USAID | DELIVER PROJECT

Malaria Seasonality and Calculating Resupply

Applications of the Look-Ahead Seasonality Indices in Zambia, Burkina Faso, and Zimbabwe

OCTOBER 2014

This publication was produced for review by the U.S. Agency for International Development. It was prepared by the USAID | DELIVER PROJECT, Task Order 7.

USAID FROM THE AMERICAN PEOPLE

PRESIDENT'S MALARIA INITIATIVE

Malaria Seasonality and Calculating Resupply

Applications of the Look-Ahead Seasonality Indices in Zambia, Burkina Faso, and Zimbabwe

USAID | DELIVER PROJECT, Task Order 7
This document was prepared by staff of the USAID | DELIVER PROJECT, Task Order 7, which is funded by the U.S. Agency for International Development (USAID) under contract number GPO-I-00-06-0007-00, order number AID-OAA-TO-11-00012, beginning on March 28, 2011. Task Order 7 is implemented by John Snow, Inc., in collaboration with 3i Infotech, Inc.; Crown Agents USA, Inc.; FHI 360; Foundation for Innovative New Diagnostics; Logenix International, LLC; The Manoff Group, Inc.; MEBS Global Reach, LC; PATH; PHD International (a division of the RTT Group); Population Services International; Social Sectors Development Strategies, Inc.; UPS Supply Chain Solutions, Inc.; and VillageReach. Task Order 7 supports USAID's goal of reducing the malaria burden in sub-Saharan Africa by procuring and delivering safe, effective, and high-quality malaria commodities; by providing technical assistance and on-the-ground logistics expertise to strengthen in-country supply systems and build capacity for managing commodities; and by improving the global supply and long-term availability of malaria commodities.

Recommended Citation
Watson, Noel, Loren Bausell, Andrew Ingles, and Naomi Printz. 2014. *Malaria Seasonality and Calculating Resupply: Applications of the Look-Ahead Seasonality Indices in Zambia, Burkina Faso, and Zimbabwe.* Arlington, Va.: USAID | DELIVER PROJECT, Task Order 7.

Abstract
This concept paper describes an approach for enhancing the simple average monthly consumption (AMC) rule to handle seasonal commodities, while maintaining some of its simplicity to continue to meet complex needs of the developing country settings where it would be used. The approach operationally involves multiplying the AMC by indices that compensate for seasonality—referred to as Look-Ahead Seasonality Indices (LSI)—before multiplying by the maximum stock level. The LSI approach is tested through three case studies, in Zambia, Zimbabwe, and Burkina Faso.

Cover photo: Zimbabwe, 2009. Photograph by Allison Belemvire.

USAID | DELIVER PROJECT
John Snow, Inc.
1616 Fort Myer Drive, 16th Floor
Arlington, VA 22209 USA
Phone: 703-528-7474
Fax: 703-528-7480
Email: askdeliver@jsi.com
Internet: deliver.jsi.com

Contents

Appendices

Figures

Tables

Acronyms

AL	artemether/lumefantrine
AMC	average monthly consumption
AQ	amodiaquine (malaria)
AS	artesunate (malaria)
CHAZ	Churches Health Association of Zambia
EMLIP	Essential Medicines Logistics Improvement Program
FDC	fixed-dose combination
GDP	gross domestic product
LSI	Look-Ahead Seasonality Index
MAPE	mean absolute percentage error
MOS	months of stock
PHCP	primary healthcare packages
PMI	President's Malaria Initiative
SOH	stock on hand
UNICEF	United Nations Children's Fund
USAID	U.S. Agency for International Development
ZIP	Zimbabwe Informed Push

Acknowledgments

The authors of this document would like to thank members of the Zambia field office for their contributions to and encouragement for a revised resupply approach for malaria commodities, particularly Wendy Nicodemus, Gamariel Simpungwe, and Jeremy Sikazwe. Thanks also for the support and valuable perspectives of the Burkina Faso and Zimbabwe field offices, particularly Parfait Edah, Ernest Wendemanagde Ouedraogo, and Tinei Chitsike. In Arlington, Virginia, thanks to Laila Akhlaghi and Kate Wolf for their technical review and Gregory Davidson for mapping assistance. Lessons drawn from the USAID | DELIVER PROJECT's experience in managing the resupply of malaria supply chains would not have been possible without these valuable contributions.

USAID contracts funded the technical assistance and assessment process that produced the material and lessons contained in this document. We are grateful to the professionals in the President's Malaria Initiative for their encouragement and commitment to improving malaria and public health programs through logistics.

Executive Summary

The simple average monthly consumption (AMC) rule for inventory replenishment estimates average consumption from the most recent months of consumption (typically the past three months). To estimate the quantity to order, AMC can be multiplied by the desired maximum stock level, then subtracting the stock on hand. The simple AMC rule has been a staple of inventory systems introduced in healthcare supply chains in many developing countries, and its simplicity and intuitive nature have undoubtedly contributed to its prevalence and popularity. Theoretically, the approach is appropriate for demand patterns exhibiting strong serial correlation from one period to the next, but can also be used for demand patterns that are generally level. However, there are healthcare commodities, such as malaria commodities, that exhibit strong seasonality, and the simple AMC rule is less effective for these commodities.

The management of malaria commodities is subject to the patterns of the disease, including seasonal fluctuations in some areas, with increased cases during the rainy season. As a result, during the rainy season, demand and use of malaria medicines and tests may increase. Therefore, because of the seasonal nature of malaria, using only the past three months of data may underestimate annual demand if it is calculated at the beginning of the malaria season, and overestimate it if calculated toward the end of the season. As a result, the inventory replenishment quantities recommended by the simple AMC rule tend to lag seasonal changes rather than precede them, which may not be appropriate for inventory management of seasonal commodities. It may be more accurate to use alternative methods for calculating resupply.

This concept paper describes an approach for enhancing the simple AMC rule to handle seasonal commodities, while maintaining some of its simplicity to continue to meet the complex needs of the developing country settings where it would be used. The approach operationally involves multiplying the AMC by indices that compensate for seasonality—referred to as Look-Ahead Seasonality Indices (LSI)—before multiplying by the maximum stock level. The LSI approach is intended to provide a practical, easy-to-use, enhanced means of calculating resupply for seasonal commodities, at all levels of the supply chain. Because of its similarity to the simple AMC rule, the LSI approach is one that should be relatively easy to introduce into healthcare settings in developing countries. Given its simplicity, it should also be relatively easy to integrate into electronic-based information systems.

The concept paper examines many of the challenges for addressing resupply of seasonal commodities. The LSI approach, at least partially if not completely, addresses all of these concerns. In three country case studies, testing antimalarial consumption data from Zambia, Zimbabwe, and Burkina Faso, the LSI approach's performance was commendable and could be recommended both as an inventory replenishment mechanism, under conditions that correspond with the healthcare settings in developing countries, as well as a forecasting method. In all country case studies, the simple AMC rule was outperformed for addressing seasonality, in comparison to the LSI approach.

For all countries, the LSI approach was evaluated in comparison to four other common models (the simple AMC rule, simple exponential smoothing, triple exponential smoothing, and triple exponential smoothing with fixed seasonality index), and evaluated in two ways: one as an inventory replenishment mechanism using the allocated inventory costs, and the other as a forecasting method using the mean absolute percentage error (MAPE). As a result of this process, each country case study provides an estimate of the seasonality index and corresponding LSI that could support application of the LSI approach, or similar approaches based on seasonality indices within these countries.

In Zambia, the LSI approach performed best in comparison with the other models, in part because the availability of monthly consumption data over nearly a two-year time period facilitated analysis. In Zimbabwe, the LSI performed well in the tests, particularly those based on allocated inventory costs. As a pure forecasting method based on its MAPE scores, the LSI performance was not as comparable to the results of the Zambia case study. However, we would still conclude that the LSI can be recommended for use in Zimbabwe. A possible limitation of the Zimbabwe data was its availability only on a quarterly basis.

A main difference between the country case studies was the issue of data availability, particularly in Burkina Faso. Special care must be taken with Burkina Faso's results since, although the LSI approach performed well in testing, it would be preferred that the supporting seasonality index had been based on more robust consumption data, and that testing of the LSI approach had been more extensive.

Because of its reliance on historical consumption data, the LSI approach might in theory be applied in other countries where the President's Malaria Initiative (PMI) or the National Malaria Control Programs (NMCPs) have access to consistent, relatively robust consumption data, and the simple AMC rule is currently in place for inventory replenishment. A limitation of the LSI approach is its capacity to look too far ahead in the future (e.g., one year). The approach may be best suited for resupply calculations, as well as shorter term forecasts, as the further ahead that the LSI is used to look ahead, the greater the chance that level of the period may have changed compared to that of the most recent consumption. An assessment of the LSI approach may also be of value, in order to test the method, its practicality, and its results in a controlled field setting, in comparison to current resupply methods.

As any quick scan of the literature reveals, there are many ways to approach both inventory replenishment and forecasting. The approach introduced and assessed here is only one such method. The discussions here should be considered a case for the legitimacy of the LSI approach for seasonal commodities, but not to the exclusion of other approaches, or that the suitability of the LSI approach should be taken for granted in all situations involving seasonal commodities. As always, the particular context of the country situation should first be carefully considered in the testing and potential application of the LSI approach.

Introduction

The simple average monthly consumption (AMC) rule for inventory replenishment estimates average consumption from the most recent months of consumption (typically the past three months). To estimate the quantity to order, AMC can be multiplied by the desired maximum stock level, then subtracting stock on hand. The simple AMC rule has been a staple of inventory systems introduced in healthcare supply chains in many developing countries. However, inventory systems that use the simple AMC rule can struggle with seasonal commodities, because of the properties of the resupply method that are implicit in the approach.

Figure 1 captures the main challenges with seasonal commodities when the simple AMC rule is used at the health facility level of healthcare systems. While this graph was created by simulation, we see a similar profile at many facilities in the field. The line chart shows actual consumption at one facility, with each period representing two months. Seasonality is apparent, with a peak season and off-peak season, with a length of six months each, or three periods. The bar chart captures inventory on hand at the end of each period.

Figure 1. Challenges for Use of the Simple AMC Rule for Inventory Replenishment of Seasonal Commodities

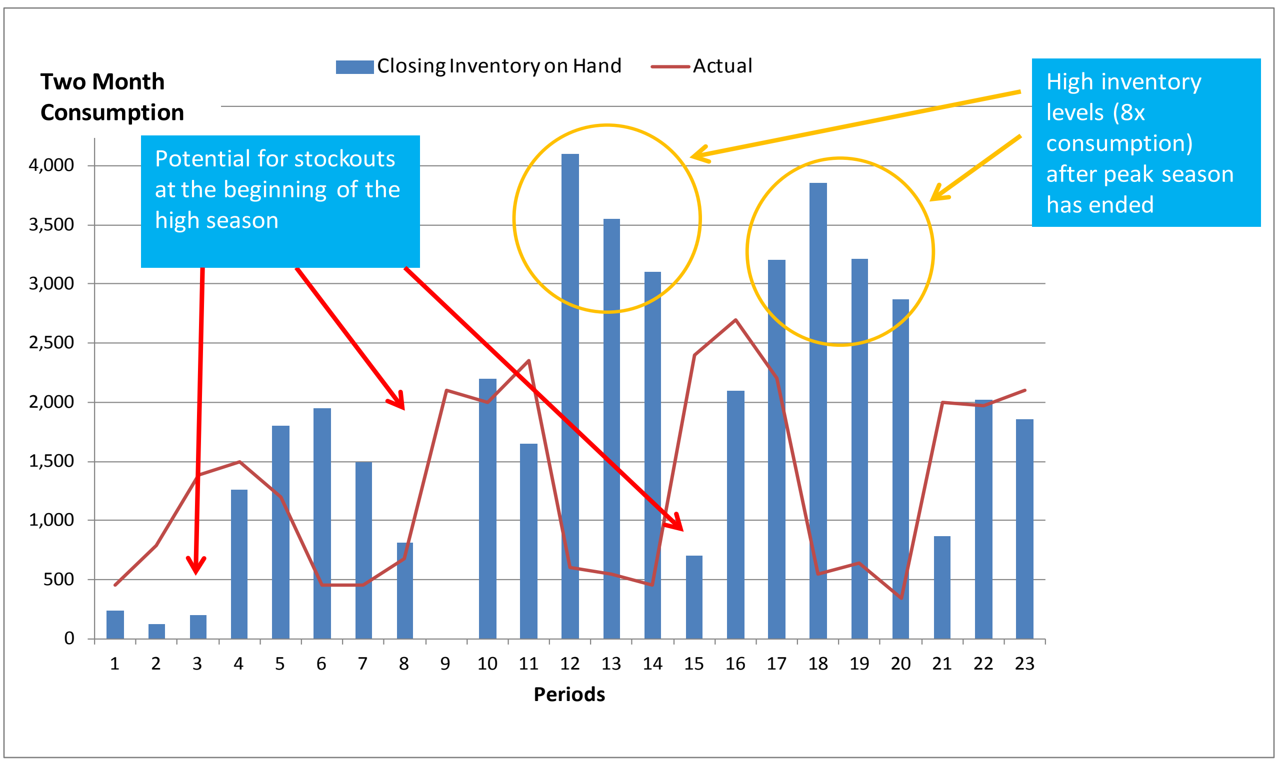

Main Challenges with the AMC Rule?

For inventory replenishment based on the simple AMC rule, there is the potential for stockouts at the beginning of the high/peak season, primarily because the simple AMC rule tends to react to demand rather than prepare for demand.

In addition to this potential for stockouts at the beginning of the peak season, there is also a tendency for high levels of inventory to result after the peak season and sustain well into the low season. In the simulation in figure 1, we see levels of inventory as high as eight times consumption levels in the low season. Such high levels of inventory are often similar to what we find in the field.

At any given level of the supply chain, the simple AMC rule incorporates two different components: a forecasting component and an inventory replenishment component. The *forecasting component* is similar to the moving average forecast, which uses a rolling forecast of the most recent periods as the forecast for consumption going forward (Makridakis, Wheelwright, and McGee 1983). The moving average is suited to demand patterns that exhibit strong serial correlation from one period to the next. The moving average can also be reasonably used for demand that has no trend, or is generally level over time. The *inventory replenishment component* of the simple AMC is an order up to mechanism, facilitated by a target created by multiplying the AMC by the maximum stock level. Ideally, order up to mechanisms are appropriate for inventory replenishment in settings where inventory is reviewed periodically, and this characterizes many of the healthcare settings in developing countries. Primarily, the simple AMC rule's challenges with seasonal commodities stem from the moving average forecasting approach inherent in its design.

This concept paper describes an approach for enhancing the simple AMC rule to handle seasonal commodities, while maintaining some of its simplicity to continue to meet complex needs. Based on experience in field settings, a rule is needed that is still intuitive enough that people can grasp, understand, and explain to others. The revised approach operationally involves multiplying the AMC by indices that compensate for seasonality, referred to as Look-Ahead Seasonality Indices (LSI), before multiplying by the maximum stock level.

There are many ways to approach both inventory replenishment and forecasting, and the LSI approach introduced and assessed here is only one such method. The discussions here should be considered an argument for the legitimacy of the LSI approach for seasonal commodities, but not to the exclusion of other approaches, or that the suitability of the LSI approach should be taken for granted in all situations involving seasonal commodities.

Inventory Replenishment and Forecasting with Seasonality

The basic inventory management approach for addressing seasonality involves understanding the timing within the seasonal cycle and how consumption is expected to change over the near term. For example, at the facility level, are we at the beginning of the peak season? Or for higher tiers in the supply chain, are we close to the peak season? In terms of execution, should we start ramping up inventory levels, and if so, by how much? Similar questions exist for the end of the peak season. The approach to forecasting is crucial to providing the understanding of seasonality that supports the inventory management approach.

Typically, approaches to forecasting seasonality are based on historical data and generate a numeric representation of the consumption's general profile over the seasonal cycle. For example, for a commodity with an annual seasonal pattern, if the monthly consumption quantities over a year are divided by the consumption in the first month, then the resulting ratios still capture the general shape of the annual consumption pattern. The resulting ratios also give the relationship of consumption in a particular month to that of the first. Such a set of ratios is usually referred to as a set of *seasonality indices*. Table 1 provides an example of the calculation process for creating seasonality indices.

Table 1. An Example of Calculating Seasonality Indices for Our Primary Health Facility

Month	Consumption (units)	Seasonality Index Calculation	Seasonality Index
Jan–Feb (reference month)	500	= 500 / 500	1.00
Mar–Apr	865	= 865 / 500	1.73
May–Jun	1,515	= 1,515 / 500	3.03
Jul–Aug	1,645	= 1,645 / 500	3.29
Sept–Oct	1,315	=1,315 / 500	2.63
Nov–Dec	500	= 500 / 500	1.00

Forecasting approaches generally make use of these seasonality indices to remove or reintroduce seasonality as needed when generating a forecast. Generally, existing consumption data would have the seasonality removed (de-seasonalized) before additional analysis and extrapolation are done to create an intermediate forecast. Seasonality would then be added (seasonalized) to create the final forecast (Makridakis et al. 1983). A simple form of de-seasonalizing data is that of dividing by the appropriate seasonality index; seasonalized data can be created by multiplying by the appropriate seasonality index. An example of a well-known forecasting approach for seasonality is Holt-Winter's Triple Exponential Smoothing (see appendix A), which incorporates trend and seasonality information using a seasonality index, but allows both the trending parameters and the seasonality index to be continuously updated as each period of consumption occurs.

Motivation for Look-Ahead Seasonality Indices

The motivation for the LSIs comes from a simple approach of using history to predict the future, or to look ahead and compare its performance to the simple AMC rule. Consider an example of the simple AMC rule, which creates a maximum stock level quantity target based on consumption in the previous two-month period, and a maximum stock level multiplier equal to two, giving four months of consumption.

If consumption follows that of the "actual" line chart in figure 2 (which is taken from the first nine periods of consumption in our example in figure 1), we can see that we will have some stockout concerns in period 9, when consumption is about 2,100 units, but the maximum stock level quantity target will only be about 1,500, since consumption in period 8 is about 750 units. A simple example of using history to look ahead would, at the beginning of the two-month period #9 (months five and six [May and June] of year two), look at consumption one year before in period #3 (months five and six [May and June] of year one). Period #3 covers the same May and June two-month period as period #9, and hopefully the same time with the seasonality cycle, but one year earlier. The top-up level target would be 2,800, which is double the consumption in period #3 (about 1,400 units).

Figure 2. Motivating Example for Look-Ahead Seasonality Indices

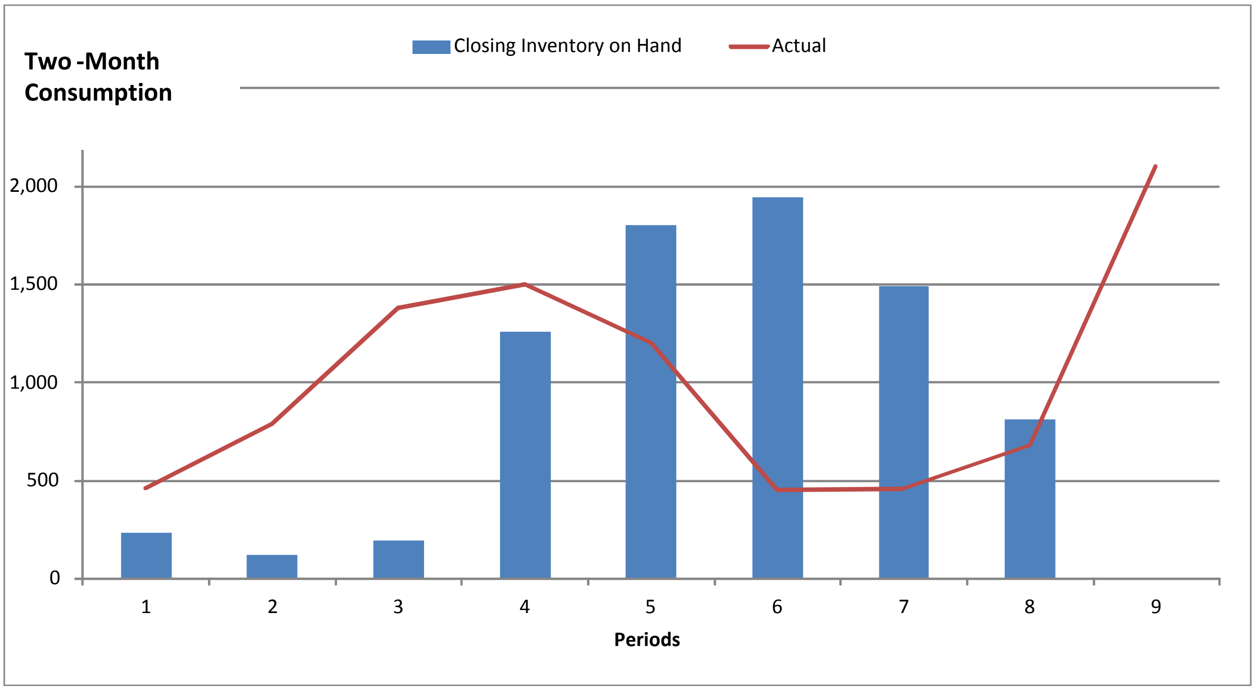

Figure 3 puts this idea of using history to look ahead and compares it to the simple rule of using the previous period's consumption for the AMC. For the approach of using history to look ahead, we use a three-period average of consumption one year ago. For both approaches, the same maximum stock level is used to multiply.

What we see first is that until we have historical data, we cannot use our new approach. (How to deal with limitations of historical data, including the case of lost consumption data due to stockouts, is addressed later in this document.) But with this history, there is still some potential for stockouts at the beginning of the peak season, although we seem to do a little better than the simple AMC rule. Still, we see in period #11 (September and October of year two) that we do stockout using the historical data to use the look-ahead approach, but not using the simple AMC rule. That's because consumption in the peak season in the current year is on average higher than consumption in the peak season one year prior, and we have no way of accounting for that. The look-ahead approach definitely seems to do better with inventory during the low season as compared to the simple AMC rule.

Figure 3. Comparing AMC Rule to Using Historical Consumption to Look Ahead

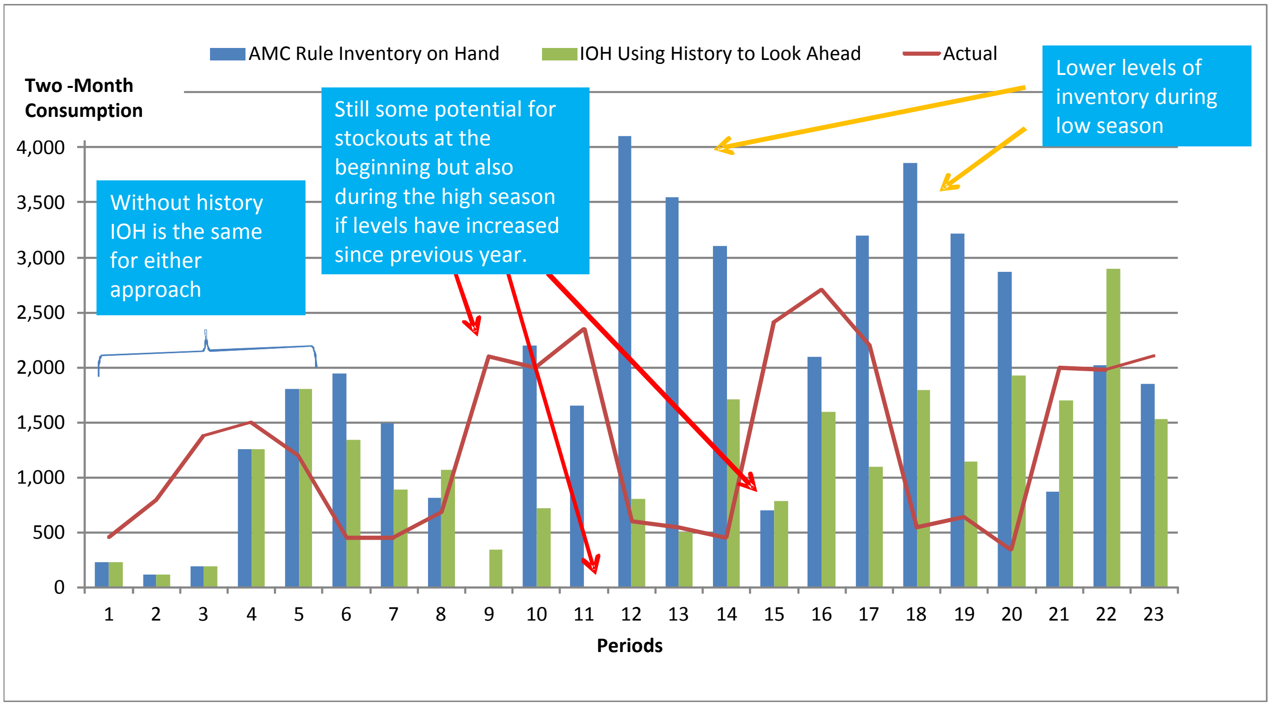

Our simple example of using history to look ahead seems to be moving in the right direction, but it still has some issues. Those challenges are listed as follows and will be used to help evaluate the LSI approach once we introduce it:

- There does not seem to be a mechanism to deal with changes in levels from one year to the next, especially when consumption increases.
- What happens with lost consumption due to stockouts?
- In figure 2, what if the two months in period #9 (May and June, year two) are at a slightly different place in the seasonal cycle (e.g., due to rainy season starting early or late)?
- Can the approach be used if we have no historical consumption data?
- What about higher tiers in the supply chain? Would they use the same rule?
- Can we use this approach to help plan for facilities that regularly get cutoffs from routine inventory replenishment?
- How would we operationalize the approach within existing healthcare systems in developing countries?

