Published in final edited form as:

Vaccine. 2011 July 26; 29(33): 5512–5518. doi:10.1016/j.vaccine.2011.05.044.

The Optimal Number of Routine Vaccines to Order at Health Clinics in Low or Middle Income Countries

Jayant Rajgopal, PhD³, Diana L. Connor, MPH^{1,2}, Tina-Marie Assi, MPH^{1,2}, Bryan A. Norman, PhD³, Sheng-I Chen, MS³, Rachel R. Bailey, PhD^{1,2}, Adrienne R. Long, BS^{1,2}, Angela R. Wateska, MPH^{1,2}, Kristina M. Bacon, MPH^{1,2}, Shawn T. Brown, PhD⁴, Donald S. Burke, MD², and Bruce Y. Lee, MD, MBA^{1,2}

¹Departments of Medicine and Biomedical Informatics, School of Medicine

²Department of Epidemiology, Graduate School of Public Health

³Department of Industrial Engineering, School of Engineering

⁴Department of Biostatistics, Graduate School of Public Health

⁵Pittsburgh Supercomputing Center, Pittsburgh, PA

Abstract

In a low or middle income country, determining the correct number of routine vaccines to order at a health clinic can be difficult, especially given the variability in the number of patients arriving minimal vaccination days and resource (e.g., information technology and refrigerator space) constraints. We developed a spreadsheet model to determine the potential impact of different ordering policies, basing orders on the arrival rates seen in the previous 1, 3, 6, or 12 sessions, or on long-term historical averages (where these might be available) along with various buffer stock levels (range: 5% to 50%). Experiments varied patient arrival rates (mean range: 1–30 per session), arrival rate distributions (Poisson, Normal, and Uniform) and vaccine vial sizes (range: one-dose to ten-dose vials). It was found that when the number of doses per vial is small and the expected number of patients is low, the ordering policy has a more significant impact on the ability to meet demand. Using data from more prior sessions to determine arrival rates generally equates to a better ability to meet demand, although the marginal benefit is relatively small after more than 6 sessions are averaged. As expected, the addition of more buffer is helpful in obtaining better performance; however, this advantage also has notable diminishing returns. In general, the long-term demand rate, the vial sizes of the vaccines used and the method of determining the patient arrival rate all have an effect on the ability of a clinic to maximize the demand that is met.

INTRODUCTION

In a low or middle income country (as defined by The World Bank) [1], ordering the correct number of vaccines at an administration clinic can be challenging when it is difficult to predict the exact demand at the clinic (i.e., the exact number of patients who will arrive). Funders such as UNICEF have reported difficulties in correctly forecasting yearly demands

© 2011 Elsevier Ltd. All rights reserved.

Corresponding Author: Bruce Y. Lee, MD MBA, Public Health Computational and Operations Research (PHICOR), University of Pittsburgh, 200 Meyran Avenue, Suite 200, Pittsburgh, PA 15213, Phone: (412) 246-6934, FAX: (412) 246-6954, BYL1@pitt.edu.

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

for vaccines [2], and at a national level, a recent epidemiologic study conducted in The Gambia found that while there was not a countrywide vaccine shortage at the time, vaccines remained unavailable to some women and children arriving at a clinic [3]. Since many clinics in low or middle income countries have limited immunization sessions (e.g. every week or every month) and it can be inconvenient for people to return at another time if vaccine is not available, ordering the appropriate number of vaccines is especially important in these settings. Variability in the number of patients arriving at a health center or clinic each day can result in the under- or over-ordering of vaccines. Ordering too many vaccines can result in vaccines going unused and occupying excessive refrigerator or freezer storage space, a major constraint in many low or middle income countries' rural clinics. Ordering too few vaccines can lead to vaccine stock-outs (i.e., running out of vaccines when vaccines are still needed) and missed vaccination opportunities (i.e., patients arriving but not receiving vaccination).

