
writeWorksheetToFile('/Users/zueberjuma/Desktop/CodeOutput.xlsx', inv_Umoja, sheet="Umoja Current Inventory",
clearSheets=F)
```

APPENDIX C: EXCEL MODEL

[See Attached File – description of file on Home page]

WORKS CITED

Choy, Murphy, and Michelle Cheong. "Identification of Demand through Statistical Distribution Modeling for Improved Demand Forecasting." (n.d.): n. pag. *Arxivhttp://arxiv.org/ftp/arxiv/papers/1110/1110.0062.pdf*. Singapore Management University. Web.

"The Gamma Distribution." *Workshop in Applied Phylogenetics*. N.p., 26 Mar. 2013. Web. 03 June 2015.

"How to Figure Out the Weighted-Average Contribution Margin." *Small Business*. N.p., n.d. Web. 1 June 2015.