Look-Ahead Seasonality Indices

By strict definition, the LSI is a set of supply chain, tier-specific indices that adjust current consumption rates to create a look-ahead forecast for consumption by incorporating seasonality and current levels of consumption.

A set of seasonality indices gives the relative difference in consumption across the season. The LSI is essentially a ratio of averages of seasonality indices: the average seasonality index for the periods that you are trying to predict, and the average seasonality index for the periods of recent consumption that you plan to use for the forecast.

The LSI is supply chain tier specific. As a result, LSI corresponds to tier-specific max/min levels. The LSI also incorporates the lead time for inventory replenishment to the specific tier, as well as its inventory review period. (An explanation of how varying lead times and review periods can be incorporated in the LSI approach will be explained in the section "Incorporating LSI at higher tiers in the supply chain," and appendix B also provides a general formula.)

Lead time is the time between when new stock is ordered from the vendor and when it is received and available for use, and the review period affects how much inventory is ordered at a time; the two concepts are independent of each other. (Lead time varies by tier of the supply chain. However, no matter what the tier, stock is considered "received and available for use" when that particular level of the supply chain can actually use them, whether that "use" may be to hold in storage or push to the next level of the supply chain [at a regional warehouse, for example] or provide commodities to patients [at the facility level].) If a tier of the supply chain has a four-month lead time, then the forecasts for this tier needs to look ahead four months when placing an order, as compared to another tier in the supply chain that might get its order immediately and only has to forecast one period ahead.

The inventory review period also varies by tier of the supply chain and is the routine length of time between periodic assessments of inventory to determine if additional stock is needed. If the review period of one tier is two months (meaning you order every two months and so your order must cover at least two months of inventory), while another is only one month, then the LSI for the former should look out over two months, while the LSIs for the latter should look out over one month. Lead times have to be less than the review period, and many facilities below the central medical store have a low lead time compared to the review period.

In most cases, lead times at the central level are greater than lead time between levels in country. For example, the central medical store may have a longer lead time, especially if procuring from other countries. Max/min levels are usually based on the lead time and review period. At the highest level, they are also for procurement planning, which is not only forecasting, but also accounting for any quantities required to fill the pipeline.

Using the LSI for inventory replenishment for different tiers is fairly straightforward once calculated. In the simple AMC rule, we would take the product of consumption and the maximum stock level multiplier, subtracting stock on hand (SOH).

Determining order quantity using simple AMC:

Order quantity = AMC × max – SOH

Using LSI, determining the maximum stock level quantity involves taking the product of average recent consumption and LSI for the current point in the seasonality profile. This provides the look-ahead forecast for consumption, which is then multiplied by the maximum stock level, subtracting SOH.

Determining order quantity using LSI:

Order quantity = AMC × LSI × max – SOH

A couple points can be made about the LSI approach already:

First, this approach incorporates the current level of the consumption, because recent consumption is used to estimate the look-ahead forecast, and the LSI is a ratio that accommodates changes in levels from one year to the next. So, if recent consumption exhibits a level change compared to history, it will be incorporated into the look-ahead forecast.

Second, the form of the LSI maximum stock level target is one that should be easy to introduce into healthcare settings in developing countries, because it is similar to the simple AMC rule. Given its simplicity, it should also be easy to integrate into electronic-based information systems. The formula can even be used for non-seasonal commodities where appropriate, by setting the LSIs equal to 1.

Creating LSIs

As an example of how to create a set of LSIs for inventory replenishment, consider creating one for a primary health facility. Here, we assume that the health facility reviews its inventory every two months and lead time for receiving inventory is fairly small compared to that review period (e.g., only a couple of days).

We start with seasonality indices, which give the relative difference in consumption across the season. (Refer to page 3 for an example of determining seasonality indices.) One way to create these indices is to take consumption for each period in the year, and divide by consumption in a reference period. An equivalent set of seasonality indices would be generated if we divide by the average consumption in a period. In our example, we take the first six two-month periods of simulated historical consumption in figure 1, and divide actual consumption by consumption in the first period.

This gives the seasonality indices charted in figure 5 and listed in table 1. We can interpret these indices in a straightforward manner. For example, we generally expect consumption in the third period to be three times consumption in the first period. In addition, we can take the ratio between any two periods, and that allows prediction of consumption of one period versus another, as opposed to always having to use the reference period. That is the basic idea of the LSI as a ratio of seasonality indices.

Figure 4. Formula for LSI in a Period for Our Primary Health Facility

Three-period average of seasonality indices centered at period *i* (for which consumption hasn't happened yet)

Additional indices add some correction if seasonality has shifted a bit earlier or later

$$LSI_i = \frac{Ave(s_{i-1}, s_i, s_{i+1})}{Ave(s_{i-3}, s_{i-2}, s_{i-1})}$$

If the formula for LSI given in figure 4 looks a bit daunting, it isn't really. The formula basically states that the LSI for a period i is a ratio, with the average of the seasonality indices for the three periods before i (i-1, i-2, i-3) in the denominator, and the average of the seasonality index for period i and the indices of the periods preceding and following period i (i-1 and i+1).

Let's take a concrete example. Suppose we're at the beginning of period #4 (July and August) at a health facility, and we want our LSI to help predict demand for that period. At the beginning of period #4 we have the seasonality indices available for the last three periods, 1 (January and February), 2 (March and April), and 3 (May and June). So, our LSI for period 4 has the average of seasonality indices for the periods #1, #2, and #3 in the denominator. It also has the seasonality index for period #4, and the two periods on either side of period #4, periods #3 (May and June) and #5 (September and October). Periods 3 and 5 are in the numerator because we want to add some correction, in case seasonality has shifted a bit early or late. The three indices in the numerator will allow us to calculate the average of recent consumption.[1] Figure 5 and table 2 show the LSIs that we would generate from our seasonality indices (calculated in table 1), and the following example also shows how to calculate the LSI for one period, July and August (period #4):

$$LSI_{April=}\frac{Ave\ (S_{i-1},\ S_i,\ S_{i+1})}{Ave\ (S_{i-3},\ S_{i-2}, S_{i-1})} = \frac{Ave\ (3.03 + 3.29 + 2.63)}{Ave\ (1 + 1.73 +\ 3.03)} = \frac{2.98}{1.92} = 1.55$$

[1] The mathematically inclined will recognize that since the LSI is a ratio of seasonality indices, it simultaneously de-seasonalizes and reseasonalizes the most recent consumption data in order to create the look-ahead forecast.

Table 2. Seasonality Index (SI) and LSI for Our Primary Health Facility

Month	SI	LSI
Jan–Feb	1.00	0.54
Mar–Apr	1.73	1.24
May–Jun	3.03	2.16
Jul–Aug	3.29	1.55
Sept–Oct	2.63	0.86
Nov–Dec	1.00	0.52

Figure 5. Seasonality Index and LSI for Our Primary Health Facility

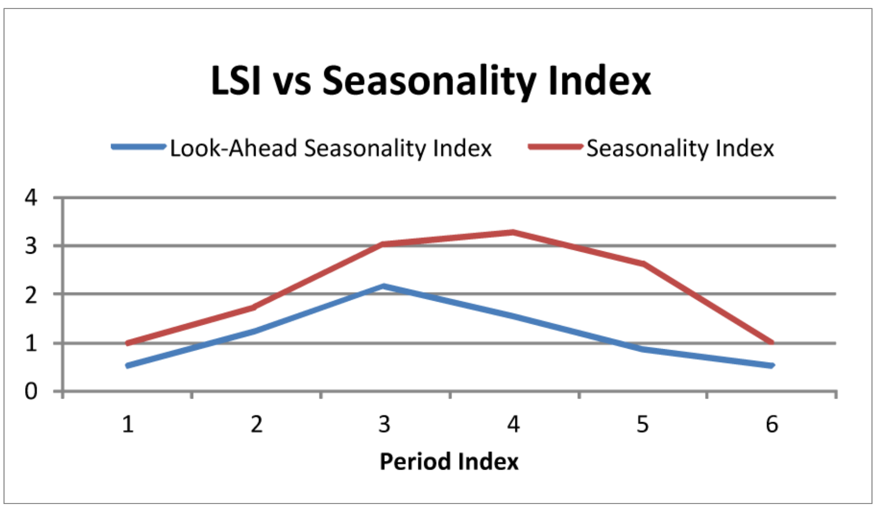

Figure 6 shows the performance of our LSI approach, instead of using history to look ahead, compared to the simple AMC rule. The LSI approach improves on avoiding stockouts compared to using historical consumption to look ahead that we used earlier (see figure 4) and not only lowers inventory in the low season, but inventory tends to ramp up toward the high season. Using the simple AMC rule, inventory tends to trend down going into the peak season, as a result of too much inventory being introduced at the end of the peak season.

Figure 6. Comparing AMC Rule to LSI Approach

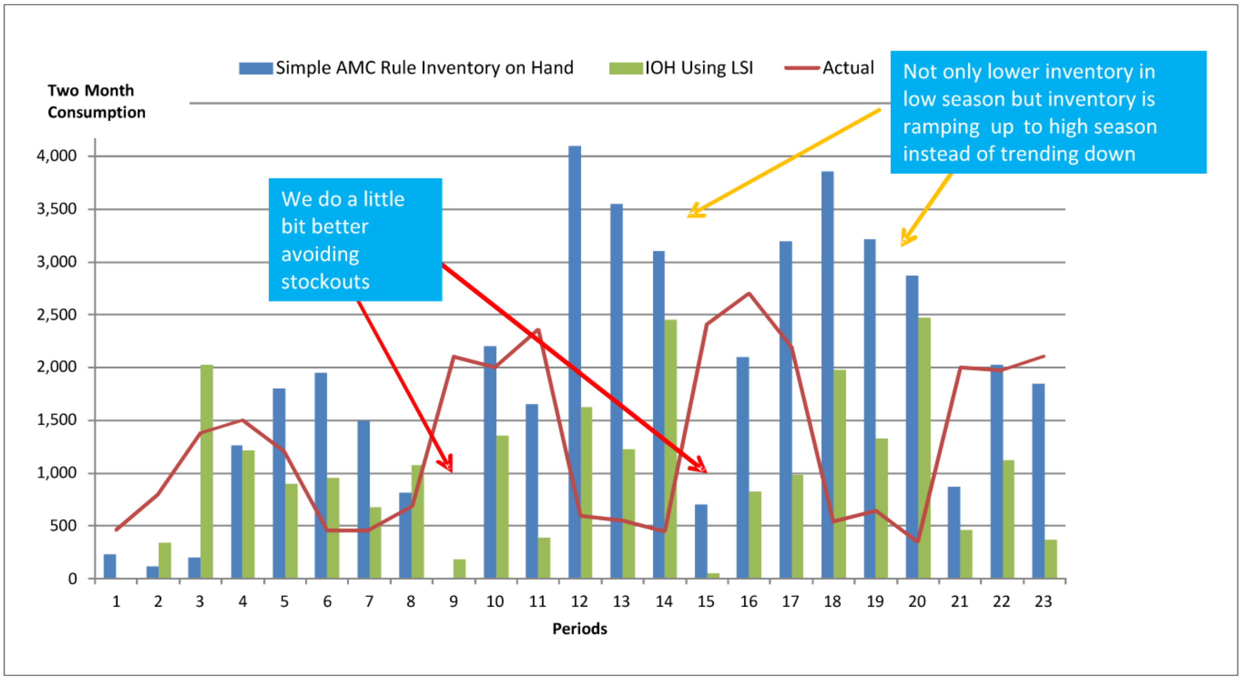

Discussion

At the end of the introduction, after introducing the idea of using history to look ahead and forecast, the following issues were raised for evaluating a general approach for addressing seasonality:

- incorporating changes in level compared to history
- accounting for lost consumption due to stockouts
- compensating for shifts in timing of seasonality
- using with poor or no historical data
- using the approach at higher tiers in the supply chain
- planning for facilities that regularly get cut off from routine inventory replenishment
- operationalizing the approach within existing healthcare systems in developing countries.

We examine in more detail in the following paragraphs the LSI approach with respect to each of these issues.

Incorporating Changes in Level Compared to History

As mentioned earlier, the LSI approach has some mechanisms for incorporating changes in the level of consumption, since it uses the most recent consumption data in making its forecasts. Being a ratio, the LSI preserves the level of the recent consumption data and only adjusts it for seasonality.

The LSI approach, however, does have a weakness if forecasts need to look too far ahead in the future (e.g., one year). The further ahead that the LSI is used to look ahead, the greater the chance that level of the period may have changed compared to that of the most recent consumption. In such situations, the forecast using the LSI approach can be augmented by a separate trend parameter. For example, assuming an annual growth rate of 3 percent, forecasts that look ahead one year can be increased by 3 percent to account for the growth rate.

Accounting for Lost Consumption Due to Stockouts

Lost consumption as a result of stockouts refers to consumption that would have occurred if there was inventory available to meet all of the consumption. When there is lost consumption, this tends to underestimate the true AMC if it is not tracked in some fashion. For example, the number of days out of stock may be recorded, or the number of patients recorded that would have received the health commodity in the event of a stockout. Lost consumption is a distinct reality for healthcare systems in developing countries and, therefore, it is important to understand how robust the LSI approach can be to the impact of lost consumption on forecasting.

Figure 7 shows the stock on hand and the average percentage of consumption that is lost at the facility as a result of stockouts for three scenarios: using the simple AMC rule (which uses the previous period's consumption as a forecast) and not tracking consumption; using LSI and tracking consumption; and using LSI but not tracking consumption.[2] Here, tracking consumption means recording the consumption that would have occurred if there had not been some stockout (e.g., recording days out of stock). Also, stock on hand is the inventory at the facility after meeting demand.

Generally, the lower the curve on this graph, the better performing the scenario, since it implies that the scenario is able to meet the same level of consumption, but with lower levels of inventory in the entire supply chain. Here, using LSI and tracking lost consumption is expected to outperform the other two methods, since tracking lost consumption reduces the underestimation of the true AMC, which should improve forecasting.

In figure 7, we see that the performance of using LSI without tracking lost consumption is similar to that of using LSI with the tracking of lost consumption. The similarity in performance here of both LSI approaches can be at least partially attributed to the use of averages of consumption over multiple periods, as the more periods that are used to calculate the AMC, the less the under-estimation effect of stockouts.

Compensating for Shifts in Timing of Seasonality

As mentioned earlier, the formula for LSI includes additional seasonality indices in the numerator (those preceding and following the periods that are being forecast), as a means of providing some correction for shifts in the timing of seasonality. Incorporating the preceding period incorporates

[2] These curves were generated by varying the maximum stock level multiplier for all three scenarios.

the potential that the seasonality is delayed in its arrival, while incorporating the following period incorporates the potential that the timing of seasonality is early.

Using LSI without Historical Consumption Data

If there is insufficient historical consumption data to generate accurate seasonality indices and LSIs, LSIs can still be generated from a rough sense of seasonality and used until sufficient consumption data are generated. For some seasonal commodities like malaria, it is usually possible to identify which periods in the season are generally considered low season and which are considered peak. In addition, it is often possible to estimate how much higher the consumption in the peak season is compared to the low season. For example, facility workers can give a rough estimate of how many patients they see per week in the rainy season as compared to the low season. This information can be enough to generate a rough seasonality index from which LSIs can be generated.

Figure 7. Comparing Inventory Performance Given Stockouts and Lost Consumption

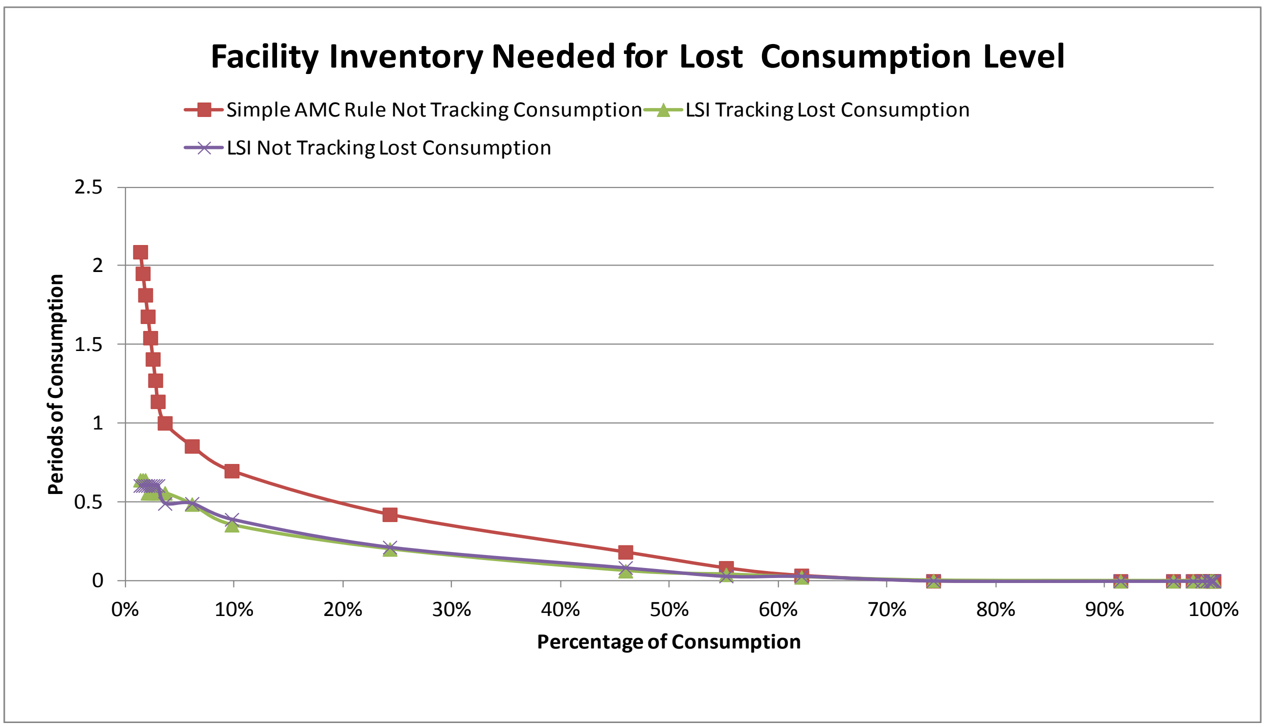

Table 3. Crude Estimates for Seasonality Index and LSI

Original SI	Original LSI	Estimated SI	Estimated LSI
1	0.538681	1	0.666667
1.730263	1.244076	1	2.057601
3.026316	2.159506	2.5	1.522136
3.289474	1.554286	2.5	0.808178
2.631579	0.859635	2.5	0.622532
0.995614	0.517157	1	0.41676

Consider our example of seasonality from table 1. Such a seasonality pattern could be approximated by assuming that consumption in the peak season is 2.5 times that of consumption in the low season and that the peak season lasts for half of the year. Such a seasonal pattern can be represented by setting the seasonality index in the low season to 1 and the seasonality index in the peak season to 2.5, as in table 2.

From this crude seasonality index, estimated LSIs can be calculated in the usual way, using the given LSI formula. Although this is a gross estimate of the true LSI, it allows an inventory management process that is built around the LSIs to operate and thus exhibit some sensitivity to seasonality, without having to wait for the year or two to generate consumption data needed for more accurate estimates. To compensate for the crudeness of the LSI estimates, it would also be appropriate to increase the maximum stock level multiplier slightly, so as to provide additional buffer inventory to cover errors in forecasting.

For example, if the multiplier is 2, it could be raised to 2.2 or 2.5. The increase in the maximum stock level would be subjective. There is no rule of thumb for how much you would raise the max, although it would be based on given country context and the feasibility of increasing maximum stock levels and by how much. Again, this subjective increase would only be applied when LSI estimates are crude. In the absence of historical data, an increase in the maximum stock level multiplier would also be appropriate, because in general, it would drive down stockouts, improve the quality of the consumption data that is being calculated, and improve the estimates of seasonality that would be inferred.

Use of LSIs by Higher Tiers in the Supply Chain

As described earlier, the LSIs are considered supply chain tier specific: They should be adjusted to match the lead times for inventory to be replenished to that tier of the supply chain and the time between the review of inventory. Again, because of its similarity to the simple AMC rule, the LSI approach is ideally intended to provide, at all levels of the supply chain, a practical and easy-to-use means of calculating resupply for seasonal commodities.

Earlier we gave an example formula for the LSI for facilities, assuming that facilities review inventory every two months and that the lead time for receiving inventory was fairly small compared to that review period of two months. Consider another example of a warehouse that has a lead time of four months and reviews its inventory every two months. The formula for its LSIs is given in Figure 8.

As mentioned earlier, review period and lead time are different concepts that affect the periods that you are trying to forecast. If supply chain tiers differ in either of these, then the LSI should reflect that. Here, the numerator is the average of the seasonality index of the period that is four months away (two, two-month periods), and the index of the period that precedes and that follows. Appendix B also provides a general LSI formula for any tier in the supply chain with a particular lead time and length of the review period.

Upstream tiers in the supply chain can use their LSIs on consumption data from downstream tiers like health facilities (if this information is shared with the upstream tier), or can use the LSI on their own issues. As discussed in appendix C, the most benefit to the supply chain comes from all tiers in the supply chain using LSI. However, even if upstream tiers do not use the LSI for their own issues but only on downstream tiers, there is still benefit, as the LSI makes orders from the downstream tiers more predictable and easier to forecast. Interestingly, the benefits from sharing consumption data from downstream tiers with upstream tiers only comes when all the tiers use a sophisticated approach like LSI and only for commodities that are generally in good supply (with consumption data available).

Figure 8. Formula for Supply Chain Tier with Lead Time of Four Months and Review Period of Two Months

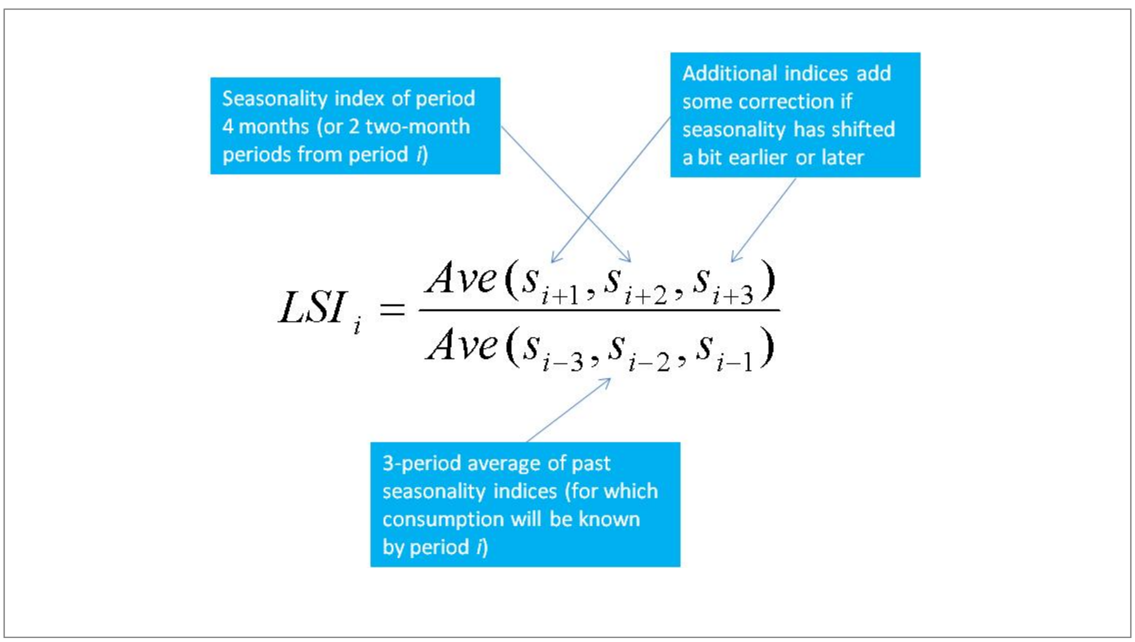

$$LSI_i = \frac{Ave(s_{i+1}, s_{i+2}, s_{i+3})}{Ave(s_{i-3}, s_{i-2}, s_{i-1})}$$

Planning for Facilities that Get Cut Off from Routine Inventory Replenishment

The current approach for dealing with facilities that are regularly cut off from routine inventory replenishment is to use the AMC and a higher maximum stock level (compared to facilities that do not get cut off), prior to the time that such facilities are usually cut off. The higher maximum stock level is intended to allow inventory to cover the period over which the facility is expected to be cut off. However, if the period that a facility is cut off coincides with the peak consumption season, and since the AMC will be based on the low season, then it is possible that the higher maximum stock level still results in insufficient inventory for the facility to support it during the period that it is cut off.