Health clinics, especially those in rural areas, often do not have access to demand forecasting information technology and therefore assume that the past predicts the future, i.e., they measure past arrival rates to generate estimates of future demand [4]. However, in doing this it is important to consider how much historical data (previous session demand, past three sessions, past six or twelve months, etc.) should be used to estimate future vaccine needs. This is especially helpful in situations where monetary resources are limited and buffer stock cannot not be ordered. However if it is economically and logistically feasible, it is also important to consider if one should have some buffer or safety stock, and if so, how much is needed. Buffer stock is an additional number of vaccines above and beyond the estimated number of arrivals at a vaccination session to account for unanticipated surges in vaccine demand (e.g., fluctuations in birth rates, immigration, or vaccination compliance) or losses of vaccine (e.g., through breakage or spoilage). Although organizations such as the Global Alliance for Vaccines and Immunization (GAVI) and World Health Organization (WHO) support the need for buffer stock and promulgate a widely used buffer stock of 25% above the estimated a patient arrival rate [5], there remains a need to further examine the optimal level of buffer stock needed at an immunization location [6]. Moreover, it is unclear how ordering policies should vary with vaccine vial size (e.g., single-dose versus five-dose versus ten-dose vials) and different levels of arrival rates [7].

Therefore, as part of the Vaccine Modeling Initiative (VMI), funded by the Bill and Melinda Gates Foundation, we developed a spreadsheet based simulation model to determine the impact of various vaccine ordering policies on missed vaccination opportunities. In developing and evaluating the ordering policies, the model considers different average arrival rates on a clinic day, the usage of different amounts of historical arrival data for forecasting future vaccine demand, different vial sizes, and different buffer stock levels.

METHODS

The model was developed in Microsoft Excel (Microsoft Corporation, Redmond, WA) and integrated with Visual Basic for Applications (VBA) coding. Figure 1 depicts the structure of the model. Note that at the end of the immunization session (bottom right on figure 1) one of two things could happen: either all incoming patients are vaccinated and there might be some leftover vials which remained unopened, or there are some patients who cannot be treated, which results in missed vaccination opportunities. The objective of the model was to evaluate the impacts of implementing different vaccine ordering policies at an administering health center. We assumed delivery lead times are insignificant in comparison to the time between successive administration sessions, and that decisions related to acquisition costs/budgets have already been made at higher levels in the supply chain. There are two basic inputs to the model:

1. *Vaccine Demand at a Session:* The session arrival rate (i.e. the number of patients requiring immunization at each session) is assumed to be a random variable. Sensitivity analyses explored the effects of using different distributions such as Normal, Uniform and Poisson, and there was little substantive difference in our results. Therefore assume that the number of arrivals follows a Poisson distribution since this distribution is commonly used to describe many arrival processes [8]. However, if one has explicit knowledge that the demand follows some other statistical distribution our overall methodology would still be applicable. We conducted our analysis for various plausible values of the mean of the distribution (λ) ranging from 1 to 30.
2. *Vaccine Vial Size:* We assume that the vaccine always comes in the same vial size, and we conducted our analysis for each of three different vaccine vial sizes (d): 1, 5 and 10 doses per vial.

Given these two inputs, eleven different vaccine ordering policies were evaluated. For ease of exposition we assume that vaccination sessions are held once every month (as happens in countries such as Thailand, Vietnam, and many administration locations in Tunisia), and that a reorder occurs between each session on a regular and on-time schedule [9]. However, the assumption about the timing of a vaccination session is not a crucial one. Depending on the country or organization's operating policies, sessions may occur every two weeks or over some other time interval and our results would still apply. The only requirement of our model is that a vaccine order placed at the end of one immunization session is available at the beginning of the next immunization session. Given that health clinics are typically in close proximity to the next higher node in the supply chain, this is reasonable. With each of our ordering policies we associate an order quantity Q and an order-up-to level R . If the number of vials of vaccine left over at the time that the order is placed is given by v_l , then the policy is to set $Q=R-v_l$, (unless of course, v_l is larger than R , in which case we do not place any order). That is, when we order, the amount ordered is enough to bring the vial inventory level back up to R at the beginning of the next session. This kind of policy (often called a (Q, R) policy) is very typical in practice. To establish the basis for determining the order-up-to level R , the following options were considered:

1. Number of arrivals in the previous session
2. Average of the number of arrivals over the past three sessions
3. Average of the number of arrivals over the past six sessions
4. Average of the number of arrivals over the past twelve sessions
5. Estimate of the mean arrival rate at the clinic (λ)
6. Estimate of the mean with 5% buffer
7. Estimate of the mean with 10% buffer
8. Estimate of the mean with 25% buffer
9. Estimate of the mean with 30% buffer
10. Estimate of the mean with 40% buffer
11. Estimate of the mean with 50% buffer

Note that demand for vaccines generated by clinic arrivals is defined in terms of doses, but all orders are for vials. The first four strategies decide on future actions purely based on what happened in the recent past. The first (and most simplistic) policy is a model that assumes there is no pattern to the arrival rate and uses the previous period's arrival rate as the estimate for the next session. The next three policies are averages of varying length;

shorter periods place more emphasis on more recent data, while longer ones are better at smoothing out random variation. The next seven policies use a long-term estimate of the arrival rate (perhaps based on data, experience or expert opinion), rather than a heuristic approach based on history over some fixed interval of time. The first of these is similar to the heuristic policies in that it does not include any buffer stock, while with the other policies we systematically increase the buffer stock to determine the incremental benefit.

Finally, the metric we use to evaluate our policies is the long-term fraction of demand that is satisfied. In other words, it is the ratio of the total number of individuals vaccinated to the total number of individuals who arrived at the clinic, evaluated over a sufficiently long interval of time.

Description of methodology

For each of the scenarios modeled, we simulated the corresponding ordering policy over a time horizon of 2000 immunization sessions. We start by fixing values for λ (the mean patient arrival rate at a session) and for d (the vial size), and we then generate a random demand stream for the entire set of 2000 sessions, where the number of arrivals at a session is drawn from a Poisson distribution with a mean of λ (e.g. for a mean arrival of 30, the demand ranges from 20 to 40 roughly 95% of the time). Note that with this, the actual number of arrivals at a session will vary, but will be around the set mean. For each session the sequence of events is as follows: patients arrive according to our random demand stream, and as many vials as required are then opened. If the number of vials in the refrigerator is insufficient to meet the needs, we track the amount of demand that is not met; patients whose demands are not met are considered missed vaccination opportunities. On the other hand, if there are enough vials to vaccinate all arriving patients, it is possible that there may be unopened vials that remain in the refrigerator after the session ends and we track this number as well; it is assumed that any (partially used) open vials are discarded (since many vaccines in low and middle income countries are lyophilized and therefore once reconstituted can only be utilized for an additional 6 hours); they cannot be used for the next session. We then place an order for an amount such that we start the next session with R vials; the value of R is determined by the specific policy being evaluated. This process continues for 2000 sessions, after which we compute the amount of demand met across the entire time period and estimate the overall percentage of demand satisfied.

To ensure that all comparisons are based on the same data, the above procedure is then applied to the same demand stream sequentially using each of the eleven policies and the performance of each policy is evaluated. The entire process is repeated ten different times (with a different random number seed for each replicate in order to generate a different demand stream), and the results for each policy are averaged across the ten replicates. We then moved on to the next combination of λ and d , and repeat the same procedure; in particular we study values of λ ranging from 1 through 30 patients, and $d=1, 5$ and 10 doses (for a total of 90 different combinations of the two inputs).

We use the following notation to denote the key variables in our model:

- Number of vaccines vials ordered (Q)
- Number of vaccine doses per vial (d)
- Estimate of the mean arrival rate at a vaccination session (λ)
- Actual number of arriving patients at the vaccination session (P)
- Number of vials opened (v_o)
- Number of missed vaccination opportunities (m)

- Number of vials leftover at the end of the period (v_l)
- The order-up-to level (R)
- The time period (t)

For ease of notation, any variable with a subscript added simply refers to its value corresponding to a specific time period referred to by the subscript.

Our model assumes that the clinic satisfies all patient demand whenever feasible. So, the total number of vials opened for an immunization session is given by:

$$v_o = \text{Minimum}\{R, \text{ROUNDUP}(P/d)\}$$

Thus, for example, if $R=4$, $d=10$ and $P=14$ then $v_o=2$; i.e., when 14 people come to the clinic we open 2 (out of the 4) ten-dose vials in stock.

Missed vaccination opportunities arise only when:

$$R < \text{ROUNDUP}(P/d)$$

The number of patients who go unvaccinated at a session is given by:

$$m = \text{Maximum}\{0, P - v_o d\}$$

Thus, in our previous example, $m=0$. However, if $R=1$ (instead of 4), then $v_o=1$ and $m=4$.