Figure 9. Formula for LSI for Facility Cut Off for Two Periods

Two periods over which the facility is expected to be cut off

Additional indices add some correction if seasonality has shifted a bit earlier or later

$$LSI_i = \frac{Ave(s_{i-1}, s_i, s_{i+1}, s_{i+2})}{Ave(s_{i-3}, s_{i-2}, s_{i-1})}$$

Three-period average of past seasonality indices (for which consumption will be known by period *i*)

Consider an example where the monthly consumption at a facility during a six-month peak season is 300 units, and monthly consumption in the low season is 100 units. The facility could be cut off for all six months of the peak season. The max level required at the beginning of the peak season to cover twice the consumption in the peak season over the six months is 3,600 units, implying a maximum stock level of 36 months of stock (MOS) on the monthly consumption in the low season. Current approaches would focus on the monthly consumption in low season, and choose a maximum stock level closer to 12 rather than 36 on the monthly consumption low season, resulting in lower inventory than appropriate for the cut-off facility.

To address this, LSIs can be generated for cut-off facilities for the period right before the facility is to be cut off. These LSIs would determine the top-up level for these facilities. For example, consider a facility that is cut off for four months (that is, two two-month periods) beginning at a certain period. The LSI for that period can be recalculated using the formula given in figure 9.

Here the formula is similar to that from figure 8, except that in the numerator, the seasonality indices that are averaged cover the two periods that the facility is cut-off, periods *i* and *i+1,* and the periods before and after the cut-off period. Recall that these additional periods add some correction, in case seasonality has shifted a bit earlier or later.

If the maximum stock level quantity target is two times the consumption over the cut-off period, then the maximum stock level multiplier would be four, since the facility is cut off for a total of two two-month periods.

The advantage of the LSI is that it adjusts for differences in consumption over the cut-off period and consumption when the inventory decision is made about what to send to the cut-off facility. As such, the maximum stock level only needs to take into account the length of the cut-off period. The max for cut-off facilities would be different compared to non–cut-off facilities, even without LSI. The problem is when the forecast in consumption for the cut-off facility doesn't take into account peak season levels but rather is based on consumption in the low season. The LSI should correct for this mistake and ensure that old max levels for cut-off facilities work better.

Operationalizing within Existing Healthcare Systems

Fit with Information Systems

Existing healthcare systems in developing countries typically are a hybrid of information systems, with primarily manual systems in place in downstream facilities, like health centers, and electronic systems becoming more prominent in upstream facilities, like central medical stores. The popularity of the simple AMC rule is at least partially due to its ease of use in manual systems, with inventory control decentralized and pushed down to downstream facilities. The LSI approach has the potential to still be used in manual systems due to its simplicity.

For example, special forms can be created with preprinted LSI information that can be used by health workers to calculate their top-up levels for each replenishment period after calculating the AMC. Orders are the difference between maximum stock levels and the inventory at the facility. So each facility would need to calculate AMC, then multiply it by LSI and the maximum stock level multiplier to get the maximum level. Then inventory would be subtracted from the maximum stock level, for the quantity that should be ordered.

Electronic information systems could easily adopt the use of LSIs given their simplicity. These systems could also allow for more sophistication in their use, such as having different sets of LSIs for different regions or types of facilities (e.g., primary versus tertiary) if the historical consumption data showed that such differences were significant. Such sophistication would be difficult to administer in more manual systems, where one set of LSIs for all facilities may sacrifice accuracy to allow for ease of implementation and ongoing administration.

Use across Products

Across products within the healthcare system for developing countries, multiple products may share the same seasonal pattern. For example, it would be reasonable to expect that most malaria-related products would share the same seasonal pattern. In the case of products that do not show a seasonal pattern, setting the LSI for those products across all periods equal to 1 is the same as using the simple AMC rule when it is more appropriate. Thus, the LSI approach can be incorporated into managing all product, whether they have a seasonal pattern or not (e.g., family planning products, which do not exhibit a seasonal pattern). Economies can be gained by using a single set of LSIs across groups of products, where appropriate.

On a side note, in some countries, the rainy season can result in an overall increase in patient volumes at health facilities, which generally drives an increase for health commodities (e.g., antibiotics for upper respiratory tract infections). The increase in volume is probably nowhere near that for commodities like malaria, and this seasonal variation may be covered sufficiently by the maximum stock level used in the simple AMC rule. However, if it becomes needed, the LSI approach can incorporate this milder seasonality profile through development of an additional LSI index, reflecting the milder seasonality profile.

Summary

The LSI approach is presented here as a mechanism for enhancing the simple AMC rule to handle seasonal commodities, while maintaining some of its simplicity to continue to meet complex needs. The approach operationally involves multiplying the AMC by indices that compensate for seasonality, before multiplying by the maximum months of stock level. The LSI approach makes use of a standard approach in forecasting so-called seasonality indices from historical consumption data.

Table 4 captures the list of challenges for an approach used to address seasonality and highlights the features of LSI that help to address those challenges. We discussed the ability to address changes in level of consumption from history to the present because of the use of ratios and the fact that we use recent consumption for forecasting. With respect to lost consumption, simulation results suggest that LSIs are reasonably robust compared to the underestimate of consumption, which occurs as a result of stockouts, partly because we use averages in our formula. That is, since stockouts lead to low levels of consumption, which when used in forecasts tend to result in forecasts that are biased low, the LSI is considered robust against these effects. This is because the bias for low forecasts in the face of stockouts is not as high as other forecasting methods; this robustness is partly the result of the use of averages of multiple periods of consumption in our formula.

Also, if lost consumption is tracked—for example, using days out of stock or tracking patients that do not receive commodities—then the implied lost consumption can be easily incorporated and used with the LSI. The LSI has some correction for potential shifts in seasonality due to the incorporation of additional seasonality indices beyond those periods that are being forecast.

In situations that lack historical demand, an inventory system can be initialized with a rough idea of seasonality (e.g., assuming that consumption during the peak season is two to three times that of the off-peak season and that the peak season is half of the year). LSIs can be built based on that rough idea and updated when historical consumption data are available.

LSIs are linked to procurement lead times and review periods of the supply chain tier where they are used, and thus are tier specific. For upstream tiers, they can be used on consumption data from downstream tiers or on issues. The LSI approach can also be used to plan for cut-off facilities, because LSIs can be created that look ahead for the period(s) that the facility will be cut off. Finally, the similarity between the AMC and LSI formulas for top-up maximum stock level targets between the simple AMC rule and the LSI approach lends itself to operationalizing both within existing healthcare systems and systems that are shifting to increased use of electronic information systems.

Table 4. Challenges for Addressing Seasonality and Features of the LSI Approach

Challenges for Addressing Seasonality	LSI Features
Are there level changes, especially when consumption is up compared to previous year?	LSI based on ratio, which accommodates level changes
What happens with lost consumption due to stockouts?	LSI is robust to lost consumption through use of averages
What if there are seasonality shifts (e.g., due to rainy season starting early or late)?	LSI has some correction for shift in seasonality
Can it be used if we have no history?	Rough seasonality index can be used until enough consumption data is captured
What about higher tiers in the supply chain? Would they use the same rule?	LSI is tier specific by being based on procurement lead time and length of review period
Can we use it to help plan for cut-off facilities?	LSIs can "look ahead" for the period that the facility will be cut off to aid planning
How do we operationalize within healthcare settings in developing countries?	Similar formula for top-up level target as simple AMC rule and can be used with commodities with no or moderate seasonality

Zambia Case Study

Whereas the previous section focused on the concepts behind the LSIs, this section examines the Zambian context for the potential utility of the LSI, as well as provides simulated results of the use of the LSI approach on historical consumption data.

Zambia

Zambia is a landlocked country of about 13.5 million people in Southern Africa (Central Intelligence Agency 2011). Per capita annual income of approximately U.S.$1,700 (2012) places the country among the world's poorest nations: Zambia was ranked 150th on the Human Development Index in 2010 (United Nations Development Programme n.d.). Zambia is also one of the most highly urbanized countries in sub-Saharan Africa, with 44 percent of the population concentrated in a few urban areas along the major transport corridors, while rural areas are sparsely populated. Rural poverty rates stand at about 78 percent and urban poverty rates at 53 percent (United Nations 2010).

Malaria is the primary public health challenge, with an estimated three million cases in 2008–2009 (World Health Organization 2010). The public sector is the largest provider of healthcare, followed by the Churches Health Association of Zambia (CHAZ) member institutions and the mine hospitals. Public sector facilities are located throughout the country, while mission facilities are concentrated in the rural areas. A small number of private facilities (hospitals, clinics, pharmacies, etc.) also provide some healthcare, but predominantly in the urban and mining areas. There are some small and informal outlets—*Kantembas*—which operate in remote areas and high-density, urban shanty compounds. A lack of human resources also impacts the delivery of services in Zambia, with medical staff leaving the country or choosing practice in the private sector. The impact of HIV and AIDS on health workers is a challenge as well.

The climate of Zambia is tropical modified by elevation. There are two main seasons: the rainy season (November to April) and the dry season (May/June to October/November). In the rainy season, the mosquito population generally rises, since mosquitoes lay eggs in standing water, and there is an expected rise in cases of malaria and the need for antimalarial treatments. Simultaneous with the increased need for malaria treatments, there may also be greater transportation difficulties in getting to remote areas (e.g., due to flooding).

Creating a Seasonality Index for Zambia

Figure 10 shows the geographic distribution of health facilities in the Essential Medicines (including malaria) Logistics Improvement Program (EMLIP) districts of Zambia; the Project serves and has access to data only in the EMLIP districts. The green dots represent the facilities with 18 or more months of data, all located in the original 16 EMLIP districts, which were used to create a seasonality estimate for Zambia. We refer to these facilities as "seasonal" facilities. Orange dots represent facilities with sufficient data, but their consumption profile seemed initially to not fit the seasonality pattern of the seasonal facilities, and we refer to these as "aseasonal." Blue facilities were facilities with sufficient data, but were neither convincingly seasonal or aseasonal. Smaller purple circles were facilities with limited data, while gray small circles are facilities not in EMLIP focus districts. The prime takeaway from this map is that the seasonal and seemingly aseasonal facilities were distributed broadly across the country, with most regions holding both aseasonal and seasonal facilities. This kind of dispersion lends support to use of a single seasonality profile for the facilities in Zambia.

Figure 10. Geographic Distribution of Health Facilities in Zambia: Facility Distribution of Seasonal and Aseasonal Facilities in Zambia from October 2010 to October 2012.

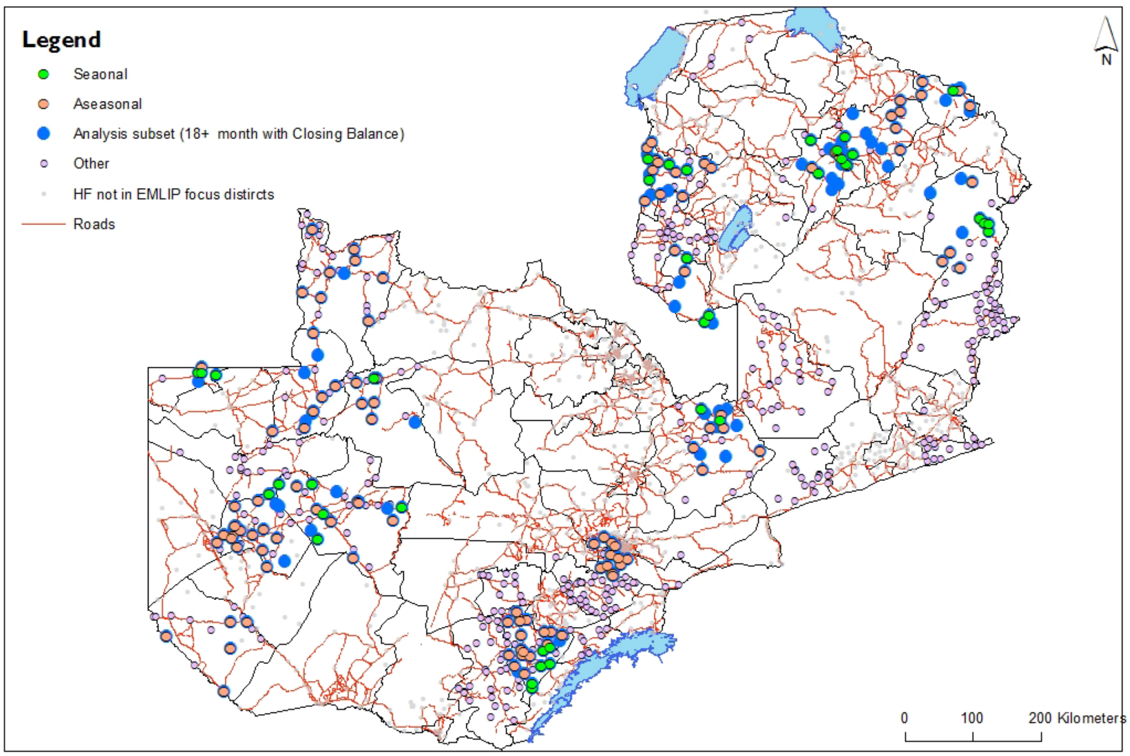

Figure 11 specifically describes the steps taken to create the seasonality profile for Zambia. We used the seasonal facilities identified earlier that had sufficient stock to avoid the effect of stockouts and close to two years of data. Because of the limited availability of data, we converted the consumption of the different presentations of artemether/lumefantrine (AL; 1x, 2x, 3x, 4x six-tablet pack) to the consumption of tablets. This conversion to the consumption of tablets overall is particularly relevant given that presentations of AL may be cut down or combined to provide the appropriate dose for a patient in stockout situations.

Figure 11. Steps for Creating a Seasonality Index in Zambia

1. Use "seasonal" facilities with good inventory presence and close to two years of data.
2. Convert consumption of different presentations of AL to consumption of tablets.
3. Generate two month consumption totals over two years.
4. Take the average of (1) consumption in a one two month period and (2) consumption one year later.

We then generated two-month consumption totals over two years and took the average of consumption in one two-month period and consumption in the same two-month period, but one year later. This averaging approach is particularly appropriate if there is a change in level across years and that change is better described as a percentage increase rather than as a constant increase in magnitude across all the months (i.e., multiplicative instead of additive).

Figure 12 shows the resulting seasonality index in blue and the LSI for forecasting at the facility level in red. The seasonality profiles reveal the low point of consumption in the August to September time period and the highest consumption in the four months of February through May. This corresponds with a general rainy season that stretches from November to April, assuming that malaria persists for some time even after the rainy season ends. In particular, the seasonality index shows that the ratio of consumption at its peak to consumption at its lowest is 2.6. Even this piece of knowledge can be helpful in inventory and resource planning, as consumption in August and September can provide four to six months of advance notice about needs for malaria the following year.

Figure 12. Seasonality Index and LSI for Zambia Seasonal Health Facilities

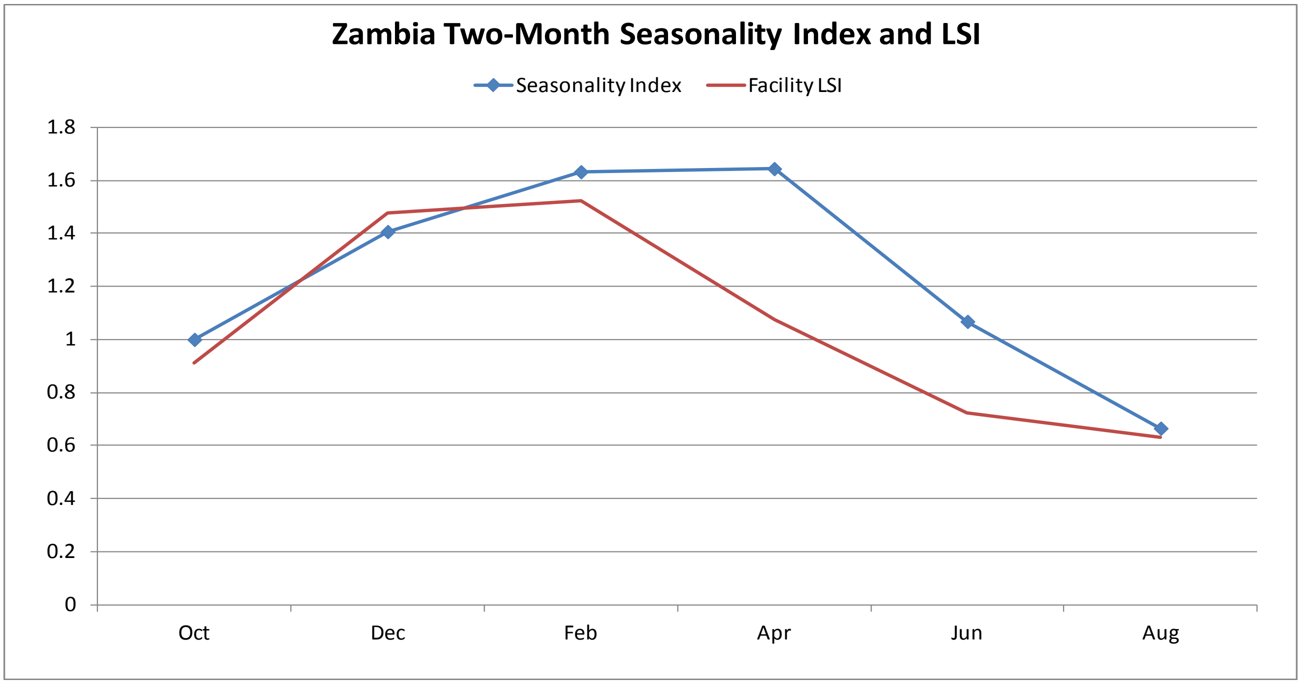

Testing LSI on Zambia Consumption Data

To test the LSI approach on Zambia data, additional forecasting models were considered for comparison. The models are described as follows.[3] In particular, these models were compared using data both from seasonal facilities and aseasonal facilities. Note that these are all forecasting methods, not reordering rules, as they predict consumption in future periods. Based on the predictions of these methods, inventory replenishment procedures can be adjusted and maximum stock levels determined based on established max/min levels.

The first forecasting method is the simple AMC rule, which takes the average of previous two periods of consumption as the forecast for the next period.[4] The second forecasting method is the simple exponential smoothing (ExS). The formula for exponential smoothing is: $F_t = \alpha x_t + (1-\alpha)F_{t-1}$, where F_t is the forecast for period t+1; x_t is the consumption in period *t*; and α is a constant that is optimized, where the alpha is varied to reduce forecasting error and to provide the best forecasting outcome. The approach involves smoothing the previous forecast by taking a weighted average of it and the most recent consumption data. This is a very famous forecasting approach, popular because of its simplicity. In creating the forecast, the value of α was optimized to give the best result on each of the datasets—seasonal facilities and aseasonal.

---

[3] Although Gallien et al. (2012) also address seasonality issues in the Zambian context, their approach does not make use of an explicit forecasting methodology that could be used for comparison here, as described in the appendix.

[4] One of the rules of thumb for evaluating a forecasting approach when an external benchmark is not available or appropriate, is to compare it to forecasting using the most recent consumption. The simple rule here is similar, but a bit more robust since it uses two periods rather than one.

The third approach is triple exponential smoothing (3xExS), which is similar to the exponential smoothing approach but involves additional parameters to incorporate trend and seasonality (see appendix A). The approach allows parameters to vary by facility, which are updated after every period of consumption. As with the exponential smoothing, the values of user-specified parameters within the forecasting approach were partially optimized to give the best result on each of the datasets.[5] Finally, since the LSI approach uses the same seasonality factor across all facilities, the last model is based on the triple exponential smoothing, but here the seasonality index was artificially fixed to be the same across all facilities (3xExSF).

Evaluating Models Using Mean Average Percentage Error (MAPE)

The indicator used here to evaluate the performance of the different forecast approaches, including the LSI, is referred to as the mean absolute percentage error (MAPE). The formula for MAPE is the following:

$$MAPE = Ave\left(\frac{abs(Forecast - Actual)}{Actual}\right)$$

The indicator involves taking the absolute error as a percentage of the actual, where the absolute error is the difference between the forecast and the actual, ignoring the sign of the error. Finally, the average or mean of all the percentage errors is taken. It follows then that the MAPE can range from zero to infinity, with a lower MAPE score indicating better forecasting performance. MAPE, like other performance indicators, can be affected by many things, including the fit of the forecasting methodology to the data, the inherent complexity of the forecasting task, (e.g., if seasonality or noise in the data are present), or if forecasting a new product compare to an existing/old product. The MAPE is also affected by level of aggregation; so for example, everything else being equal, the MAPE should be lower for national forecasts than for facility forecasts, and the same for annual versus weekly forecasts or family product versus individual product forecasts.

As a result of the variation that we can find in MAPE scores due to all of these factors, it is difficult to have an external benchmark for forecasting performance, so there are some rules for alternatives to the external benchmark. For example, 10 percent to 50 percent can be considered good, depending of course on the context. In lieu of an external benchmark, some recommendations are to compare the MAPE of the forecasting methodology to the simple forecast, which uses the consumption in the previous period, or to compare the forecast methodology to itself but, over time, expecting to see an increase in effectiveness as learning occurs.

Figure 13 shows how the forecasting methods compare based on the MAPE for 102 aseasonal facilities. Since the MAPE can be driven by outliers arising from consumption actuals that are very small, the x-axis captures a cut-off threshold that results in ignoring the forecast error for any consumption levels that are lower than the threshold. For example, at a cut-off consumption level

---

[5] The parameters for the triple exponential smoothing were only partially optimized as there was not enough historical data to optimize all of them, in particular, the parameters for smoothing the seasonality indices.

threshold of 200, forecasts are not calculated and, therefore, are not included in the MAPE score for all consumption levels below 200.

In forecasting aseasonal facilities, the LSI approach uses the seasonality index that was derived from the seasonal facilities. Our first observation is that the relative performance of the forecasting methods is roughly the same across the different cut-off thresholds. For aseasonal facilities, the triple exponential smoothing has the lowest error, while LSI is similar to simple exponential smoothing and triple exponential smoothing with fixed seasonality index across facilities. For reference, at a cutoff of 200, which is about 11.5 percent of the average two-month consumption, the MAPE score for triple exponential smoothing is 44 percent, while for LSI the MAPE score is about 56 percent. The worst performer is the simple rule, with 68 percent.

Figure 13. Forecasting Performance Comparison for Aseasonal Facilities

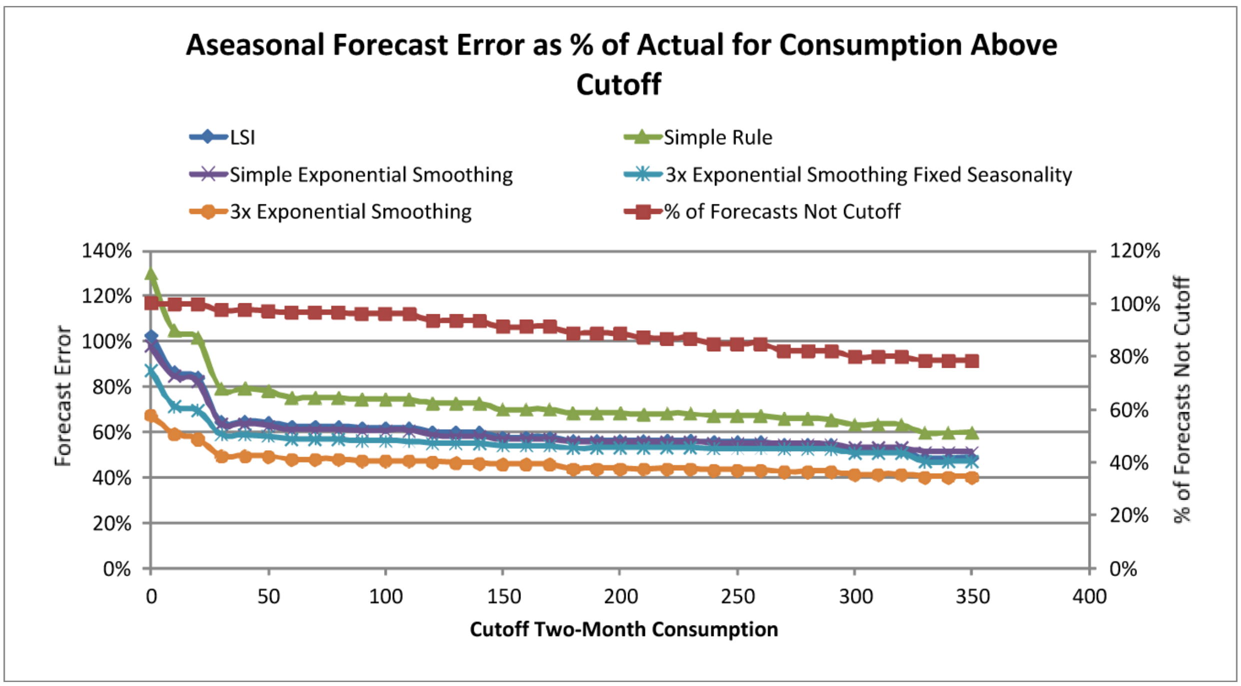

For the 41 seasonal facilities, we see similar results in figure 14, although the cut-off threshold has more of an effect on the relative ranking of the forecasting methods. For reference, at a cutoff of 300, which is about 11.8 percent of the average two-month consumption for the 41 seasonal facilities, the MAPE score for triple exponential smoothing is 32 percent; for triple exponential smoothing with fixed seasonality index, the MAPE score is 42 percent; while for LSI it is 55 percent. The worst performer again is the simple rule, at 85 percent.

Figure 14. Forecasting Performance Comparison for Seasonal Facilities

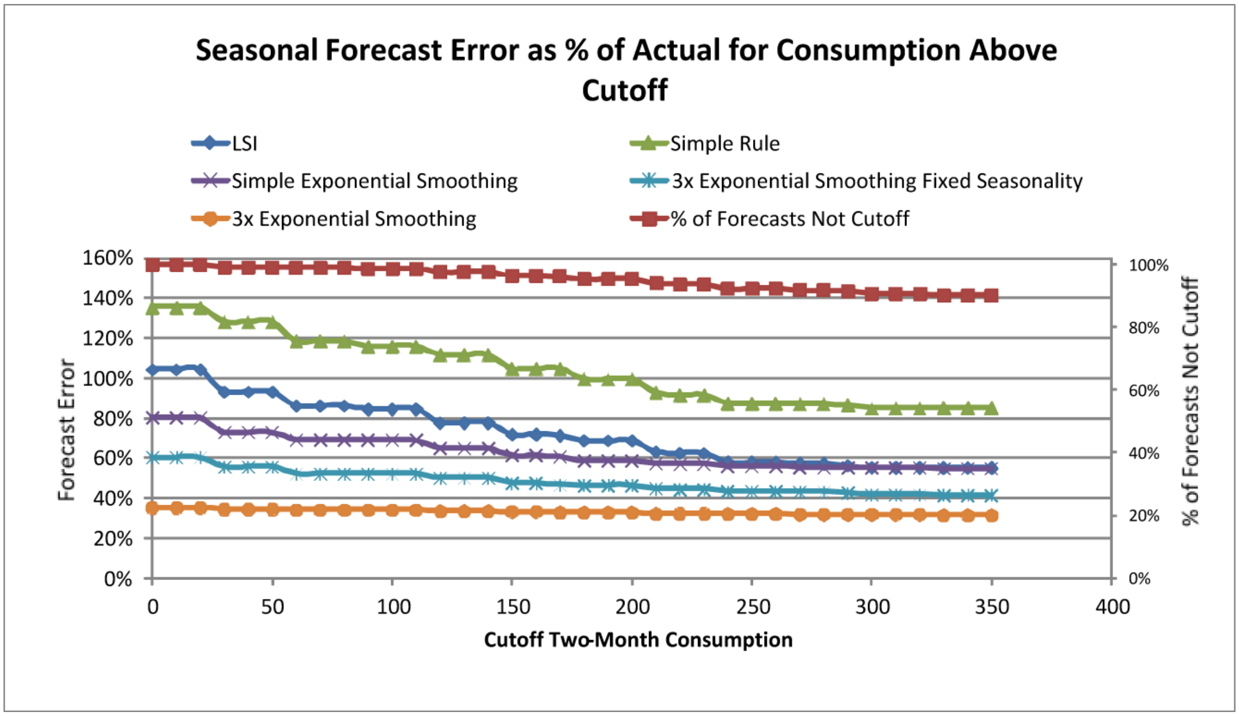

Evaluating Models Using Allocated Inventory Cost Methodology

In evaluating a forecasting approach for inventory replenishment, measures like the MAPE do not always translate into good inventory replenishment performance. For example, we can get the same MAPE score if a forecast methodology were consistently above the actual and if another were consistently below the actual by the same amount. However, for antimalarial commodities, we would prefer to have excess inventory rather than an equal number of units of missed/lost consumption. The timing of the forecasting error with respect to seasonality matters as well. The forecasting error during a peak season tends to bias the MAPE score lower than the forecast error during the low season, due to the higher levels of the actual consumption, but generally, the errors in the peak season are the ones that are more significant for health outcomes and inventory management response. Finally, decisions made in one period can have lasting effects on subsequent periods. For example, too much inventory in one period can carry over to later periods, and lost consumption in one period can affect forecasting in subsequent periods, especially if it is not tracked.

The alternative is a standard approach used to evaluate inventory replenishment approaches and involves simulating aggregate inventory costs by allocating a cost to excess inventory and a separate cost to lost consumption. In our approach, we normalize the cost of an excess unit of inventory by setting it to one. Then, we allow the cost of one unit of lost consumption to vary from zero (or less costly than an excess unit of inventory) to a large number (that is more costly than an excess unit of inventory). The latter is more the norm and, in the case of donated commodities, the relative cost of lost consumption to cost of excess inventory can be quite high, with very conservative estimates of this ratio generally above 10. We consider separately the cases where lost consumption is tracked versus not tracked. When lost consumption is tracked, although there are stockouts, we still know what consumption would have been and can use this information in our forecasting method. In the cases where the lost consumption is not tracked, we do not know what consumption would have been.

Figure 15 shows how for aseasonal facilities and tracking lost consumption, the inventory costs vary according to the cost of lost consumption, which can vary from zero to a large number and, in the case of the graph, to 20. The lower the inventory costs, the better the performance of the forecasting method. Generally, costs should rise as cost of consumption increases; however, some methods may have a steeper slope if they generally have more stockouts than another method. What we see here is that the LSI performance tracks closely to the triple exponential smoothing and does so more than any other method, including the triple exponential smoothing with fixed seasonality index, for all facilities. For reference, at a lost consumption cost of 10, the LSI and triple exponential smoothing have a difference in cost of 2 percent, while the triple exponential smoothing with fixed seasonality index for all facilities is 16 percent higher. Another interesting thing to note here is the performance of the simple exponential smoothing. It performs disastrously, showing a tendency for stockouts as shown by the increase in costs as the cost of lost consumption increases. Yet, simple exponential smoothing still had a similar MAPE score to that of LSI and triple exponential smoothing with a fixed seasonality index for all facilities.

Figure 16 shows the performance for seasonal products when lost consumption is tracked. Again, the LSI approach tracks closest to the triple exponential smoothing, while the triple exponential smoothing does very well at seasonal facilities, given that its costs barely rise as the cost of lost consumption does. This results from the fact that, in this instance, the forecasting method primarily results in an oversupply of inventory, with few stockouts and lost consumption. Therefore, its costs do not rise much as the costs of lost consumption increase. For reference, at a lost consumption cost of 10, the LSI and triple exponential smoothing have a difference in cost of 12 percent while the triple exponential smoothing with fixed seasonality index for all facilities is 24 percent higher than the triple exponential smoothing.

Figure 15. Inventory Replenishment Performance Comparison for Aseasonal Facilities and Tracking of Lost Consumption

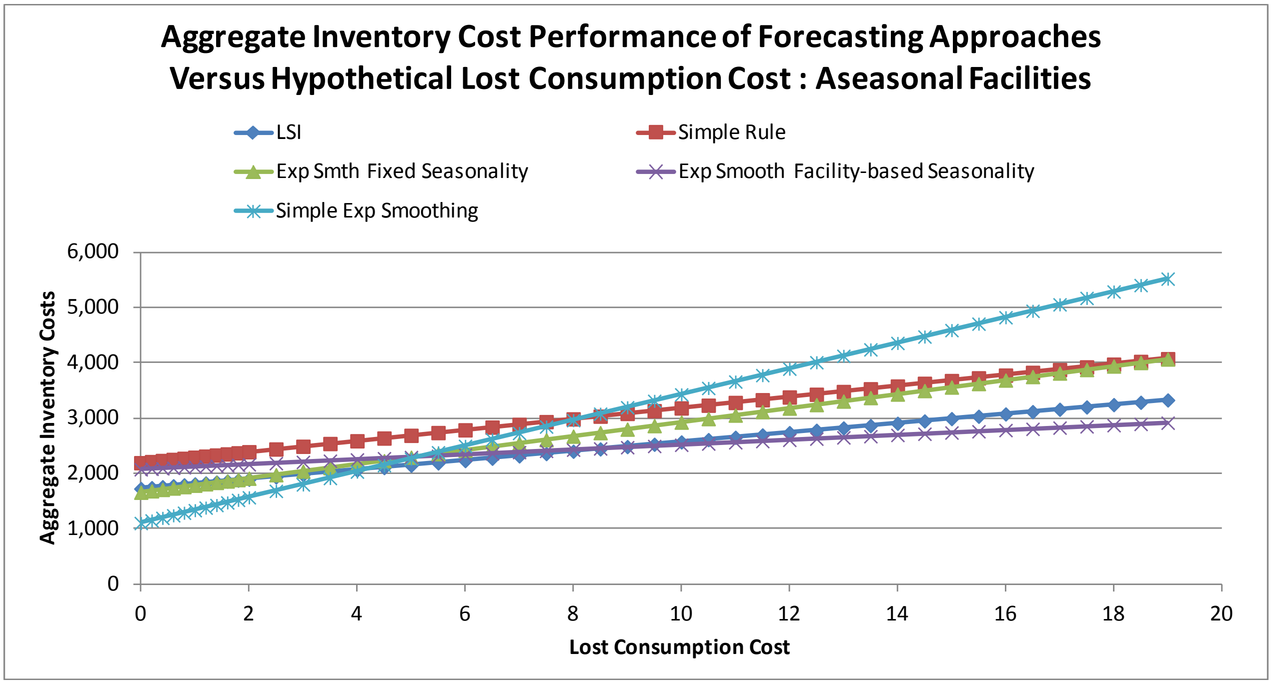

Figure 16. Inventory Replenishment Performance Comparison for Seasonal Facilities and Tracking of lost Consumption

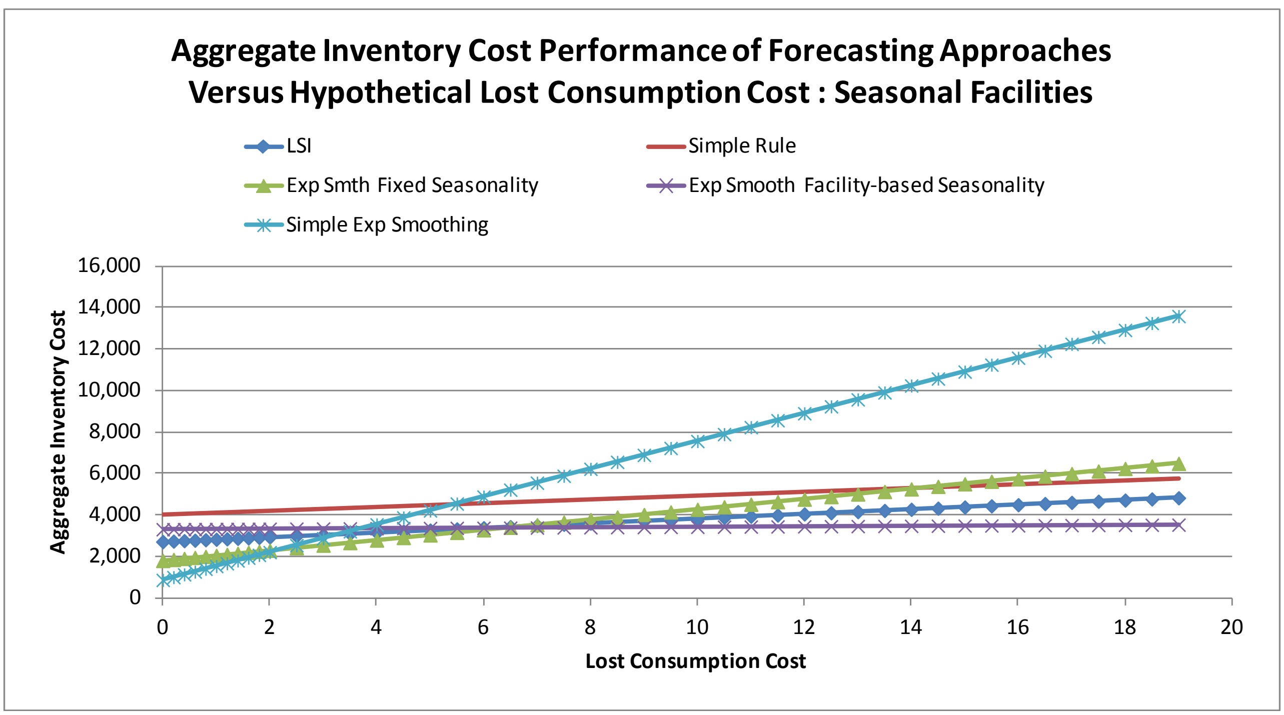

Figure 17 shows the performance of the forecasting method for seasonal when lost consumption is not tracked, and so the forecasting method does not have the information of what consumption would have been if there was no stockout. Again, LSI tracks the closest to triple exponential smoothing. For reference, at a lost consumption cost of 10, the LSI and triple exponential smoothing have a difference in cost of 15 percent, while the triple exponential smoothing with fixed seasonality index for all facilities is 44 percent higher than the triple exponential smoothing. The outperformance of the triple exponential smoothing might be exaggerated here. As mentioned earlier, when tracking lost consumption (see figure 16) the forecasting method tended to oversupply and suffer little stockouts. Therefore, it follows that it would be minimally affected if lost consumption was not tracked.

Figure 17. Inventory Replenishment Performance Comparison for Seasonal Facilities and No Tracking of Lost Consumption

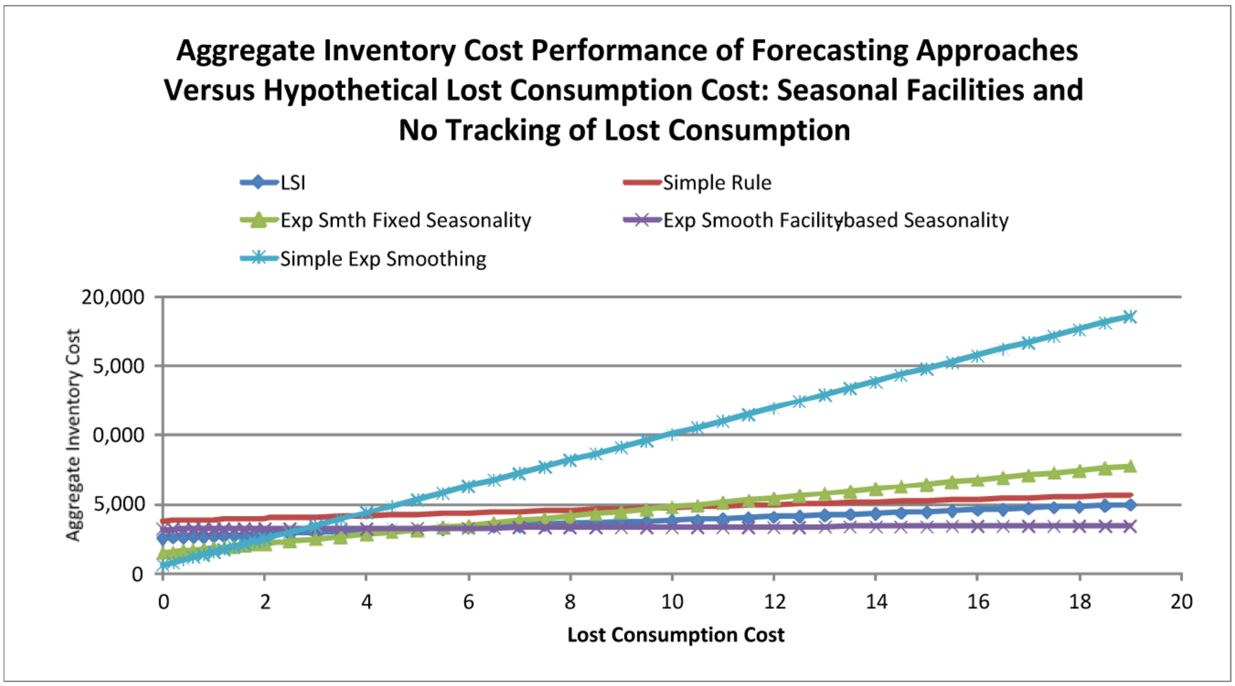

Examining the results on the aseasonal facilities, figure 18 shows different and very interesting results, especially since the triple exponential smoothing does suffer more significant stockouts, and therefore, can be affected more significantly if lost consumption is not tracked. Here the LSI is the one model that outperforms triple exponential smoothing and is followed in performance by none other than the simple AMC rule. Recall that the simple AMC rule had consistently the worst MAPE scores. For reference, at a lost consumption cost of 10, the LSI and simple rule have a difference in cost of just 16 percent, while the triple exponential smoothing is 24 percent higher than the LSI.

The results also are not terribly surprising because the exponential smoothing forecasting methods place a lot of weight on the most recent consumption data. If those data are poor, the forecasts are corrupted. The LSI and the simple rule use averages of consumption from multiple periods, which provides some robustness to the effect of lost consumption that is not tracked. Both methods, at a lost consumption cost of 10, are slightly worse when lost consumption is not tracked: 4 percent higher for LSI and 1 percent higher for the simple rule. However, the exponential methods have costs increasing by 64 percent or more when lost consumption is not tracked. In fact, the simple rule comes second in performance only because the other methods suffer so significantly.

Figure 18. Inventory Replenishment Performance Comparison for Aseasonal Facilities and No Tracking of Lost Consumption

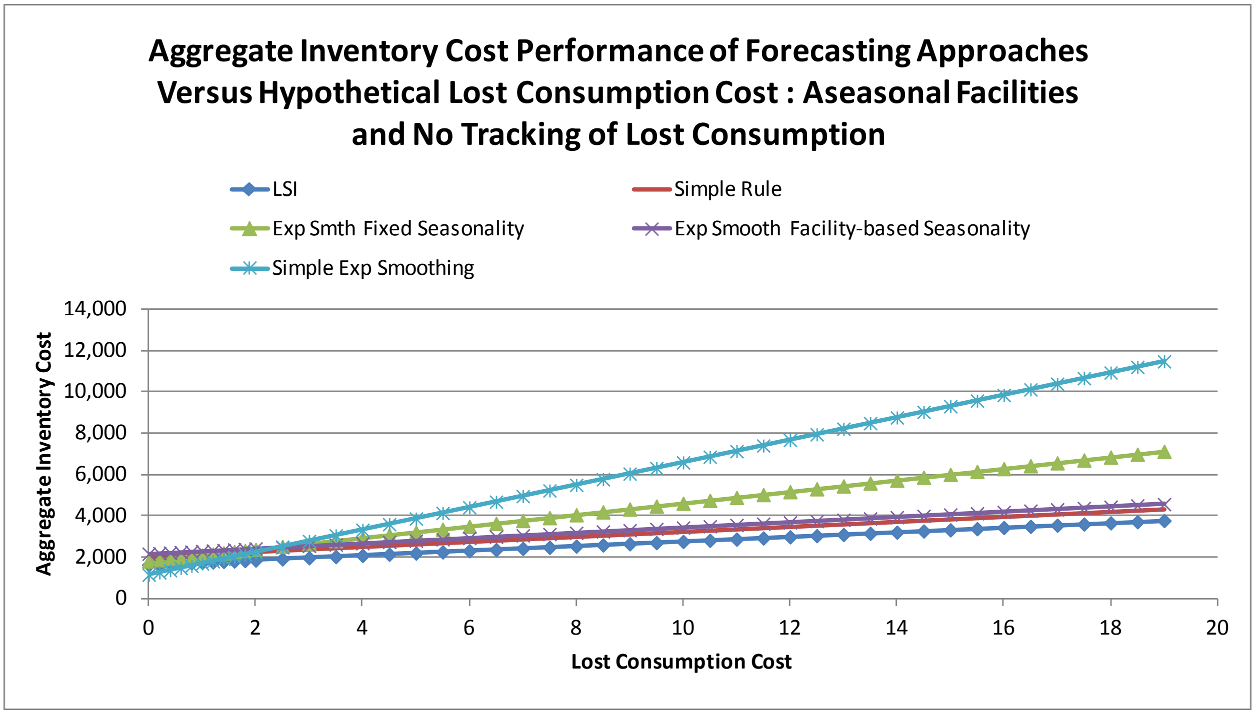

Tables 5 and 6 capture all of the results that were individually referenced for each comparison. Generally, LSI on the more important inventory cost method of performance evaluation is second only to the sophisticated 3xExS and has superior performance when lost consumption is not tracked. In addition, the MAPE scores for LSI are adequate.

Table 5. Summary of Forecasting and Inventory Replenishment Performance Comparison I

	MAPE (aseasonal)	MAPE (seasonal)	Inventory Costs (aseasonal, tracking)	Inventory Costs (seasonal, tracking)
3xExS	44%	32%	Lowest	Lowest
LSI	56%	55%	+2%	+12%
3xExSF	53%	42%	+16%	+24%
ExS	55%	55%	+36%	+185%
Simple AMC	68%	85%	+26%	+44%

Table 6. Summary of Forecasting and Inventory Replenishment Performance Comparison II

	Inventory Costs (aseasonal, not tracking)	Inventory Costs (seasonal, not tracking)
3xExS	+25%	Lowest
LSI	Lowest	+15%
3xExSF	+67%	+44%
ExS	+138%	+203
Simple	+16%	+44%

Zambia Analysis Summary

The availability of consumption data at health facilities with good inventory support for close to two years facilitated our analysis in this case study. A geographic analysis of so-called seasonal facilities revealed no discernible pattern to their distribution across Zambia, lending support for considering their seasonal pattern as being representative for Zambia as a whole. Tests on seemingly aseasonal facilities using this seasonal pattern seemed to confirm its suitability.

The Zambia consumption data allowed for a considerable test of the LSI approach. Given its simplicity, ease of use, and ease of explanation, the LSI performs admirably in the tests here. It was tested both as an inventory replenishment mechanism under conditions that correspond to healthcare settings in developing countries and as a forecasting method. Although its performance was not consistently superior to all of the models used for comparison, its performance for both seasonal and aseasonal facilities was still commendable given how closely it tracked the consistently superior model (i.e., triple exponential smoothing). The LSI even outperformed triple exponential smoothing when lost consumption was not tracked (a common phenomenon in many countries). In addition, some of the methods that the LSI was evaluated against were far more sophisticated and optimized for the data being forecast.

Zimbabwe Case Study

This section examines the Zimbabwe context for the potential utility of the LSI approach and also to provide simulated results of the use of the LSI approach on historical consumption data.

Zimbabwe

Zimbabwe is a country of approximately 13.8 million people in Southern Africa (Central Intelligence Agency 2013). With a gross domestic product (GDP) per capita of approximately U.S.$600 (2012), Zimbabwe was ranked 172nd out of 186 countries on the Human Development Index in 2012 (United Nations Development Programme 2013). Zimbabwe is a relatively urbanized country in sub-Saharan Africa, with 39 percent of the population concentrated in a couple urban areas.

Approximately 50 percent of the population of Zimbabwe is at risk for developing malaria, which is the third most common cause of illness and death and accounts for 30 percent of outpatient visits at clinics. Annually, there are approximately 1.5 million malaria cases and 1,000 deaths from the disease. From 2004 to 2007, there was a marked decline in malaria incidence; however, it is unclear whether this resulted from intensified malaria control interventions or other factors, such as a weakened surveillance system and poor reporting.

Malaria transmission in Zimbabwe is largely unstable, with five strata that can be identified: malaria free; sporadic malaria; low and seasonal malaria; moderate and seasonal malaria; as well as high and seasonal malaria transmission areas (figure 19 Ministry of Health and Child Welfare Zimbabwe 2008 2008–2013). With a tropical climate modified by elevation, most of the country has conditions conducive to malaria transmission, particularly during the rainy season, during the November to April time period. Rainfall varies, with malaria transmission higher in high rainfall provinces and lower transmission in low rainfall areas. Out of the 62 rural districts, 45 are classified as having malaria transmission and, of these, 30 are classified as heavy burden districts (Ministry of Health and Child Welfare Zimbabwe 2008).

Antimalarial commodities are distributed with essential medicines and medical supplies (procured by the United Nations Children's Fun [UNICEF]) through the harmonized Zimbabwe Informed Push (ZIP)/Primary Healthcare Packages (PHCP) system. The collection, aggregation, and analysis of malaria logistics data are automated through AutoDRV software, and essential logistics data are readily available, on a quarterly basis, for decisionmaking. As compared to the other reporting systems in Zambia and Burkina Faso, the Zimbabwe data are unique in reporting on a quarterly basis.