The missed vaccination opportunities are then converted to % of demand met:

$$\% \text{demand met} = ((P - m) / P) * 100\%$$

Given the above calculations, the following ordering policies were evaluated for determining the value of R for any period t :

(1) Previous session's arrival rate: R is set to the minimum number of vials required to meet demand equal to the previous session's demand at that health center:

$$R = \text{ROUNDUP}(P_{t-1} / d)$$

(2) Average of the past three sessions' arrival rates: R is set to the minimum number of vials required to meet demand equal to the average of the previous three sessions' demand at that health center:

$$R = \text{ROUNDUP}(\text{Average}\{P_{t-1} + P_{t-2} + P_{t-3}\} \div d)$$

(3) Average of the past six sessions' arrival rates: Same as (2) but using the average from the previous six sessions:

$$R = \text{ROUNDUP}(\text{Average}\{P_{t-1} + P_{t-2} + \dots + P_{t-6}\} \div d)$$

(4) Average of the past twelve sessions' arrival rates: Same as (2) but using the average from the previous twelve sessions:

$$R = \text{ROUNDUP}(\text{Average}\{P_{t-1} + P_{t-2} + \dots + P_{t-12}\} \div d)$$

(5) Estimate of the mean arrival rate: R is set to the number of vials needed to meet a fixed estimate of the arrival rate at the clinic (λ)

$$R = \text{ROUNDUP}(\lambda \div d)$$

(6) 5% buffer stock: R is set to the number of vials needed to meet the expected demand (λ), plus 5% extra doses.

$$R = \text{ROUNDUP}(\{\lambda \times 1.05\} \div d)$$

(7) 10% buffer stock: Same as (6) but using 1.10 instead of 1.05

(8) 25% buffer stock: Same as (6) but using 1.25 instead of 1.05

(9) 30% buffer stock: Same as (6) but using 1.30 instead of 1.05

(10) 40% buffer stock: Same as (6) but using 1.40 instead of 1.05

(11) 50% buffer stock: Same as (6) but using 1.50 instead of 1.05

Finally, note that the number of vials ordered for period t (at the end of period $t-1$) is always given by:

$$Q = \text{Maximum}\{0, R - v_{l,t-1}\}$$

RESULTS

Overall Impact

Our results show that the best methods for predicting the number of vaccines to order for a vaccination session use information from more than just the most recent immunization session. Averaging the number of arrivals over several sessions provides an increased benefit; however, benefit is not on a linear scale. As more averages are taken over time, less benefit is seen. This indicates that there is a benefit from using enough data from previous sessions to get an estimate of demand which is closer to actual demand, but using too much historical data could result in potential errors when changes occur (such as population changes, population redistribution among geographic regions, or increases in demand). Furthermore, adding a buffer to a reasonably accurate estimate of the patient arrival rate at a session (perhaps from long-term historical data) results in even better performance; however the benefit from additional buffers diminishes as the average demand increases and larger vial sizes are utilized. For ease of exposition, we will refer to policies 1 through 4 of the previous section as *averaging* policies and policies 5 through 11 as *estimate-based* policies.

Varying Patient Arrival Rate

Averaging Policies—

- Diminishing returns (with respect to meeting demand) were seen after adding more than 6 session's worth of data. Single-dose vials had the greatest amount of benefit from using the average of 6 sessions over the average of 3 sessions (1% versus 0.6% and 0.4% for the 5-dose and 10-dose vials respectively). (Figure 2)
- A policy that averages the past 12 sessions' demand always performs the best and simply using the past session's demand is the worst. (Figure 2)
- In general, the performance of all policies improves as the mean arrival rate at a session increases. (Table 1)
- Differences in performance across different vial sizes tend to be smaller when the mean arrival rates are high and more pronounced when the rates are relatively small (under 10 per session). (Table 1)

Estimate-based Policies—

- Having an additional buffer – even a small one of 5% of the estimated arrival rate – yields much better performance in all cases. Policies with larger buffers perform better but with significantly diminishing returns. Going above a buffer of about 25% does not seem to provide much return in terms of performance; even doubling this buffer to 50% buys relatively little in terms of absolute performance. (Figure 3)
- In all vial-size scenarios, it appears that more than a 25% buffer is unnecessary.
- As demand increases, the benefit of additional buffer stock decreases.
- All policies display significant sensitivity with respect to whether the mean arrival rate is a multiple of the vial size. This sensitivity is lower when the buffers are larger (say 25% or more) and also when mean arrival rates are high.
- Once again, larger mean arrival rates yield better performance in general.