Figure 19. Zimbabwe Map of Malaria Transmission Strata

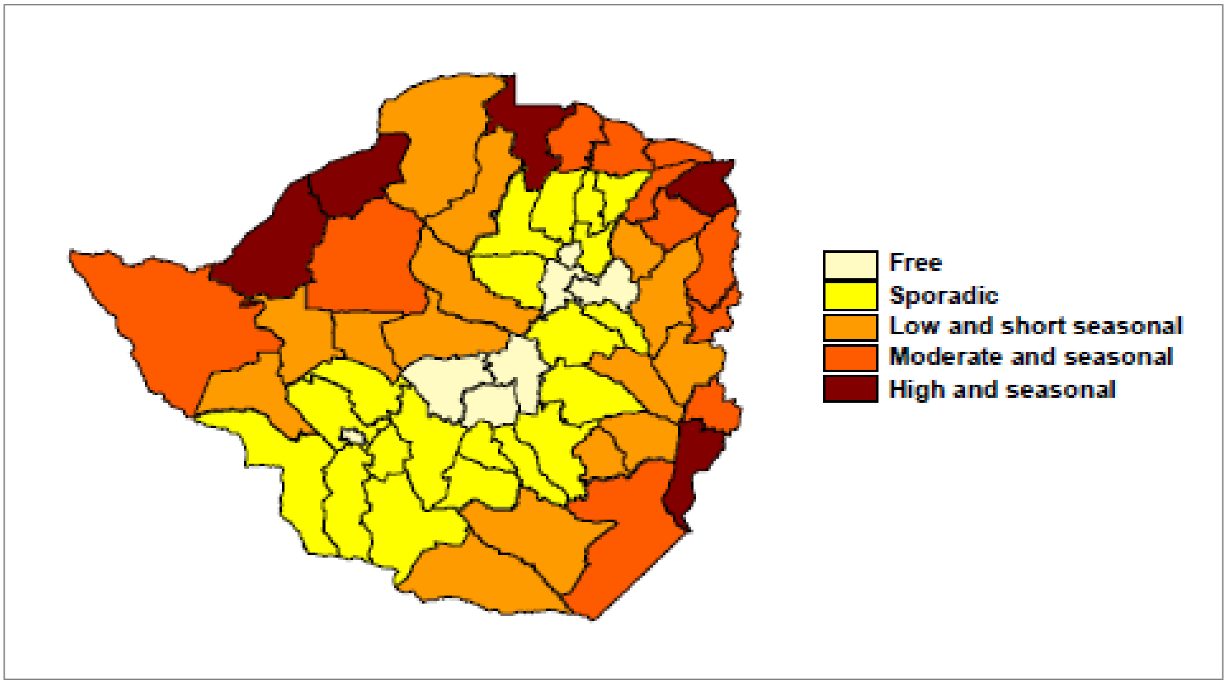

Creating the Seasonality Index for Zimbabwe

Figure 20 shows the geographic distribution of seasonal selected health facilities in Zimbabwe. All facilities receiving product through the ZIP/PCHP system were considered for analysis, because all facilities receiving antimalarial products were assumed to serve populations susceptible to malaria. The blue dots represent facilities with at least seven (out of eight) quarters of data, over the course of 2011–2012, and sufficient amounts of stock. For purposes of this analysis, an "insufficient" amount of stock was defined as a facility having, in that particular quarter, 30 days or more of AL stockouts. (Thirty days per quarter was determined as insufficient based on 10 days of stockout per month, which was a category created across all case studies.) These were excluded because of the misleading effect of stockouts on consumption patterns. Facilities without stock represent possible lost consumption, and facilities without sufficient stock over time may be less likely to attract clients in general, artificially lowering consumption further if the public perceives that a facility does not have medicines on hand and as a result does not seek facility services.

Figure 20. Geographic Distribution of Seasonal Selected Facilities, Zimbabwe

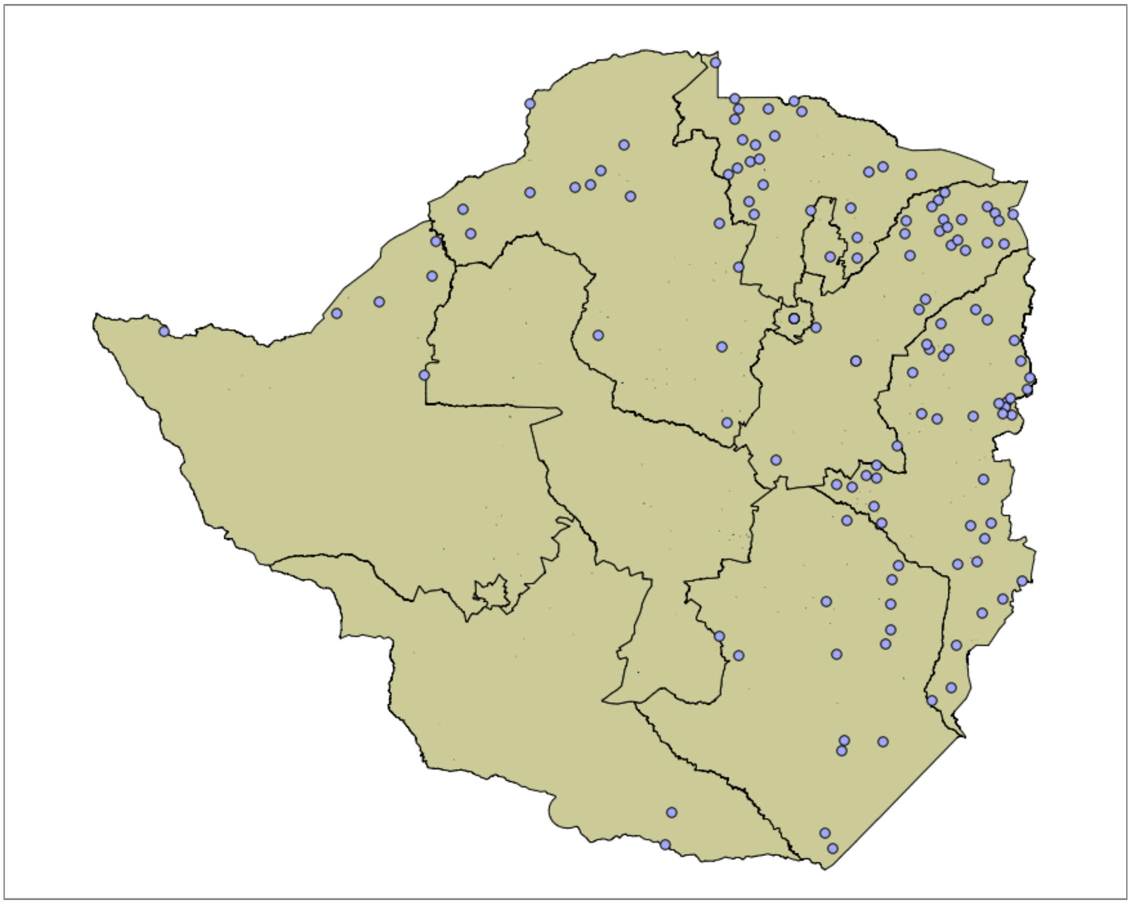

The other health facilities with sufficient data generally did not fit the seasonality pattern of the seasonal facilities, with either no discernible trend or, less frequently, a relatively flat trend over time. Of approximately 700 facilities with sufficient data, 159 were selected as having a seasonal trend. All quarters of the two years had relatively good reporting rates of at least 90 percent.

Figure 20 describes the steps taken to create the seasonality profile for Zimbabwe. As described, we used the seasonal facilities that had sufficient stock to avoid the effect of stockouts and at minimum seven quarters (close to two years) of data. As in Zambia, first line malaria treatment used in Zimbabwe is AL, with co-formulated tablets containing both artemether (20 mg) and lumefantrine (120 mg) and packaged in blister packs of four presentations: 1x6 tablets (5–14 kg); 2x6 tablets (15–24 kg); 3x6 tablets (25–34 kg); and 4x6 tablets (>34 kg). To investigate trends in AL consumption overall, we converted the consumption of the different presentations to the consumption of tablets overall. (The decision was made to calculate and combine consumption of tablets, rather than investigate individual presentations because of the limited availability of data.) This is particularly relevant given that presentations of AL may be cut down or combined to provide the appropriate dose for a patient in stockout situations.

Figure 21. Steps for Creating a Seasonality Index for Zimbabwe

1. Use "seasonal" facilities with sufficient stock and at least seven quarters of data.
2. Convert consumption of different presentations of AL to consumption of tablets.
3. Generate three-month aggregate consumption totals over two years.
4. Take the average of consumption across the same quarter in each of the two years.
5. Divide the average consumption for the four quarters of the year by the consumption in the Jan–Mar quarter to create seasonality indices.

Figure 22 and table 7 show the resulting seasonality index and the one-quarter ahead LSI for forecasting at the facility level. The seasonality profiles reveal the low point of consumption in the July–September time period, which matches with the general rainy season that stretches from November to April. In particular, the seasonality index shows that the ratio of consumption at its peak to consumption at its lowest is 3.06. Even this piece of knowledge can be helpful in inventory and resource planning, as consumption from July to September can provide six to nine months advance notice about the needs for malaria the following year.

The LSI for forecasting at the facility level used here is defined by the formula:

$$LSI(i) = \frac{Average(s_{i-1}, s_i, s_{i+1})}{Average(s_{i-3}, s_{i-2}, s_{i-1})}$$

Where s_i is the seasonality index for quarter *i,* since facilities generally have a very low lead time between their quarterly orders and delivery of inventory.[6] Through the rolling warehouse mechanism, deliveries are made immediately after data are collected and resupply is calculated.

Table 7. Zimbabwe Consumption, Seasonality Index and LSI

	Jan Mar	Apr–Jun	Jul Sept	Oct Dec
Average Consumption (2011–2012)	1104.64	1281.10	419.19	627.05
Seasonality Index	1	1.1597406	0.3794779	0.5676535
LSI	1.2945229	1.3040817	0.7724854	0.7668231

[6] Given the quarterly review period, we also considered alternative definitions of LSI, specifically defining LSI as the ratio of only one forward-looking seasonality index and one backward-looking seasonality index, and then the ratio of the average of two forward-looking and backward-looking seasonality indices. In comparisons, the current definition of LSI outperformed the alternatives.

Figure 22. Zimbabwe Seasonality and LSI

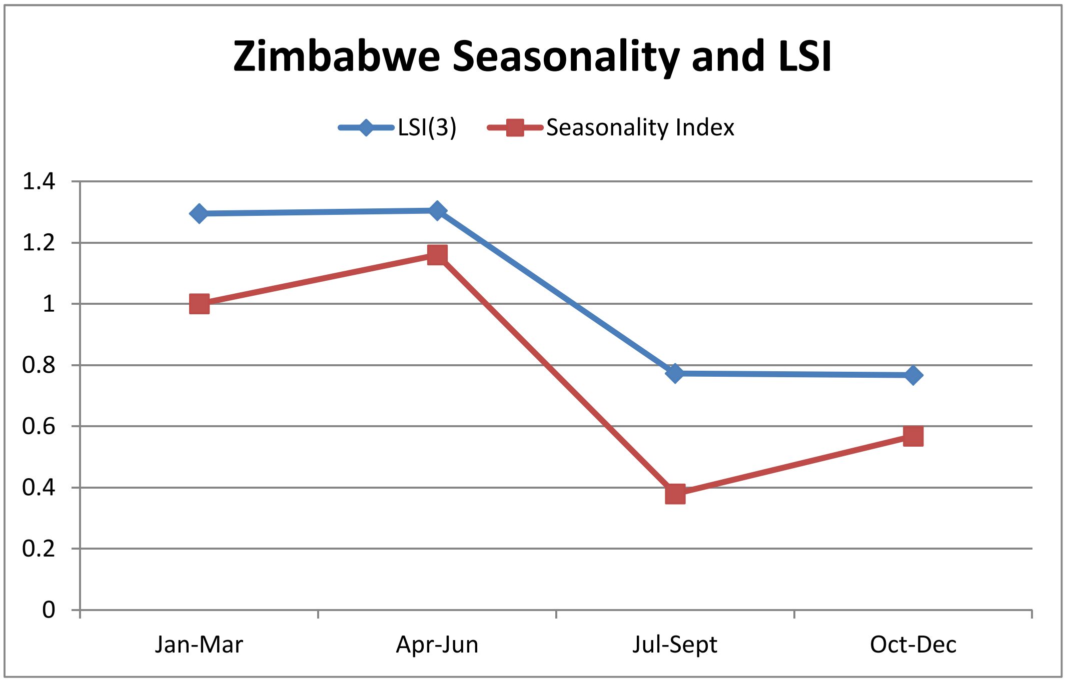

Testing LSI on Zimbabwe Consumption Data

To test the LSI approach on Zimbabwe data, additional forecasting methods were considered for comparison, and the models are all described in detail in the previous Zambia case study section of this report. Based on the predictions of these methods, inventory replenishment procedures can be adjusted and maximum stock levels determined based on established max/min levels. As with the Zambia case study, these models were compared using data from both seasonal facilities and aseasonal facilities.

The first forecasting method used is the simple rule, which in this case takes the previous quarter's consumption as the forecast for the next period. This rule is currently practiced in Zimbabwe, where AMC is calculated from consumption data from the last three months and adjusted for any days out of stock. The second forecasting method is the simple exponential smoothing. The third approach is triple exponential smoothing, which allows parameters to vary by facility and updated after every period of consumption. Finally, because the LSI approach uses the same seasonality factor across all facilities, the fourth and last model is based on the triple exponential smoothing, but where the seasonality index is artificially fixed to be the same across all facilities.

Evaluating Models Using Mean Average Percentage Error (MAPE)

As with the Zambia case study, the performance indicator used here to evaluate the performance of the different forecast approaches, including the LSI, is referred to as the mean absolute percentage error, or MAPE.

Figure 23 shows how the forecasting methods compare based on the MAPE for 539 aseasonal facilities. Since the MAPE can be driven by outliers arising from consumption actuals that are very small, on the x-axis we vary a cut-off threshold. This results in the MAPE score ignoring the forecast error for actual consumption that is lower than the threshold. In forecasting aseasonal facilities, the LSI approach uses the seasonality index that was derived from the seasonal facilities.

Our first observation is that the relative performance of the forecasting methods is roughly the same across the different cut-off thresholds. For aseasonal facilities, the triple exponential smoothing has the lowest error, while LSI is similar to simple exponential smoothing and triple exponential smoothing with fixed seasonality index across facilities. For reference, at a cutoff of 100, which is about 10.8 percent of the average quarterly consumption, the MAPE score for triple exponential smoothing is 60.5 percent, while for LSI the MAPE score is about 113 percent. The worst performer is the simple exponential smoothing with 122 percent.

Figure 23. Forecasting Performance Comparison for Aseasonal Facilities

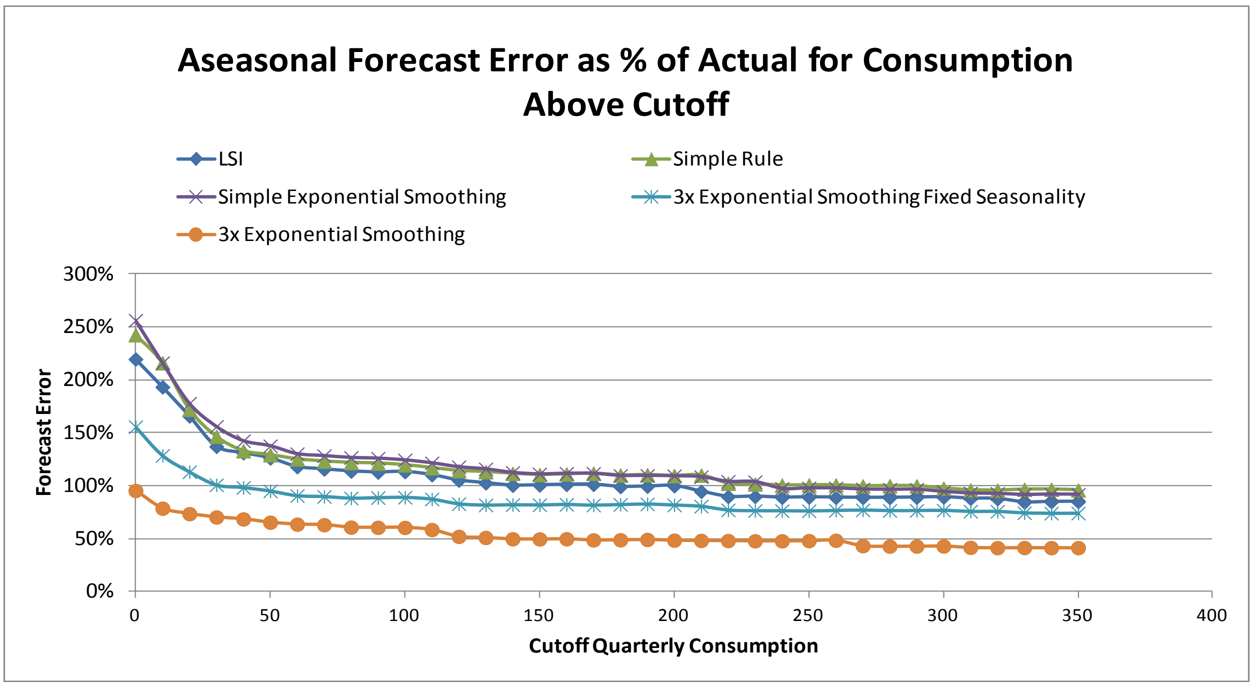

For 159 seasonal facilities in figure 24, we see similar results. For reference, at a cutoff of 90, which is about 10.42 percent of the quarterly consumption, the MAPE score for triple exponential smoothing is 29 percent; for triple exponential smoothing with fixed seasonality index, the MAPE score is 64 percent; while for LSI, it is 82 percent. The worst performer here is the simple rule, at 124 percent.

Figure 24. Forecasting Performance Comparison for Seasonal Facilities

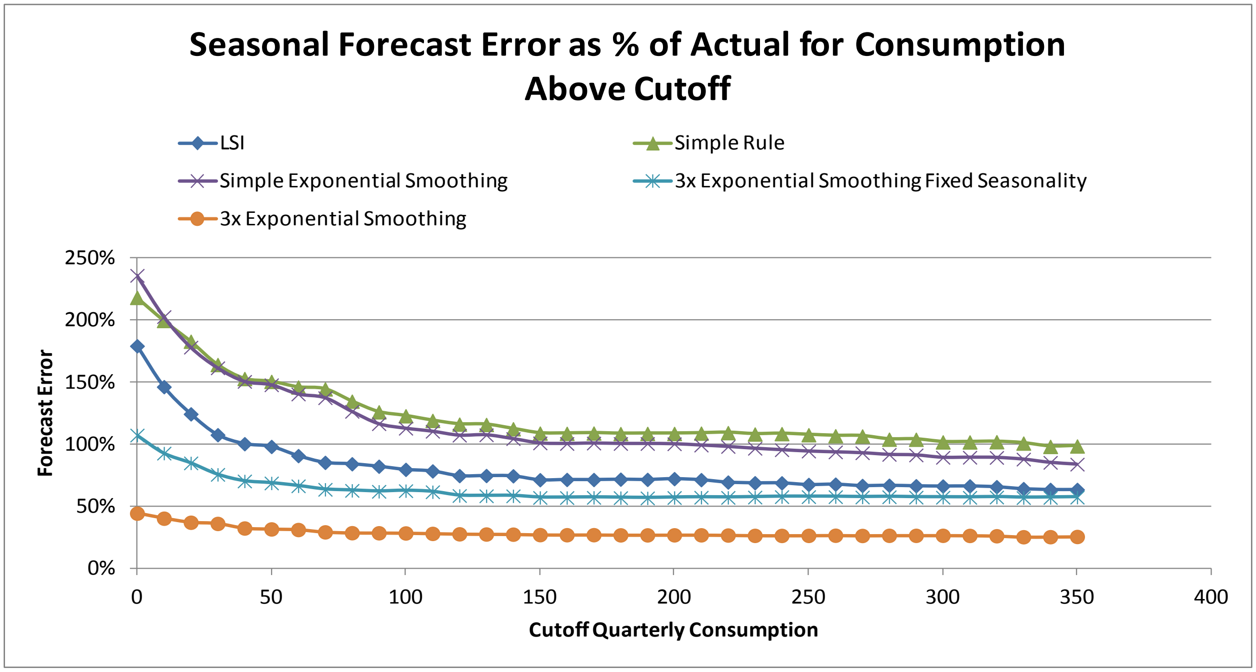

Evaluating Models Using Allocated Inventory Cost Methodology

The second approach to evaluating LSI in relation to other approaches involves simulating aggregate inventory costs by allocating a cost to excess inventory, plus a separate cost to lost consumption. In our approach, we normalize the cost of an excess unit of inventory by setting it to one, and then allowing the cost of one unit of lost consumption to vary from zero (or less costly than an excess unit of inventory) to a large number (that is more costly than an excess unit of inventory). We consider only the case where lost consumption is tracked, so that although there are stockouts, we still know what consumption would have been and can use this information in our forecasting method.

Figure 25 shows how for aseasonal facilities and tracking lost consumption, the inventory costs vary according to the cost of lost consumption, which can vary from zero to a large number: in the case of this graph, to 20. The lower the inventory costs, the better the performance of the forecasting method. Generally, costs should rise as cost of consumption increases; however, some methods may have a steeper slope if they generally have more stockouts than another method. What we see here is that the LSI performance tracks closely to the triple exponential smoothing with fixed seasonality index, which both perform closest to the triple exponential smoothing. For reference, at a lost consumption cost of 10, the triple exponential smoothing and triple exponential smoothing with fixed seasonality index for all facilities have a difference in cost of 49 percent, LSI has a difference of 54 percent, simple exponential smoothing has a difference of 69 percent, and the simple rule has a difference of 84 percent.

Figure 26 shows the performance for seasonal products again when lost consumption is tracked. Again, the LSI approach tracks closely to the triple exponential smoothing with fixed seasonality index, both of which perform closest to the LSI triple exponential smoothing. This results from the fact that in this instance, the forecasting method primarily results in an oversupply of inventory and little stockouts and lost consumption. For reference, at a lost consumption cost of 10, the triple exponential smoothing and triple exponential smoothing with fixed seasonality index for all facilities have a difference in cost of 50 percent, LSI has a difference of 50 percent, simple exponential smoothing has a difference of 70 percent, and the simple rule has a difference of 99 percent.

Figure 25. Inventory Replenishment Performance Comparison for Aseasonal Facilities and Tracking of Lost Consumption

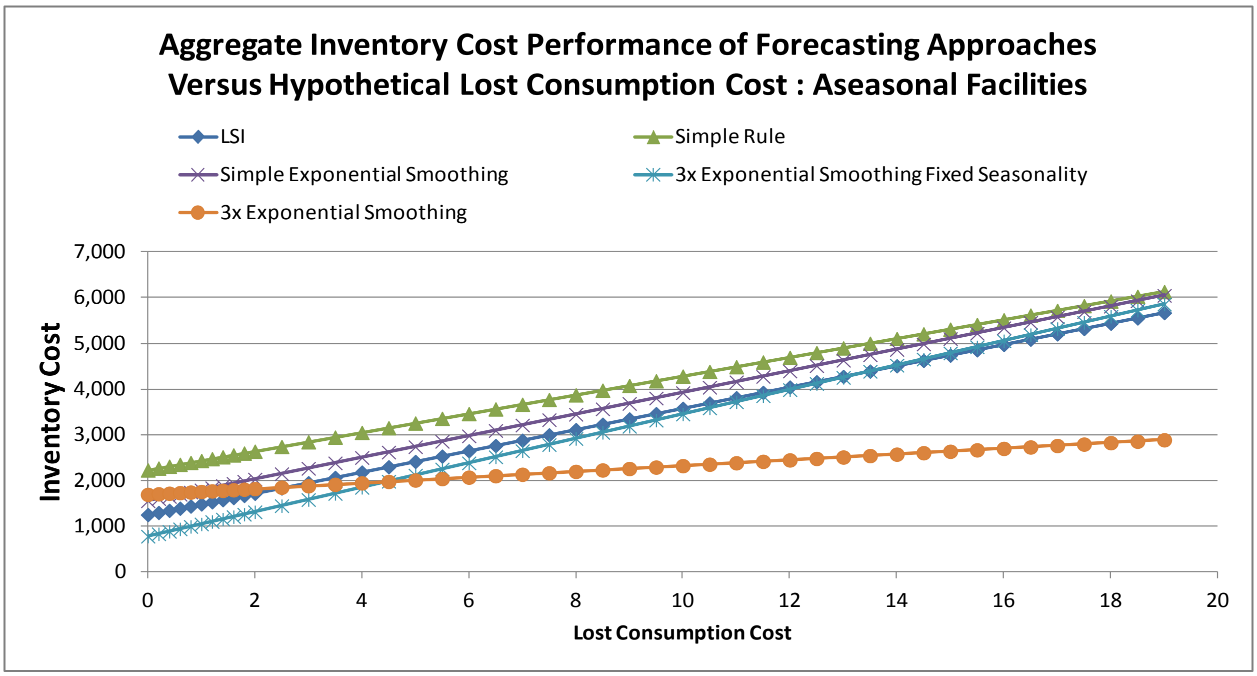

Figure 26. Inventory Replenishment Performance Comparison for Seasonal Facilities and Tracking of Lost Consumption

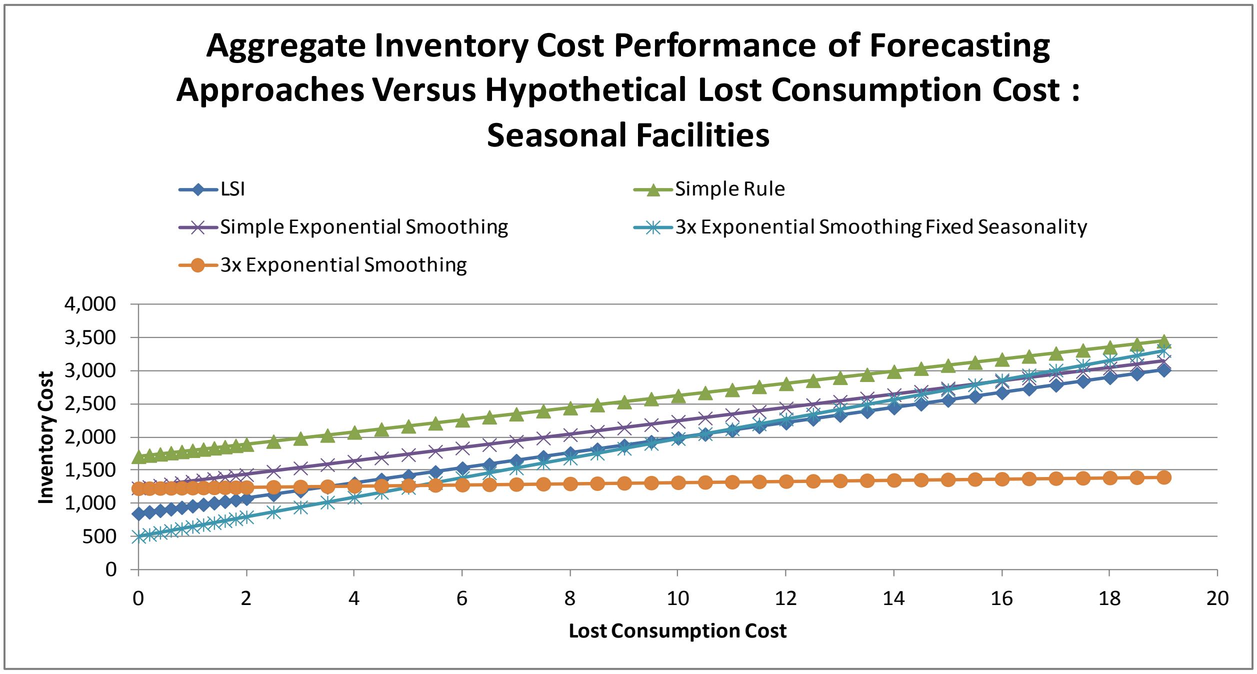

Table 8 captures all of the results that were individually referenced for each comparison. Generally, the LSI is as close to the sophisticated triple exponential smoothing as triple exponential smoothing with fixed seasonality on the more important inventory cost method of performing evaluation. However, the MAPE scores for LSI are somewhat high.

Table 8. Summary of Forecasting and Inventory Replenishment Performance, Comparison I

	MAPE (aseasonal)	MAPE (seasonal)	Inventory Costs (aseasonal, tracking)	Inventory Costs (seasonal, tracking)
3xExS	60%	29%	Lowest	Lowest
LSI	113%	82%	+54%	+51%
3xExSF	89%	62%	+49%	+50%
ExS	124%	117%	+69%	+70%
Simple	120%	126%	+84%	+99%

Zimbabwe Analysis Summary

The availability of consumption data at health facilities with good inventory support for close to two years facilitated our analysis here. A geographic analysis of so-called seasonal facilities revealed no discernible pattern to their distribution across Zimbabwe, lending support for considering their seasonal pattern as being representative for Zimbabwe as a whole. In addition, the seasonal facilities covered the various strata in malaria transmission that have been identified for Zimbabwe (see figure 19).

When we look at the distribution of facilities that were selected as having seasonal patterns, they covered most of the areas of Zimbabwe, including those with high rates of malaria and those with low rates of malaria. The incidence rates of malaria aren't so important for seasonality; what is important is that malaria has the same pattern through the year in terms of when there are peaks and relative heights of the peak versus the low season. (Consumption at the facility level is used to estimate forward consumption. Using consumption at the facility adjusts for the local incidence rates, whether rates are high or low.) So seasonal patterns can be representative, even in a country that has different incidence rates in different regions.

Tests on seemingly aseasonal facilities using this seasonal pattern seemed to confirm its suitability. There is the possibility that seasonality patterns developed for each of the various strata for malaria consumption could improve forecasting and inventory performance further, but this was not attempted in this analysis.

The Zimbabwe consumption data allowed for a similar testing of the LSI approach as with Zambia. Given its simplicity, ease of use, and ease of explanation, the LSI performed well in the tests, particularly those based on allocated inventory costs, having similar performance to the triple exponential smoothing with fixed seasonality.

As a pure forecasting method based on its MAPE scores, the LSI performance was not as comparable to the triple exponential smoothing with fixed seasonality as was the Zambia case study. However, we believe that the testing based on the allocated inventory costs should be given preference. In evaluating the performance of the LSI and along with its expected robustness in the event that consumption is not tracked, we would still conclude that the LSI can be recommended for use in Zimbabwe.

As is the case with the other countries, the AMC simple rule was not ideal for addressing seasonality in Zimbabwe. To some extent this is to be expected, as the simple AMC rule is reactive to demand by only looking at past consumption instead of proactive (i.e., anticipating demand). Generally, addressing seasonality usually requires some proactive activities in the supply chain. A quarterly review period was also probably not as helpful; theoretically, you would only realize demand was changing after three months, when a more frequent review period would allow for earlier detection of the high or low season and a faster reaction. Because data were only available by quarters, we could not test the impact of the quarterly review (e.g., by simulating what would happen if reviews were monthly or bimonthly).

Zimbabwe should benefit from LSI because it is fundamentally more sound than the AMC rule, as the LSI anticipates demand. Simulations of the resulting levels of inventory and stockouts also suggest LSI should outperform the AMC rule: lower levels of inventory, especially in the low season, and lower levels of stockouts in the peak. LSI generally also has some robustness to poor consumption data resulting from stockouts and does not have to be 100 percent accurate to provide benefits.

Burkina Faso Case Study

This section examines the Burkina Faso context, for the potential utility of the LSI approach and also to provide simulated results of the use of the LSI approach on historical consumption data.

Burkina Faso

Burkina Faso's population is estimated to be 17.8 million people in July 2013, with an annual growth rate of 3.1 percent (Central Intelligence Agency 2013). Forty-four percent of the population lives under the poverty line, with a GDP per capita in 2009 estimated to be U.S.$512 (Institut National de la Statistique et de la Démographie et ICF International 2012).

Malaria is one of the leading public health challenges in Burkina Faso. In 2011, malaria represented about 45 percent of all medical consultations at the health districts and 14 percent of all hospitalizations at the referral hospitals. Children under age five disproportionately suffer from the disease: Malaria accounts for 54 percent of consultations in health districts, and of the total deaths due to malaria, 79 percent are estimated to be among children under age five. Between 2010 and 2011, both the number of malaria cases reported and reported malaria deaths decreased slightly, with the number of malaria cases falling from 5,723,481 to 5,024,697, and the number of deaths falling from 9,034 to 7,001 (Direction Générale de l'Information et des Statistiques Sanitaires 2011).

The rainy season lasts approximately four months, from May/June to September and is shorter in the north of the country. There are three traditional malaria endemic zones, with the length of transmission seasons increasing from the south to the north of the country, as shown in figure 27. However, recent malaria incidence data from the National Malaria Control Program does not follow this epidemiological stratification, and the traditional malaria endemic zones may need to be reassessed. Typically, the northernmost Sahelian region has a short transmission season and is prone to malaria epidemics, whereas the Sudano-Sahelian region has a longer, seasonal malaria transmission season. The Sudano region in the south of the country experiences year-round malaria transmission. The majority of the population is at risk of malaria transmission, either seasonally or year round.

Figure 27. Malaria Transmission Zones, Burkina Faso

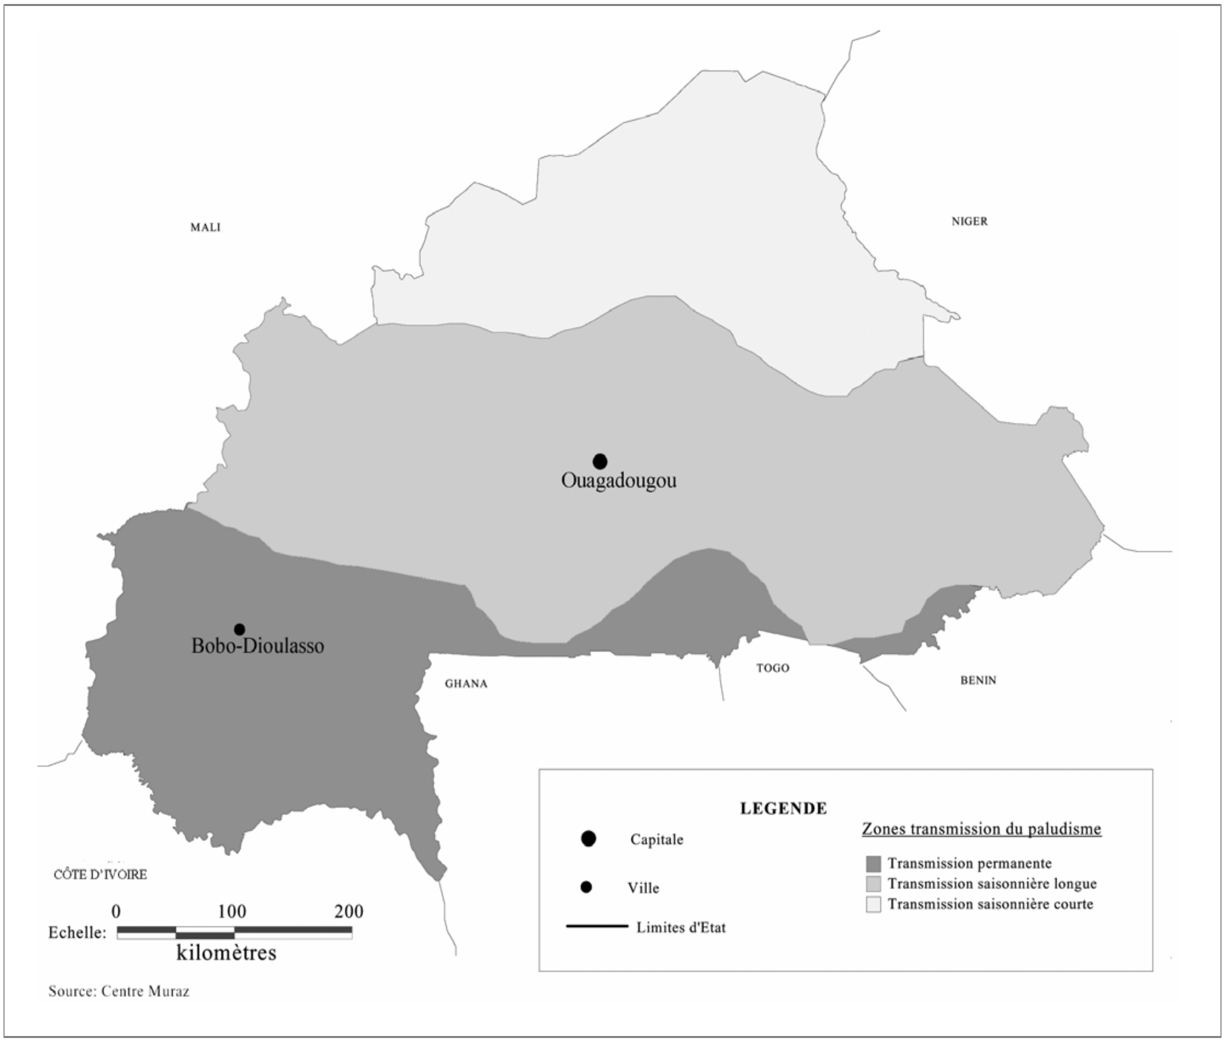

First line malaria treatment used in Burkina Faso is the fixed dose combination (FDC) artesunate/amodiaquine (AS/AQ). Four presentations are used for treatment: 25 mg/67.5 mg (3 tablets, 5–11 months); 50 mg/135 mg (3 tablets, 1–5 years); 100 mg/270 mg (3 tablets, 7–13 years); and 100 mg/270 mg (6 tablets, >13 years). Unlike AL, during a stockout situation, presentations for higher weight bands of FDC AS/AQ are generally discouraged from being cut down for children in lower weight bands, as cutting or crushing the tablets could lead to inaccurate dosages. However, the practice is common in countries with AS/AQ, and there is a precedent in country forecasting for AS/AQ to combine presentations and forecast in aggregate, rather than at the presentation level. In part due to data quality challenges encountered, in our analysis here we take advantage of this precedent by combining AS/AQ presentations for testing the LSI approach in inventory replenishment and forecasting.

Creating the Seasonality Index for Burkina Faso

Figure 28 shows the approximate geographic distribution of an initial set of health facilities that were selected as having seasonal trends in Burkina Faso. The seasonal facilities identified had at least eight months of data and sufficient stock to avoid the effect of stockouts. A month of data was excluded from analysis if the facility had insufficient stock, or was stocked out for 10 days or more during that month. Unlike the other two countries studied in this analysis, facility-level latitude and longitude data were not available. However, district-level geospatial coordinates were available and used to illustrate the approximate location of selected facilities. The selected 257 facilities represented 11 of 13 regions and 32 of 301 districts in Burkina Faso. One year of data was analyzed, although due to data quality issues, data were missing in April 2012, and only one facility had available data in January 2012.

Figure 28. Seasonal Selected Facilities, Burkina Faso

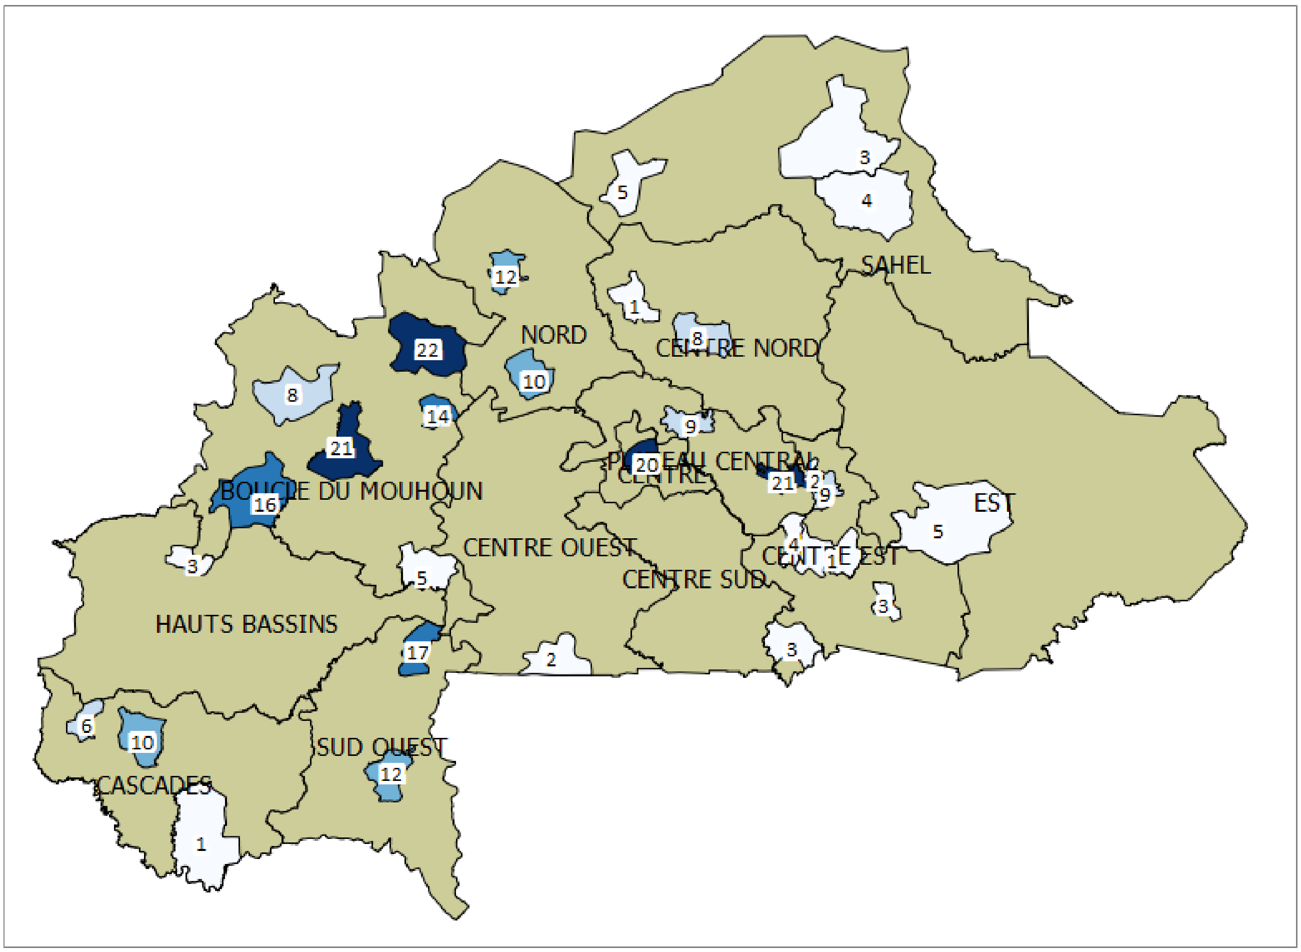

Although this initial set of facilities represented a significant number of regions, beyond the data issues in January and April 2012, there was also significant variation of the number of facilities with consumption data across months. For example, 159 facilities had sufficient data in December 2012, as compared to 257 facilities in July 2012 (see table 9).

Table 9. Number of Identified Seasonal Facilities with Consumption Data in 2012

Month	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Number of Facilities	1	227	242	0	243	255	257	257	252	207	203	159

This variation in the number of facilities with consumption data across the months in 2012, along with the data omissions in January and April 2012, necessitated refining the set of seasonal facilities further. From the initial set of seasonal facilities, facilities that had data for all of the remaining months (February, March, May–December) were selected. This resulted in 57 facilities out of the original 257. The consumption from these 57 facilities seemed to suggest two distinct set of seasonality indices, as captured in figure 29.

Figure 29. Burkina Faso Short and Long Seasonality Indices

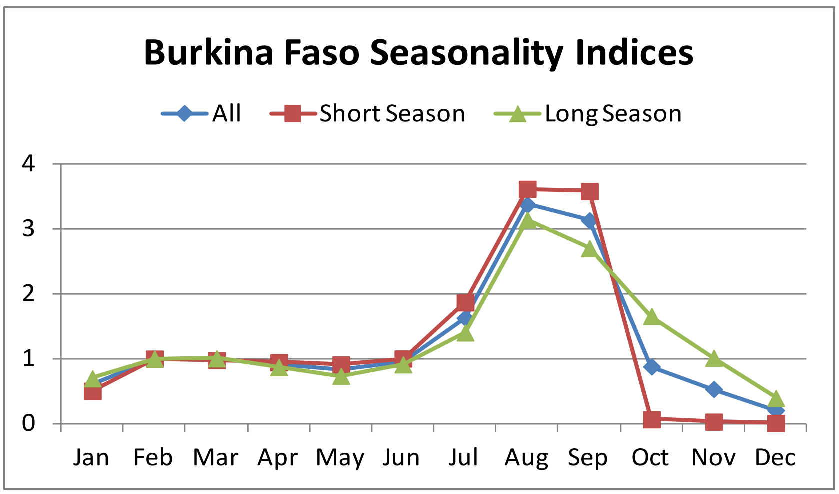

Figure 30. Steps for Creating Seasonality Indices in Burkina Faso

1. Use "seasonal" facilities with good inventory presence and at least 10 months of data.
2. Consumption patterns of individual presentations of AS/AQ were aggregated to create monthly consumption totals for 10 months (February, March, May–December).
3. Facilities with extremely low or zero consumption in October were grouped together into the short seasonality group, and the rest formed the long seasonality group.
4. Monthly consumption across all facilities in the same seasonality group was averaged.
5. Since consumption data for January and April were absent, monthly averages for those months were estimated as the average of the surrounding months: February and December for January, and March and May for April.
6. Seasonality indices were created by dividing monthly averages by the February monthly average.

Of the two sets of seasonality indices created, one seasonality index could be described as "short" and was based on 26 facilities, while the other could be described as "long" and was based on the remaining 31 facilities. Whereas both seasonality indices showed a peak in the months of August and September, the short seasonality showed a significant drop-off in consumption in the rest of the months of the year. The long seasonality showed high consumption in July and October, along with the peaks in August and September. This peak season for the "long" seasonality coincides with the Burkina Faso rainy season. Figure 31 shows the approximate geographic distribution of these 57 facilities.

Figure 31. Location of Facilities for "Short" (in Pink) and "Long" (in Red) Seasonality Indices

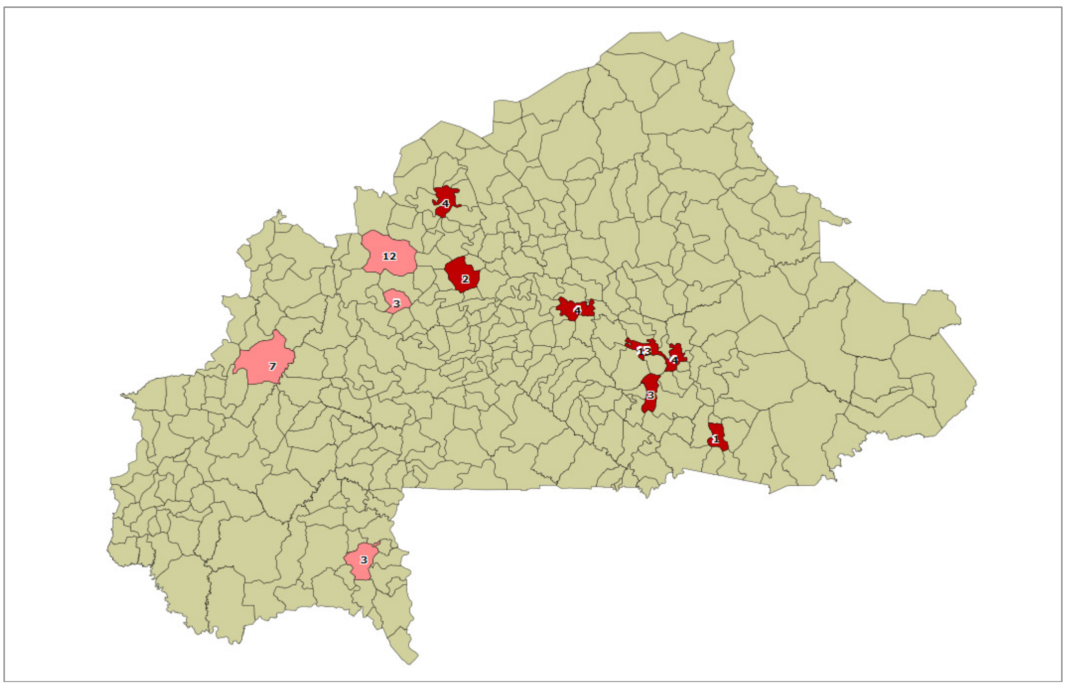

Due to the unlikelihood of consumption in the last three months of year being very low, we focused on the "long" set of seasonality indices. Results from a comparison of the long seasonality index with the short seasonality index for forecasting all available seasonal facilities supported this choice.

Figure 32 and table 10 show the resulting long seasonality index and the one quarter ahead LSI for forecasting at the facility level. The seasonality index shows that the ratio of consumption at its peak (August) to consumption at its lowest (December) is 8.

The LSI for forecasting at the facility level used here is defined by the formula:

$$LSI(i) = \frac{Average(s_{i-1}, s_i, s_{i+1})}{Average(s_{i-3}, s_{i-2}, s_{i-1})}$$

Where s_i is the seasonality index for month *i,* since facilities generally have a very low lead time between their monthly orders and delivery of inventory.