General conclusions—Using more data for the *averaging* policies or more buffer for the *estimate-based* policies improves performance. Ideally, 12 sessions of data are desirable if an *averaging* policy is followed but even 6 sessions provides acceptable performance across the board. This indicates that smoothing demand over longer periods of time is preferable to placing more importance on what happened most recently. If reasonably reliable estimates of the long-term mean arrival rate are available and an *estimate-based* policy is used, then a buffer of 25% of this estimate seems to provide very good performance. Additional buffers can improve the performance but of course, this will come at the cost of requiring more storage space. Performance in all cases is better when the mean demand rate at a session is larger. When comparing the two classes of policies, the *estimate-based* policies perform better as long as some buffer is added when ordering.

Varying Vaccine Vial Size

Next we examine how vial size plays a role in the performance of the various policies. For this purpose we group the mean arrival rate into three categories: low, medium and high with 1–10, 11–20 and 21–30 patients per session, respectively. In order to evaluate the various policies, for each range of values, performance (percentage of demand met) is averaged across the ten mean arrival rates for that range. Table 1 presents the results for the *averaging* policies while Table 2 does the same for the *estimate-based* policies. Similar to what was observed in the previous section, across all vial sizes, performance improves as the arrival rate increases and become similar at higher arrival rates.

Averaging Policies (Table 1)—

- Regardless of whether the mean arrival rates are low, medium or high, larger vial sizes always yield better performance. This is because as long as we order the correct amounts, larger vial sizes have a “natural” buffer built in when we round up the estimated requirements.
- Single dose vials are much more sensitive to the mean demand rate than 5 and 10-dose vials, and performance improves by approximately 10% when going from low to high demand rates. On the other hand, performance is very stable for all averaging policies with 5 and 10 dose vials: roughly 94–95% for 5-dose and 97–98% for 10-dose vials.
- Averaging policies that consider more than just the previous session’s demand are always better.
- Using more data for averaging generally provides better performance but with significantly diminishing returns, especially when arrival rates are higher.

Estimate-based Policies (Table 2)—

- Adding a buffer to the mean demand improves performance.
- This improvement is quite sensitive to the vial size: single dose vials benefit significantly from more buffer while the benefits are much lower for multi-dose vials.
- With 5 and 10 dose vials, even a 5% buffer appears to be adequate, while single dose vials need higher buffers at all levels of demand.

DISCUSSION

All methods of averaging (prior 12, 6, or 3 sessions) have an advantage (of up to 7%) over simply using the previous session’s demand, irrespective of vial size and for all mean arrival rates of more than 4 patients. However, the gain in performance from averaging the prior 12 sessions over averaging the previous 6 is less than 1% and this suggests that there is relatively little benefit over averaging the 6 prior sessions. It is important to find the point at which only minimal gain accrues from using additional past data since one has to evaluate the benefit of the additional data versus the cost and/or difficulty of obtaining or maintaining that additional data. Vial size plays an important role, especially when buffer stock percentages are low and when the demand is low. The mean arrival rate also has an impact in all cases; as arrival rates increase, the impact of different policies on performance is less pronounced especially with larger vial sizes. Overall ability to meet demand is best when the arrival rates are higher.

Larger vial sizes result in a higher percentage of demand being met because of the built-in buffer that is inherent in larger vials. This allows for more demand being met even when the actual value is more than estimated. For example, if the estimated arrival rate was fourteen, there would be fourteen single-dose vials, three five-dose vials, or two ten-dose vials ordered. Suppose that in fact, sixteen patients arrive. Then if only one dose vaccines are available, two patients would not be served. However, only one individual would not be served if five-dose vials happened to be in use, and with the ten-dose vials all sixteen individuals could be vaccinated (since a total of twenty doses are available). In fact, the percentage of demand met would remain at 100%, even if six more individuals arrived in this last scenario. Having vials with more doses allows more demand to be met when arrival rates vary.