Table 10. Burkina Faso Long Seasonality Index and LSI

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Long Seasonality	0.70	1.00	1.00	0.87	0.73	0.91	1.40	3.14	2.70	1.66	1.01	0.39
LSI	0.68	1.28	1.37	0.97	0.87	1.17	2.17	2.38	1.37	0.74	0.41	0.39

Figure 32. Burkina Faso Long Seasonality and LSI

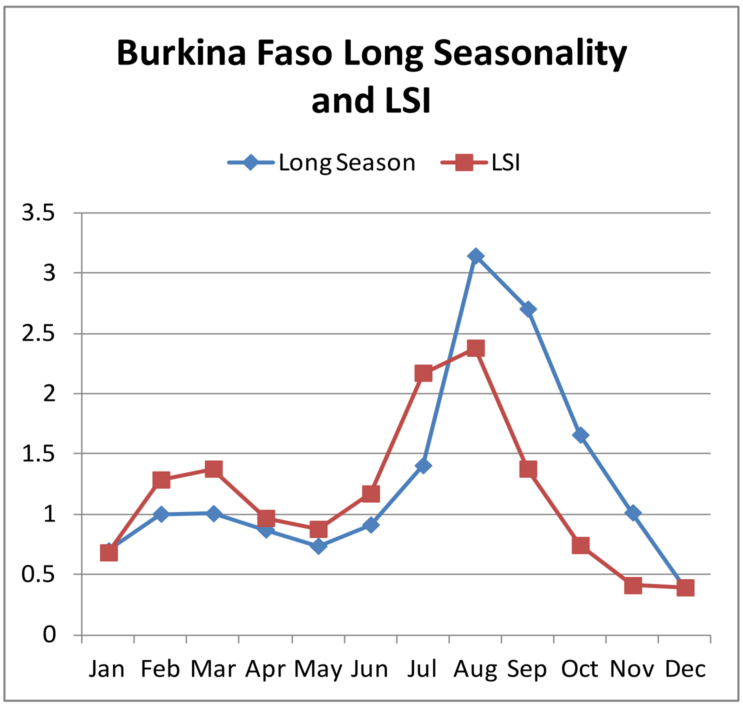

Testing LSI on Burkina Faso Consumption Data

To test the LSI approach on Burkina Faso data, additional forecasting methods were considered for comparison, and the models are described in detail in the Zambia case study section of this report. Unlike the Zambia and Zimbabwe case studies, these models were only compared using data from seasonal facilities.

The first forecasting method used is the simple rule, which in this case takes the average of the previous three months of consumption as the forecast for the next period. The second forecasting method is the simple exponential smoothing. The third approach is triple exponential smoothing, and the fourth is the triple exponential smoothing but where the seasonality index is artificially fixed to be the same across all facilities.

Evaluating Models Using Mean Average Percentage Error (MAPE)

As with the Zambia case study, the performance indicator used here to evaluate the performance of the different forecast approaches including the LSI is referred to as the mean absolute percentage error, or MAPE.

Figure 33 shows how the forecasting methods compare based on the MAPE for 257 seasonal facilities. Since the MAPE can be driven by outliers arising from consumption actuals that are very small, on the x-axis we vary a cut-off threshold. This results in the MAPE score ignoring the forecast error for actual consumption that is lower than the threshold. In forecasting seasonal facilities, the LSI approach uses the seasonality index that was derived from the seasonal facilities.

For seasonal facilities, the triple exponential smoothing has the lowest error, while LSI is similar to triple exponential smoothing with fixed seasonality index across facilities. For reference, at a cutoff of 30, which is about 11.45 percent of the average monthly consumption, the MAPE score for triple exponential smoothing is 17 percent, while for LSI the MAPE score is about 59 percent. The worst performer is the simple rule, at 76 percent.

Figure 33. Forecasting Performance Comparison for Seasonal Facilities

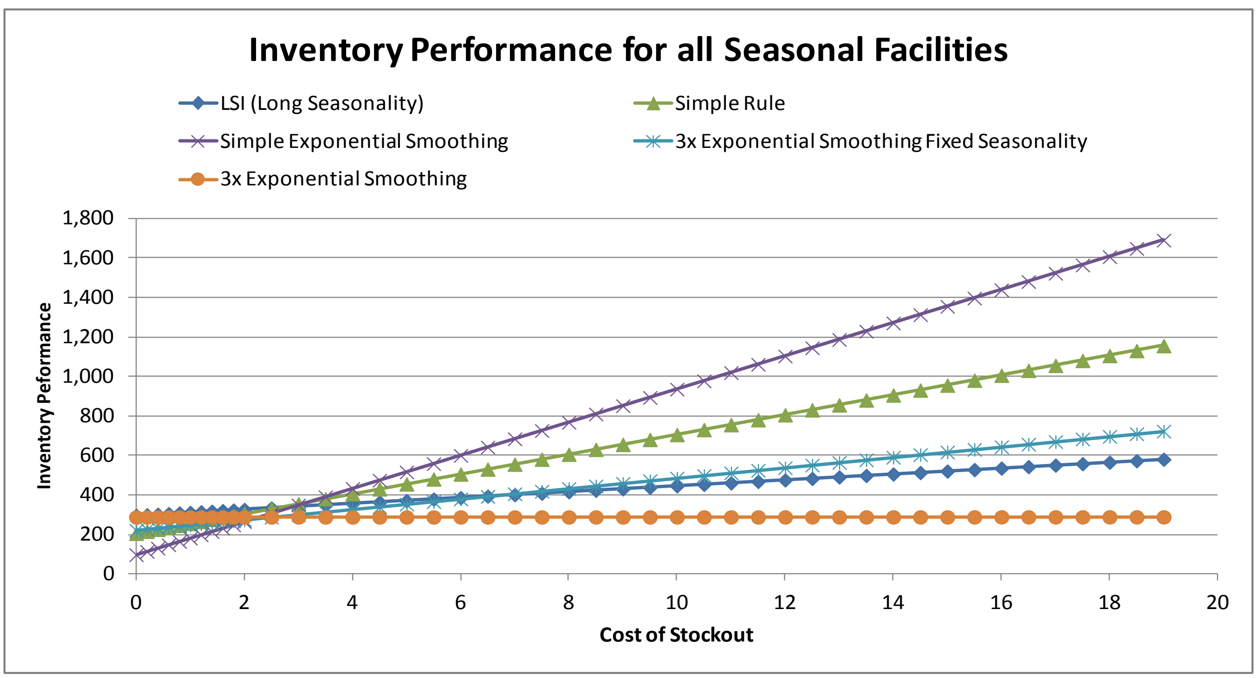

Evaluating Models Using Allocated Inventory Cost Methodology

The second approach to evaluating LSI in relation to other approaches involves simulating aggregate inventory costs, by allocating a cost to excess inventory and a separate cost to lost consumption. In our approach, we normalize the cost of an excess unit of inventory by setting it to one, and then allowing the cost of one unit of lost consumption to vary from zero (or less costly than an excess unit of inventory) to a large number (that is more costly than an excess unit of inventory). We consider only the case where lost consumption is tracked, so that although there are stockouts, we still know what consumption would have been and can use this information in our forecasting method.

Figure 34 shows how for seasonal facilities and tracking lost consumption, the inventory costs vary according to the cost of lost consumption, which can vary from zero to a large number—in the case of this graph, to 20. The lower the inventory costs, the better the performance of the forecasting method. Generally, costs should rise as cost of consumption increases; however, some methods may have a steeper slope if they generally have more stockouts than another method.

What we see here is that the LSI performance is closest to the triple exponential smoothing and outperforms the other forecasting methods. For reference, at a lost consumption cost of 10, the triple exponential smoothing and LSI have a difference in cost of 55 percent; triple exponential smoothing with fixed seasonality index for all facilities, a difference of 68 percent; the simple rule, a difference of 145 percent; and simple exponential smoothing, a difference of 225 percent.

Figure 34. Inventory Replenishment Performance Comparison for Seasonal Facilities and Tracking of Lost Consumption

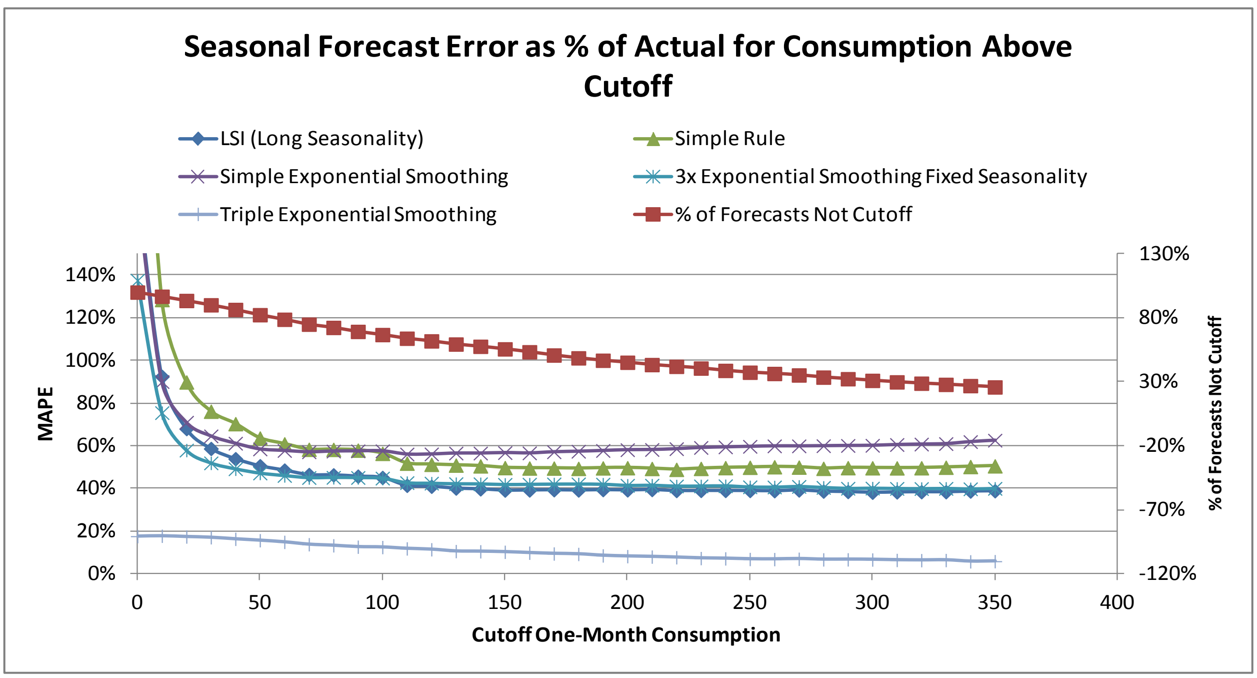

Table 11 captures all of the results that were individually referenced for each comparison. Generally, LSI on the more important inventory cost method of performance evaluation is second only to the sophisticated triple exponential smoothing. In addition, the MAPE scores for LSI are adequate.

Table 11. Summary of Forecasting and Inventory Replenishment Performance Comparison

	MAPE (seasonal)	Inventory Costs (seasonal, tracking)
3xExS	17%	Lowest
LSI	59%	+55%
3xExSF	52%	+68%
ExS	65%	+225%
Simple	76%	+145%

Burkina Faso Analysis Summary

In Burkina Faso, the availability of consumption data at health facilities with good inventory support was not ideal for our analysis here. Data were only available for 10 months of the year and only for some facilities. Regarding those that suggested little to no consumption in the last three months of the year, the data were considered suspect. This resulted in the analysis of a very small set of 31 facilities, located in a narrow band running northwest to southeast across the middle of the country. These facilities provided the seasonality profile used to evaluate the LSI approach. Interestingly, the seasonality profile given by these 31 facilities provided both good forecasting and allocated inventory cost performance on the original set of 257 seasonal facilities, which were more broadly distributed across the country.

In particular, as a pure forecasting method based on its MAPE scores, the LSI performance was as comparable to the triple exponential smoothing with fixed seasonality as was the Zambia case study. In addition, similar to the Zambia case study, based on allocated inventory costs, the LSI approach was outperformed only by the triple exponential smoothing.

While these results do lend some support for the use of the LSI approach in Burkina Faso, some caution should be expressed given that the seasonality index used here has not been developed from more robust data and also has not been tested on additional facilities(i.e., aseasonal facilities) from Burkina Faso. In addition, when data availability improves, it would be appropriate to evaluate the LSI approach on the various presentations of AS/AQ, given that there are ongoing attempts to discourage substituting across presentations when administering medication.

Conclusion

This concept paper describes an approach for enhancing the Simple AMC rule to handle seasonal commodities, while maintaining some of its simplicity to continue to meet complex needs. The approach operationally involves multiplying the AMC by indices that compensate for seasonality—referred to as Look-Ahead Seasonality Indices—before applying the maximum months of stock level. The LSI approach is intended to provide a practical, easy-to-use, enhanced means of calculating resupply for seasonal commodities, at all levels of the supply chain. Because of its similarity to the simple AMC rule, the LSI approach is one that should be relatively easy to introduce into healthcare settings in developing countries. Given its simplicity, it should also be relatively easy to integrate into electronic-based information systems.

Many of the challenges for addressing the resupply of seasonal commodities have been examined here, and the LSI approach at least partially, if not completely, addresses all of these concerns. In all country case studies that tested consumption data on antimalarials from Zambia, Zimbabwe, and Burkina Faso, the LSI approach's performance was commendable and could be recommended both as an inventory replenishment mechanism, under conditions that correspond with the healthcare settings in developing countries, as well as a forecasting method.

For all countries, the LSI approach was evaluated in comparison to four other common models (the simple AMC rule, simple exponential smoothing, triple exponential smoothing, and triple exponential smoothing with fixed seasonality index) and was evaluated in two ways: (1) as an inventory replenishment mechanism using the allocated inventory costs and (2) as a forecasting method using the MAPE. (Due to the intended need for LSI, to aid resupply in healthcare settings in developing countries, the testing of inventory replenishment based on the allocated inventory costs is most important and should be given preference.) As a result of this process, each country case study provides an estimate of the seasonality index and corresponding LSI that could support application of the LSI approach or similar approaches based on seasonality indices within these countries.

In Zambia, the LSI approach performed best in comparison with the other models, in part because the availability of monthly consumption data over nearly a two-year time period facilitated analysis. The LSI performance for both seasonal and aseasonal facilities was commendable given how closely it tracked the consistently superior (and sophisticated) model, triple exponential smoothing. The LSI even outperformed triple exponential smoothing when lost consumption was not tracked (a common phenomenon in many countries).

In Zimbabwe, the LSI performed well in the tests, particularly those based on allocated inventory costs, having similar performance to the triple exponential smoothing with fixed seasonality. However, as a pure forecasting method based on its MAPE scores, the LSI performance was not as comparable to the Zambia case study. However, we would still conclude that the LSI can be recommended for use in Zimbabwe. A possible limitation of the Zimbabwe data was its availability only on a quarterly basis; theoretically, you would only realize demand was changing after three months, when a more frequent review period would allow earlier detection of the high or low season and a faster reaction.

A main difference between the country case studies was the issue of data availability, particularly in Burkina Faso. Special care must be taken with Burkina Faso's results since, although the LSI approach performed well in testing, it would be preferred that the supporting seasonality index had been based on more robust consumption data and that testing of the LSI approach had been more extensive. Data were only available for 10 months of the year and only for some facilities. When data availability issues are addressed, the LSI approach should be tested with additional months of data and should also be tested on aseasonal facilities (as was the case with Zambia and Zimbabwe). It would also be appropriate to analyze the approach using individual presentations of AS/AQ given ongoing attempts to discourage substituting across presentations when administering medication.

Because of its reliance on historical consumption data, the LSI approach might, in theory, be applied in other countries where the President's Malaria Initiative (PMI) and the National Malaria Control Programs (NMCPs) have access to consistent, relatively robust data, and the simple AMC rule is currently in place for inventory replenishment. A limitation of the LSI approach is its capacity to look too far ahead in the future (e.g., one year). The approach may be best suited for resupply calculations, as well as shorter term forecasts, as the further ahead that the LSI is used to look ahead, the greater the chance that level of the period may have changed compared to that of the most recent consumption. An assessment of the LSI approach may also be of value in order to test the method, its practicality, and its results in a controlled field setting in comparison to current resupply methods.

In all countries, the simple AMC rule was outperformed for addressing seasonality in comparison to the LSI approach. To some extent this is to be expected, as the simple AMC rule is reactive to demand by only looking at past consumption, instead of proactive (i.e., anticipating demand as in the case of the LSI approach). Generally, addressing seasonality usually requires some proactive activities in the supply chain.

As any quick scan of the literature reveals, there are many ways to approach both inventory replenishment and forecasting. The approach introduced and assessed here is only one such method. The discussions here should be considered a case for the legitimacy of the LSI approach for seasonal commodities, but not to the exclusion of other approaches, or that the suitability of the LSI approach should be taken for granted in all situations involving seasonal commodities. As always, the particular context of the country situation should first be carefully considered in the testing and potential application of the LSI approach.

References

Central Intelligence Agency. 2011. *The World FactBook: Africa: Zambia*. Available at https://www.cia.gov/library/publications/the-world-factbook/geos/za.html (accessed January 18, 2011).

Central Intelligence Agency. 2013. *The World FactBook: Africa: Burkina Faso*. Available at https://www.cia.gov/library/publications/the-world-factbook/geos/uv.html (accessed September 25, 2013).

Central Intelligence Agency. 2013. *The World FactBook: Africa: Zimbabwe*. Available at https://www.cia.gov/library/publications/the-world-factbook/geos/zi.html (accessed September 25, 2013).

Direction Générale de l'Information et des Statistiques Sanitaires. 2011. *Annuaire statistique 2011*.

Gallien, J., Z. Leung, and P. Yadav. 2012. "Rationality and Transparency in the Distribution of Essential Drugs in Sub-Saharan Africa: Analysis and Design of an Inventory Control System for Zambia." London Business School Working Paper, February 7, 2012.

Institut National de la Statistique et de la Démographie (INSD) et ICF International. 2012. *Enquête Démographique et de Santé et à Indicateurs Multiples du Burkina Faso 2010*. Calverton, Md: Institut National de la Statistique et de la Démographie (INSD) et ICF International.

Makridakis, S., S. C. Wheelwright, and V. E. McGee. 1983. *Forecasting: Methods and Applications*. New York, NY: Wiley.

Ministry of Health and Child Welfare Zimbabwe. n.d. *National Malaria Control Strategy 2008-2013*. Zimbabwe: Ministry of Health and Child Welfare Zimbabwe.

NIST/SEMATECH. 2012. *NIST/SEMATECH e-Handbook of Statistical Methods*. Available at http://www.itl.nist.gov/div898/handbook/ (accessed July 13, 2013).

United Nations, Department of Economic and Social Affairs, United Nations Statistics Division. 2010. *Millennium Development Goals Indicators: Population Below National Poverty Line, Rural, Percentage*. Available at http://mdgs.un.org/unsd/mdg/SeriesDetail.aspx?srid=583&crid=894 (accessed January 18, 2011).

United National Development Programme. n.d. Human Development Reports, International Human Development Indicators. Available at http:// http://hdr.undp.org/en/countries (accessed September 18, 2013).

Winters, P. R. April 1960. "Forecasting Sales by Exponentially Weighted Moving Averages." *Management Science*. 6(3): 324–342.

World Health Organization. 2010. *World Malaria Report*. Available at http://www.who.int/malaria/world_malaria_report_2010/en/.

Appendix A

Holts-Winter Triple Exponential Smoothing

The Holts-Winter Triple Exponential smoothing is a forecasting approach that incorporates both trend and seasonality (Winters 1960). The method is referred to as triple exponential smoothing because a general smoothing technique involving a weighted average of the old forecast and the new data is used for the (1) seasonality indices, (2) the trend parameter, and (3) the overall forecast. The basic equations of the method is given by:

$$S_t = \alpha \frac{y_t}{I_{t-L}} + (1-\alpha)(S_{t-1} + b_{t-1}) \qquad \text{OVERALL SMOOTHING}$$
$$b_t = \gamma(S_t - S_{t-1}) + (1-\gamma)b_{t-1} \qquad \text{TREND SMOOTHING}$$
$$I_t = \beta \frac{y_t}{S_t} + (1-\beta)I_{t-L} \qquad \text{SEASONAL SMOOTHING}$$
$$F_{t+m} = (S_t + mb_t)I_{t-L+m} \qquad \text{FORECAST}$$

Where

y is the observation

S is the smoothed observation

b is the trend factor

I is the seasonal index

F is the forecast at m periods ahead

t is an index denoting a time period, and

α, β, and γ are constants that must be estimated in such as way that the mean square error is minimized (NIST/SEMATECH 2012).

Appendix B

General Formula for LSI

The general formula for LSI for a tier of the supply chain that has lead time l and review period p is given in Figure 35. Here the numerator is the average of the seasonality indices of the p periods starting when the inventory would arrive, which is l periods after period I, along with the indices of the periods that precede and follow.

Figure 35. LSI Formula for General Supply Chain Tier

Seasonality indices of *p* periods starting when the inventory would arrive, which is *l* periods after period *i.*

Additional indices add some correction if seasonality has shifted a bit earlier or later

$$LSI_i = \frac{Ave(s_{i+l-1}, \overbrace{s_{i+l}, \ldots, s_{i+l+p}}, s_{i+l+p+1})}{Ave(s_{i-3}, s_{i-2}, s_{i-1})}$$

Three period average of past seasonality indices (for which consumption will be known by period *i*)

Appendix C

Use of LSI in Higher Tiers of the Supply Chain

Higher tiers of the supply chain have a choice in terms of the information to which it will respond when managing its inventory: either responding to its own consumption information (i.e., the quantities that it ships) or consumption information from stages further down in the supply chains (e.g., patient consumption at the facility level). It's tempting to think that this information should be the same. But, generally, it turns out that this is not the case. Figure 36 shows the orders from a facility that uses a simple AMC rule that takes the average of the last two periods and compares it to the actual consumption at the facility. It shows that the orders from the facility can at times overestimate, while at other times underestimate the actual consumption at the facility. An objective measure of the variability over time—the standard deviation—of orders and of consumption shows that orders are almost twice as variable as consumption. In some cases of an underestimate of demand, orders are actually zero while consumption never actually falls to zero. In such a setting, using orders from the facility can give different results in managing inventory than using consumption at the facility. The driving factor in this decision is usually the lack of visibility into consumption at the facility. In many developing country settings, this information is incorrect due to errors in the field or lack of recording of this data or, if the data is indeed captured, there is usually a significant delay before the information is made available to upstream tiers in the supply chain.

Figure 36. Comparison of Orders and Actual Consumption for a Facility Using Simple AMC Rule

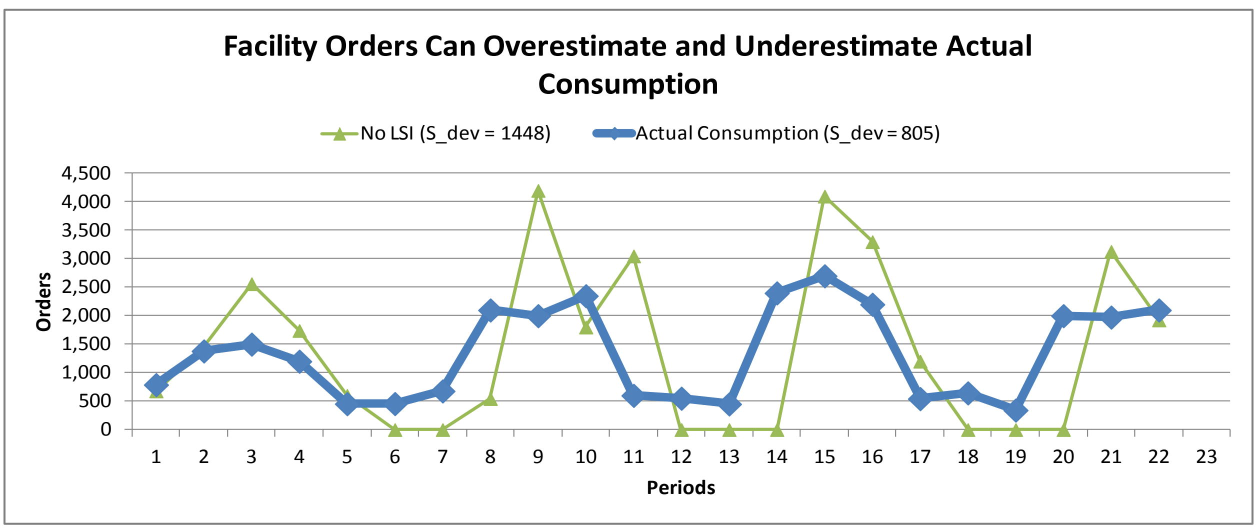

We can consider two settings, one where consumption at the facility is visible to upstream tiers in the supply chain and another where it is not visible. We use simulation to get a sense of what the effect of performance would be in each of these settings. Primarily, we assume a simple supply chain with two levels and only one entity at each level. The lowest level in this situation will be designated the Facility while the upstream level will be designated Central.[7] Although the typical supply chain will have multiple facilities and in some cases more than two tiers in the supply chain, the insights derived from this simple model will generally carry over to more complex settings. Under either visibility assumption we consider three possibilities: (1) both entities use LSI, (2) only the Facility uses LSI, or (3) neither entity uses LSI but instead use a version of the simple AMC rule.[8] In any setting, consumption is visible or not visible to Central. When the Facility is not on LSI, it uses consumption in the previous period as the AMC. When the Facility is on LSI, it uses the average of the most recent three periods of consumption as the AMC and multiplies this by the one period LSI to get a forecast for future consumption.[9] We also assume that in the event of any stockouts at the Facility, the consumption that would have occurred if inventory was available is recorded and used in forecasting.

Our expectation is that the use of LSI should improve performance, especially when both tiers use it.

Consumption Is Not Visible to Upstream Stages

Here we assume that when Central is not on LSI, it uses an average of the four most recent orders from facility and shipments to the Facility, like the AMC. When Central is on LSI, it use this AMC multiplied by one period LSI to get a forecast for future orders.[10] The first observation from the results is seen in figure 37, namely, that the use of the LSI by the Facility reduces the variability of orders from the Facility. Whereas the simple AMC rule increased the standard deviation of orders by 80 percent over that of consumption, the LSI approach only increases the standard deviation by 24 percent. In addition, orders from the Facility using the LSI approach are less likely to be zero compared to the simple AMC rule. This should aid the forecasting of orders from the Facility by Central since orders are less of an underestimate and/or overestimate of actual consumption.

---

[7] We assume that the Facility orders at the beginning of the period and that orders are replenished immediately from inventory at Central. Central then procures and receives its inventory before the beginning of the next period.

[8] When the Facility is not on LSI, it uses the consumption in the previous period as the AMC. When Central is not on LSI, it uses an average of the four most recent orders from the Facility and shipments to the Facility.

[9] The top-up multiplier for the facility is 2.

[10] We allow the top-up multiplier of Central to vary in the simulation to change the level of inventory in the supply chain and to see the effect on stockouts and lost consumption.

Figure 37. Comparison of Orders from Facilities Using an LSI Approach and Using the Simple AMC Rule

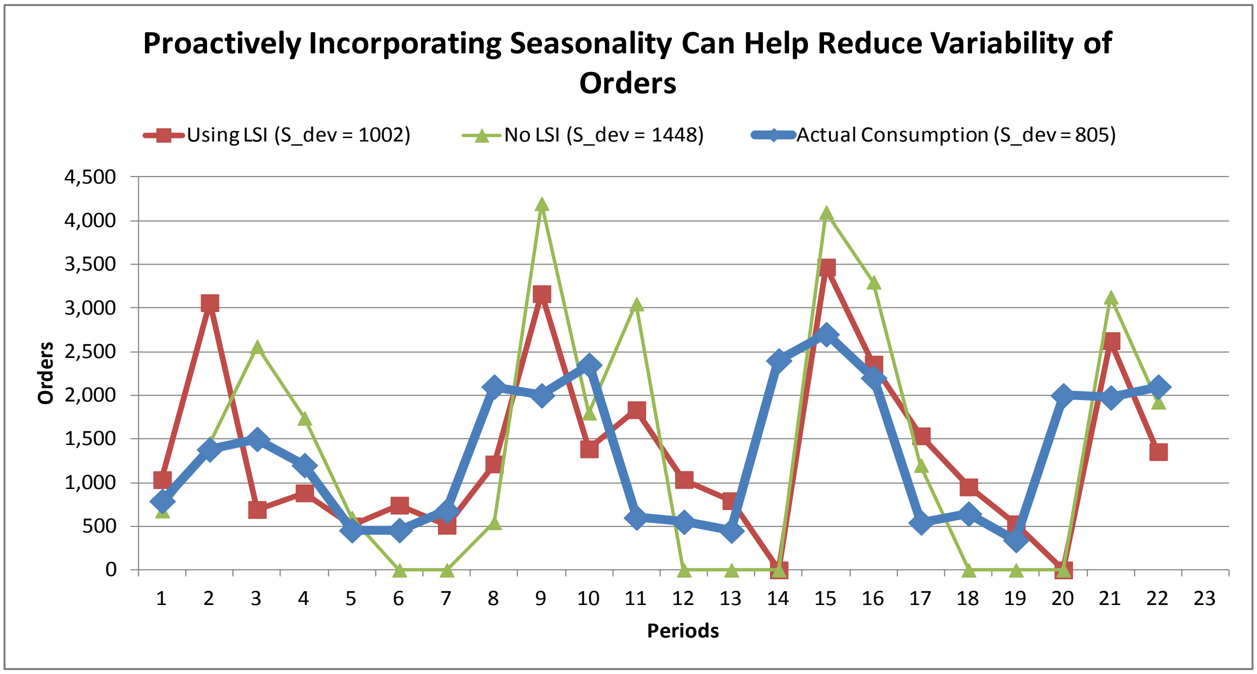

Figure 38 shows the system inventory and the average percentage of consumption that is lost at the Facility as a result of stockouts for each of the three scenarios: (1) both entities use LSI, (2) only the facility uses LSI, or (3) neither entity uses LSI.[11] Here, the system inventory is the inventory at Central after sending inventory to the Facility as well as the inventory at the Facility after meeting demand. Generally, the lower the curve on this graph, the better performing the scenario, since it implies that the scenario is able to meet the same level of consumption but with lower levels of inventory in the entire supply chain. The results show that, except for low levels of lost consumption, the scenario where both entities use LSI outperforms the other scenarios, followed by the scenario where the facility only uses LSI.

In summary, in a setting where there is a lack of visibility into facility consumption, it is always beneficial for the Facility to use LSI even if Central does not. This results partly from the reduction in variability in orders that result from using the LSI at the Facility that makes orders more representative of consumption and thus simulates some visibility.

[11] These curves were generated by varying the top-up multiplier for Central, which changed the level of inventory in the supply chain, which then influenced stockouts and lost consumption.

Figure 38. Inventory Service-Level Trade-off Comparison for Two-Level Supply Chain When Consumption Data Is Not Shared Upstream

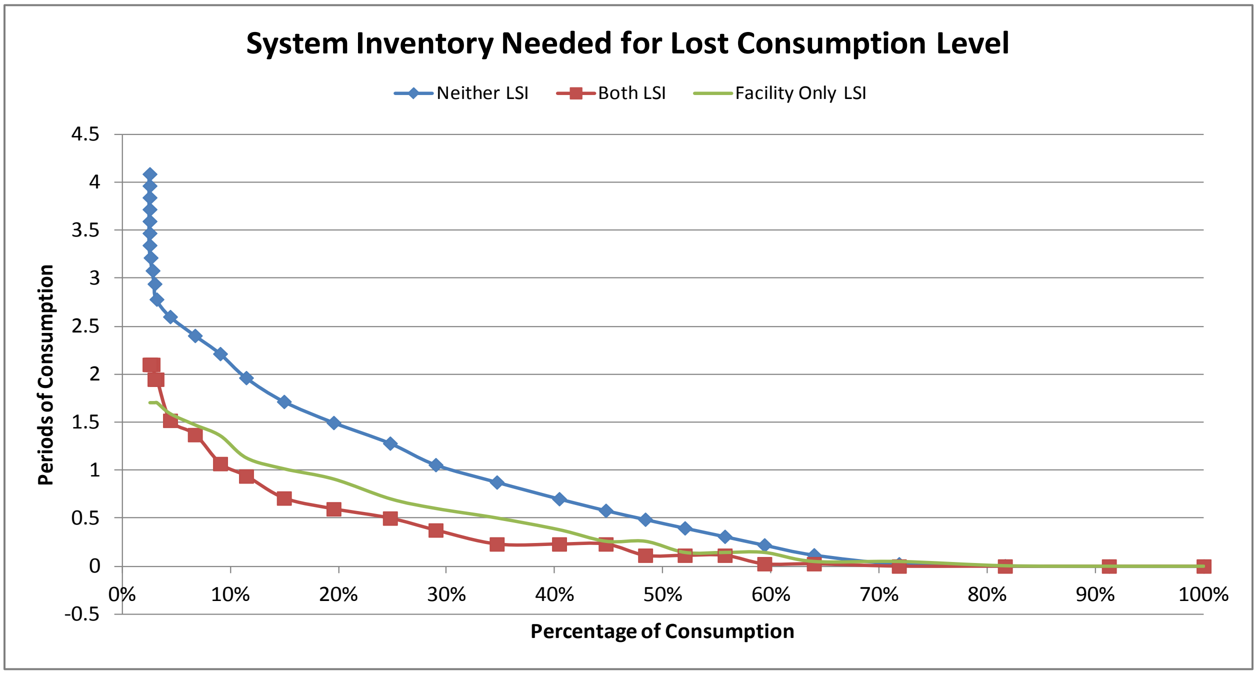

Consumption Is Visible to Upstream Stages

When Central is not on LSI, it uses Facility consumption in the previous period as the AMC. When Central is on LSI, it uses the average of the most recent three periods of consumption as the AMC and multiplies this by the one period LSI to get a forecast for future consumption. As with the setting where consumption is not visible to Central, we see in figure 39 that the scenario where both entities use LSI outperforms the other scenarios, followed by the scenario where the facility uses LSI. Interesting though, at moderate to high levels of stockouts and lost consumption, in the scenario in which only the Facility uses LSI, the Facility performs similarly to scenario in which neither Facility uses LSI.

Figure 39. Inventory Service-Level Trade-off Comparison for Two-Level Supply Chain When Consumption Data Is Shared Upstream

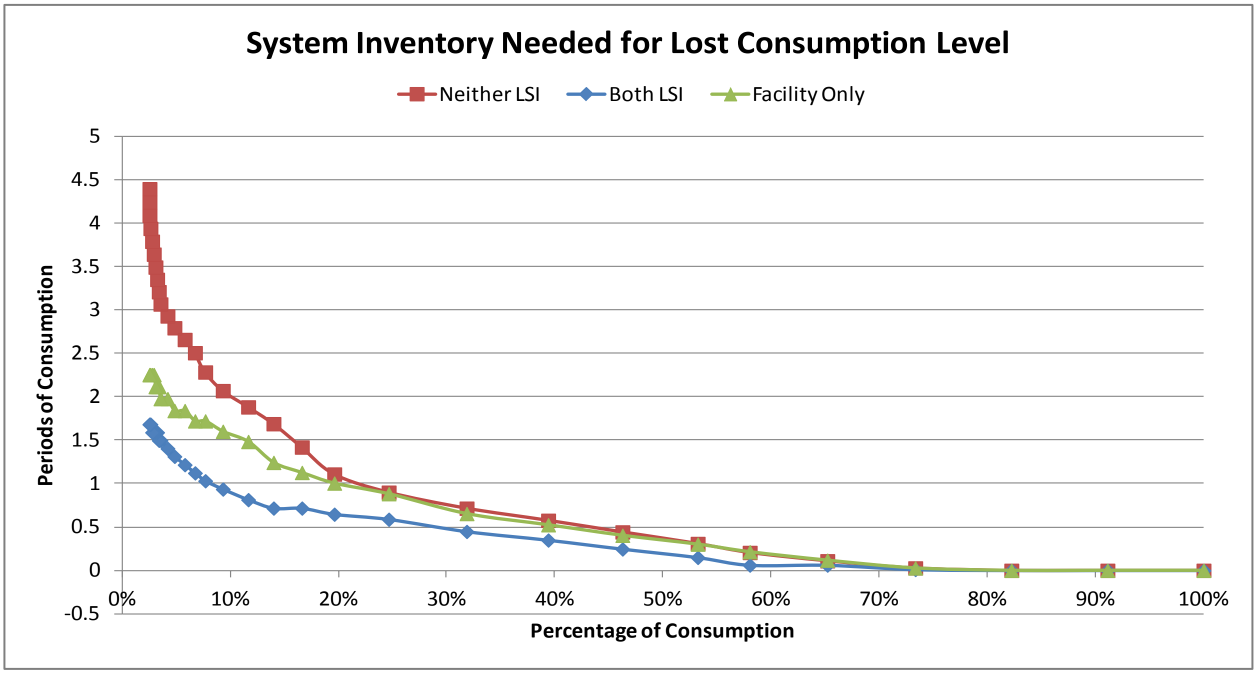

Benefits of Sharing Consumption

Although the previous analysis shows a clear recommendation for both the Facility and Central to use LSI irrespective of whether consumption information is made visible to Central, the analysis can be repurposed to provide insights into when sharing information is recommended. Figure 40 shows rather startling results. In the figure, the three possible scenarios—(1) either both entities use LSI, (2) only the Facility uses LSI, (3) or neither entity uses LSI—are compared individually under the two visibility conditions. Visibility into consumption levels or, more precisely, the use of consumption data for managing inventory is not a clear winner in any scenario and is completely rejected in the setting where only the Facility uses the LSI. For the scenario in which both Central and the Facility use the LSI, visibility is preferred when the level of stockouts and lost consumption is low; otherwise, it is not preferred. This finding is somewhat reversed for the setting where neither Central nor the Facility uses LSI.

Although there are multiple factors contributing to the results here, including the fact that the ordering mechanisms used for Central and the Facility are simple and not necessarily optimized for the particular scenario and/or make best use of information available, a contrast with the clear recommendation of the previous analyses seems clear. Visibility into consumption seems only beneficial if both Central and the Facility employ some level of sophistication in interpreting and forecasting (e.g., using the LSI) and additionally only if inventory levels are high enough to keep stockouts low. The last condition has more recently come to characterize products that receive external support by well-funded donor organizations with such commodities being labeled as in “full supply.”

Figure 40. Differential Benefits of Sharing Consumption Data Upstream

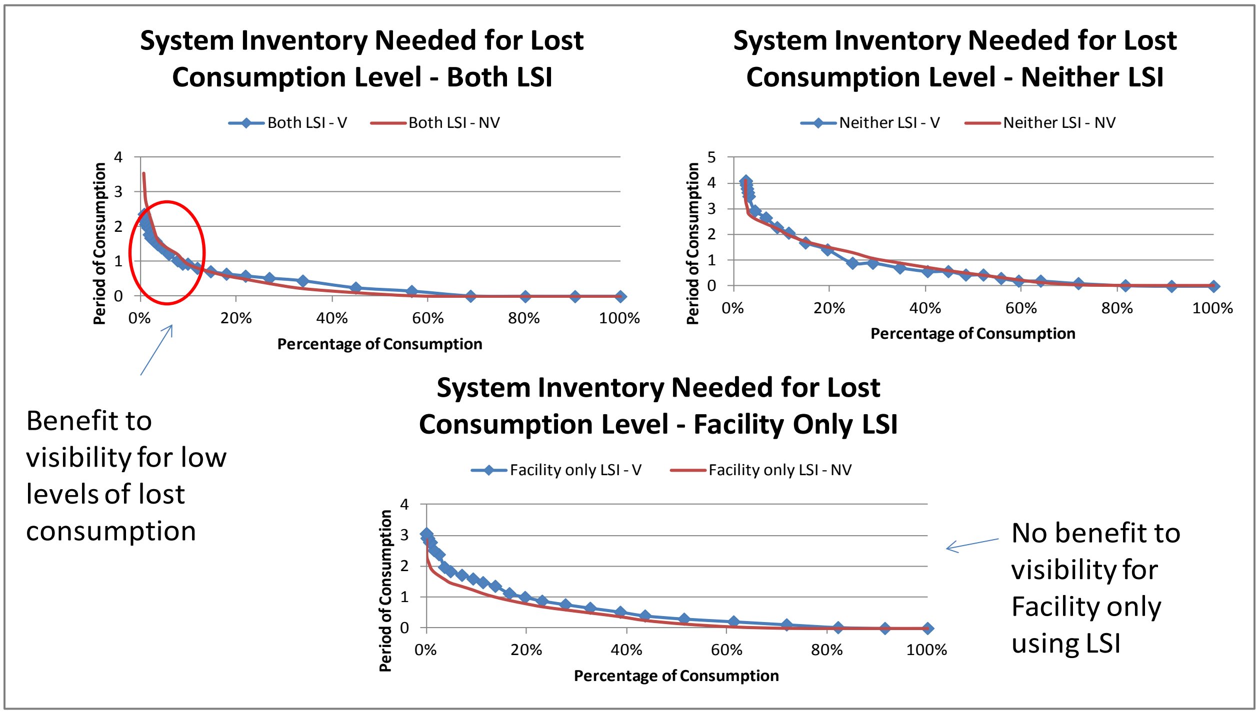

Appendix D

Decoding Gallien et al. 2012

The Gallien, Leung, and Yadav (2012) article entitled "Rationality and Transparency in the Distribution of Essential Drugs in Sub-Saharan Africa: Analysis and Design of an Inventory Control System for Zambia" and the research efforts of its authors have become a benchmark of sorts for conceptualizing the management of seasonal commodities in Zambia. Here we examine their article to have a better understanding of what they have done so that the LSI approach can be understood in relation to that work.

The Gallien et al. article seeks to promote a particular inventory replenishment/control methodology. By replenishment we mean: How does one decide how much inventory to distribute within the supply chain and where? The research activities in the Gallien article that are most relevant for our purposes are first, to test their inventory replenishment/control methodology against the simple AMC rule, especially in the presence of seasonality and supply lead times with variability. They also describe a methodology for historical demand estimation.

Comparison

A look at the objective features of the Gallien inventory replenishment approach (hereafter referred to as the Gallien approach) and the Simple AMC rule not only sets the expectations for what we would see in a comparison of the two approaches, but also identifies some of the additional steps that need to be taken in order to fairly test one approach to the other.

Gallien's approach is primarily optimization based at the individual decision level. The approach involves identifying the costs of getting incorrect the inventory replenishment decision of moving inventory from one level to the next. In this case one can either have too much inventory or too little and we can assign a cost to each of those situations. With these costs, Gallien's approach can search for an inventory replenishment solution that gives the lowest overall cost. It should also be pointed out that the Gallien approach does not have an inherent forecasting approach. It accepts forecasts as input data to their inventory replenishment approach. More precisely, these forecasts are taken as representing the true demand that the inventory replenishment must seek to satisfy. In contrast, the simple AMC rule has no optimization at the individual decision level, although the argument can be made that the top-up level target and associated minimum inventory levels are not arbitrary but have been chosen with some logic in mind. In addition, whereas the Gallien approach has no inherent forecasting approach, the simple AMC does have an inherent forecasting approach that is suited to demand patterns that exhibit strong serial correlation from one period to the next. It can be reasonably used for demand that has no trend; that is, it is generally level over time.

So, to compare, the simple AMC is going to do pretty poorly against Gallien's approach for malaria commodities in Zambia, which we expect to have a seasonal pattern, in part because it does not optimize for every decision but, more importantly, because the inherent forecasting used is geared toward a different kind of product.

Comparing the Two Approaches

To actually compare the two approaches, there is a bit of work that needs to be done to facilitate the comparison. One detail is that of a scoring method for comparing the different approaches. One of the scoring methods utilizes the same costs of getting the inventory decision incorrect that the Gallien approach optimized. Accordingly,

distribution requirement parameters.

The final activity needed to support their comparison is to incorporate forecasting performance. As mentioned earlier, Gallien's approach needs a forecast with which to operate. Instead of describing a forecasting methodology that they use on their realistic demand, the authors approach forecasting in a different way. Without specifying the details of any forecast methodology, the authors simulate what the output would be from a "good" but generic forecast methodology and what the output would be from an "ok" but, again, still generic forecasting methodology. The authors then incorporate these "good" and "ok" forecasts into the simple AMC rule approach and test the effect of forecasting on its performance.

How can one simulate forecasting performance without having a methodology? Here is the essence of how it can be done. Assume we can get some benchmarks of MAPE scores for good and ok forecast methodologies, say 35 percent and 50 percent. In the Gallien article, the authors approach some other academics who have done work in forecasting for similar benchmarks; however, the authors do not report which benchmarks each approach can be compared based on the costs incurred for the inventory replenishment decisions that result from using the approach. The higher the costs incurred, the lower the approach is considered to perform. We use a similar approach in the article for comparing the LSI approach against other forecasting models.

Since the authors did not have access to consumption data across a number of facilities, a second detail is that of having realistic facility consumption data to test the approaches against. The authors generated this consumption data by extracting a basic demand pattern from 18 facilities that they followed closely in their study and extrapolating this basic pattern to about 184 facilities with tweaks based on facility features (e.g., coverage area). Technically, this is not forecasting, although the mechanics of a forecasting approach could be distilled from what they do here. Rather, the authors simulate consumption that, in reality, has already taken place in lieu of collecting this data on consumption.

Other activities needed for supporting the comparison of different approaches include modeling supply lead time performance and inventory flow through the supply chain. Specifically, lead times from the central medical store were modeled as constant while lead time from the districts to health facilities were modeled as being variable. In addition, for potentially cut-off facilities, a probability of a facility being cut off was explicitly incorporated. With respect to inventory flow, for the Simple AMC rule, two approaches were considered: cross-docking at the district and alternatively staging inventory at the district. Gallien's approach would choose between these approaches based on which one was optimal for their particular set of cost and are used.

Next, in table 12, we've provided some examples of forecasts that have a MAPE score of 35 percent. Here we've cheated a bit and made all the actuals equal to 100. Since MAPE ignores the sign of the error—that is, the difference between the actual and the forecast—a forecast methodology that consistently outputs 65 as the forecast, in this case F1, gets the same MAPE score as one that consistently outputs 135 as the forecast, in this case F2. Similarly, a forecast that alternates between 65 and 135 also gets the same MAPE score. To get these MAPE scores we did not use actual forecasting but instead generated a set of numbers that when compared to actuals, give MAPE scores equal to 35 percent. Similarly, if we had put the forecasts at 50 or 150, then we would receive a MAPE score of 50 percent. The approach that Gallien used to generate their set of forecast numbers was probably more sophisticated; for example, there was probably more variation in their forecasts, but the basic idea of creating numbers from actuals that when compared to actual numbers would give a desired MAPE score is consistent with their description of their approach in the article. With these simulated forecasts from a generic "good" and "ok" methodology, the authors can then compare their approach using either forecast methodology with the Simple AMC rule, and additionally, the Simple AMC rule can be tweaked to incorporate both the good and ok methodologies.

Table 12. Examples of Forecasts and MAPE Scores

Facility	Actual	F1	F2	F3
1	100	65	135	65
2	100	65	135	135
3	100	65	135	65
4	100	65	135	135
5	100	65	135	65
MAPE		35%	35%	35%

Reflecting on Comparison Results

The most relevant results from their analysis are the following. The first is that forecasting is important. Every approach generally benefits when going from the generic "ok" forecasting methodology to the generic "good" one. The second is that Gallien's approach with good forecasting performance outperforms everything else. Third is that cross-docking with the Simple AMC rule but incorporating "good" forecasting can be comparable to Gallien's approach coming within 4 percent of its cost-based score.

Reflections on the article raise the following points. The first is that the results of the article don't reject the idea of searching for alternative approaches to Gallien's. The question of forecasting methodology is still an open one, since they don't propose an actual forecasting methodology and their inventory control approach needs one. Second, the Simple AMC rule with generic "good" forecasting in their simulation has comparable results.

Finally, the choice between Gallien's approach and the Simple AMC rule with good forecasting seems to concern choosing an appropriate trade off. On one hand, we have the Gallien approach, which is optimal for each facility, but that potentially requires significant effort to program and maintain in an electronic logistics management information system. It can practically be considered a black box approach, which has implications for acceptance and sustainability. As is common with optimization approaches, although we can rest assured that the approach has searched for the best solution to the problem, we may not have a lot of insight into why the solution looks the way it does and how it will change as parameters change. This lack of insight into what the approach will recommend at any given time can make it difficult to build a supply chain management process around it. On the other hand, we have a Simple AMC rule with (1) good forecasting that gives ok results for any individual facility, (2) that, depending on the forecasting approach, may be easy to program and maintain, and (3) that can be considered less of a black box compared to the Gallien approach. Accordingly, it may be easier to build a supply chain management process around the approach.

However, we must acknowledge the level of speculation in many of these comments concerning the trade off between the Gallien approach and the Simple AMC rule with good forecasting by stating there is still a lot we don't know about each approach when it comes to practice. What strengths, weakness, and opportunities do they confer on the real world systems that would attempt to use them? In particular, what are the challenges to acceptance and of a distinct orientation toward centralization that accompany an approach like Gallien's on system design and implementation? How would the supply chain management processes or infrastructure differ under each approach, if at all? How would the information system management processes or infrastructure differ under each approach, if at all? There are valid arguments for both approaches and, as such, in the interest of knowledge and learning, we are more inclined toward opportunities that allow for both approaches to be tested and evaluated. In particular, if in Zambia, Gallien's approach has generated enough political will to support its implementation, then Zambia can be the test case for their approach while another country like Tanzania can be considered for an alternative approach.

For more information, please visit deliver.jsi.com.

USAID | DELIVER PROJECT
John Snow, Inc.
1616 Fort Myer Drive, 16th Floor
Arlington, VA 22209 USA
Phone: 703-528-7474
Fax: 703-528-7480
Email: askdeliver@jsi.com
Internet: deliver.jsi.com