Larger buffer capacities naturally result in a higher percentage of demand being met. However, it is not feasible to consistently order large quantities of vaccines due to the higher cost and greater potential for wastage from expiration. Additionally, ordering excessive vaccines to ensure that 100% of demand is met will most likely result in a breakdown in the vaccine supply chain due to overburdening both the storage and transport capacities. Our results show that for a location with low predicted arrival rates that uses vaccine vial sizes of 5, any buffer capacities of 5% up to 25% result in approximately the same percentage of demand being met, and this increases only by about 1% by adding a 30% buffer or higher.

If a location uses single dose vial vaccines, effects from ordering policy and effects from additional buffer stock are much greater. This is an important finding, as there are a multitude of single-dose vials approved by the WHO, such as measles-mumps-rubella, diphtheria-tetanus-pertussis-hepatitis B, rabies, and rotavirus, for example [10]. By introducing more single-dose vials into the supply chain, there must be a more generous ordering policy in order to account for the loss of inherent buffer provided by the multi-dose vials. As more countries look towards combating the problem of open vial wastage, they need to consider both their ordering policy as well as the capacities of their supply chain.

The amount of previous sessions' demands required to make "the best guess" relies heavily on the vial size of vaccines utilized at these locations. When single-dose vials are used, there is a greater need to have the best estimate of future demand and a better estimate is provided by averaging the demand of more past sessions. We have demonstrated that when a location is ordering vaccines for the next session, they need to consider what types of vaccine vial sizes they are using, the range of demand (low, middle, or high), and using these indicators they can determine the number of past sessions data to use to best estimate the upcoming demand. It is important to keep in mind that using the most data (12 in our model) will not always necessarily be the best. In particular, if there are factors that result in significant changes in the average demand level, using too much old data could result in inaccuracies. In this situation, less past demand data should be utilized thereby minimizing potential impact factors that have occurred between the past and upcoming sessions. In general, we found that when the vial size is small and demand is low, it is more important to use additional data to predict future demand.

The results of our model highlight the substantial impact that vaccine ordering policy has on the percentage of demand that one is able to meet. The World Health Organization states that the most common method of ordering vaccines on a yearly basis is by estimating the amount required based on the amount of vaccine used in the previous year, yet it does not take into account factors such as vial sizes and variations in patient arrival rates[11]. The information from our model can help guide vaccine purchasers at the lowest levels of the supply chain (e.g., vaccine administration sites) when deciding on the amount of vaccines to order, depending upon the vaccine vial size used, budgets and the availability of buffer capacity at a particular administration site.

LIMITATIONS

Every computational model is only a representation of real life [12–13]. Our model does not consider immunization programs that use a mixture of vial sizes at a given site. We also assumed that the supply chain would contain enough capacity to accommodate all ordering decisions considered in our experimental scenarios, and would fill orders in a timely manner. We did not assign value to a missed vaccination. This study is based on determining the best policy for a vaccine administration level to order supply from the upper level of the supply chain. Future analyses will need to be performed to determine optimal ordering

policies at the higher levels, including the national level, for a yearly basis, as many funding agencies request orders from a country for the next year's vaccine needs.

CONCLUSIONS

The results of our model highlight the substantial impact vaccine ordering policy has on the resulting percentage of demand that is met and missed vaccination opportunities, especially as vaccine vial sizes and the number of patients arriving for vaccination vary. Vaccine ordering policies can have a great impact on the availability of vaccine for all patients, and need to be assessed in accordance with the parameters of that specific region's vaccine usage in order to determine which policy offers the highest level of vaccine availability. Vaccination locations that have relatively small numbers of patients to be vaccinated and that use vials with a smaller numbers of doses are impacted the most.

Acknowledgments

This study was supported by the Vaccine Modeling Initiative (VMI), funded by the Bill and Melinda Gates Foundation and the National Institute of General Medical Sciences Models of Infectious Disease Agent Study (MIDAS) grant 5U54GM088491-02. The funders had no role in the design and conduct of the study; collection, management, analysis, and interpretation of the data; and preparation, review, or approval of the manuscript.

REFERENCES

1. The World Bank. [cited 2011 March 14] How we Classify Countries. 2011. Available from: <http://data.worldbank.org/about/country-classifications>
2. Institute of Medicine. The Children's Vaccine Initiative: Continuing Activities. Washinton, D.C.: National Academy Press; 1995.
3. [cited 2011 March] Ariva Appui au Renforcement de l'Indépendance Vaccinale en Afrique. Missed Opportunity Study for Vaccination in health Facilities in The Gambia. 2003. Available from: <http://www.ariva.bf/etude/Gambie/STUDY%20ON%20Missed%20opportunity%20study%20report.pdf>
4. World Health Organization. [cited 2010 November 29] Immunization service delivery and accelerated disease control: Vaccine forecast and needs estimation. 2010. Available from: http://www.who.int/immunization_delivery/systems_policy/logistics/en/index2.html
5. World Health Organization. Geneva: 2005. Vaccine Introduction Guidelines: adding a vaccine to a national immunational programme: decision and implementation.
6. GAVI Alliance. [cited 2010 December 3] Glossary. 2010. Available from: http://www.gavialliance.org/media_centre/glossary/index.php
7. Lee BY, Norman BA, Assi TM, Chen SI, Bailey RR, Rajgopal J, et al. Single versus multi-dose vaccine vials: an economic computational model. *Vaccine*. 2010 Jul 19; 28(32):5292–5300. [PubMed: 20566395]
8. Ross, SM. Stochastic Processes. 2nd ed. Wiley; 1996.
9. Nettleton, S. [cited 2011 March 15] Monthly immunization days save live and build babies' health. At a glance: Viet Nam. 2006. Available from: http://www.unicef.org/infobycountry/vietnam_30867.html
10. World Health Organization. [cited 2011 January] Vaccine Management and Logistics: Vaccine Volume Calculator. 2009. Available from: http://www.who.int/immunization_delivery/systems_policy/logistics/en/index4.html
11. World Health Organization. [cited 2010 November 29] Demand Forecasting for Routine Immunization. 2010. Available from: http://www.who.int/immunization_supply/introduction/forecasting/en/index.html
12. Lee BY. Digital decision making: computer models and antibiotic prescribing in the twenty-first century. *Clin Infect Dis*. 2008 Apr 15; 46(8):1139–1141. [PubMed: 18444847]

13. Lee BY, Biggerstaff BJ. Screening the United States blood supply for West Nile Virus: a question of blood, dollars, and sense. PLoS Med. 2006 Feb 3.2:e99. [PubMed: 16420099]

Vaccine Ordering Policy **Immunization Session**

Figure 1.
Schematic for various ordering policies

Figure 2.
Effect of various averaging procedures on the percentage of met demand

Figure 3.
Change in percentage of met demand by varying buffer capacities

Table 1Performance of *averaging* methods by vial size and mean arrival rate

	Single-dose Vial	Five-dose Vial	Ten-dose Vial
Low mean arrival rates (1–10)			
Previous Session	81%	93%	98%
Prior 3 session average	83%	94%	98%
Prior 6 session average	84%	95%	98%
Prior 12 session average	85%	95%	98%
Medium mean arrival rates (11–20)			
Previous Session	86%	91%	94%
Prior 3 session average	89%	94%	96%
Prior 6 session average	90%	94%	97%
Prior 12 session average	91%	95%	97%
High mean arrival rates (21–30)			
Previous Session	89%	92%	95%
Prior 3 session average	91%	94%	97%
Prior 6 session average	92%	95%	97%
Prior 12 session average	93%	95%	97%

Table 2Performance of *estimate-based* methods by vial size and mean arrival rate

	Single-dose Vial	Five-dose Vial	Ten-dose Vial
Low mean arrival rates (1–10)			
Estimate of mean arrival	81%	94%	97%
5% buffer	90%	97%	98%
10% buffer	90%	97%	98%
15% buffer	92%	97%	99%
25% buffer	93%	97%	99%
30% buffer	94%	99%	100%
40% buffer	95%	99%	100%
50% buffer	96%	99%	100%
Medium mean arrival rates (11–20)			
Estimate of mean arrival	90%	94%	97%
5% buffer	93%	96%	98%
10% buffer	95%	97%	99%
15% buffer	96%	98%	99%
25% buffer	98%	99%	99%
30% buffer	99%	99%	100%
40% buffer	99%	100%	100%
50% buffer	100%	100%	100%
High mean arrival rates (21–30)			
Estimate of mean arrival	92%	95%	97%
5% buffer	95%	97%	98%
10% buffer	97%	98%	99%
15% buffer	98%	99%	99%
25% buffer	99%	99%	100%
30% buffer	99%	100%	100%
40% buffer	100%	100%	100%
50% buffer	100%	100%	100%