CHALLENGES IN MODELING DEMAND FOR INVENTORY OPTIMIZATION OF SLOW-MOVING ITEMS

Frank Grange

Andersen Consulting, LLP
1225 17th St., Suite 3200
Denver, Colorado 80202-5532 U.S.A.

ABSTRACT

This paper presents the concept of inventory optimization and the role of demand distributions. Slow-moving item demand receives special emphasis. Simulations of several popular discrete distributions illustrate the difficulties of probability modeling with the quantity and quality of demand history typically available. Results of experiments with a state-of-the-art probability modeling tool (ExpertFit™) highlight the practical difficulties of data-fitting. Finally, inventory optimization over a simulated data set with the “wrong” assumed demand distribution suggests a business case for accurately identifying demand distributions.

1 INTRODUCTION

Inventory optimization is an emerging practical approach to balancing investment and service-level goals over a very large assortment of stock-keeping units (SKUs) typified by automobile dealers, industrial equipment distributors, or telecommunications network service centers. In contrast to traditional “one-at-a-time,” marginal stock level setting, inventory optimization simultaneously determines *all* SKU stock levels to fulfill total service and investment constraints or objectives. In many published applications, inventory optimization has offered attractive, comprehensible management controls that a business owner (for example, an automobile dealer principal), financial manager, or inventory manager can manipulate.

However, in contrast to standard marginal stock level setting, inventory optimization requires a probability distribution of SKU demand. Familiar classical methods emphasize SKU forecasts. As a result, in work on behalf of a major American automobile manufacturer, the author has encountered some concerns among interested inventory management practitioners about how to model the demand distribution of slow-moving items, and about what risks might result from inaccurate or inappropriate probability modeling.

2 INVENTORY OPTIMIZATION

In the pursuit of improved but affordable customer service levels, many organizations have adopted service parts inventory optimization. Traditional optimal inventory analysis establishes order quantity, stock level guidance (maximum stock), and reorder point on a part-by-part basis, for example, with Economic Order Quantity or (s,S) models (Azoury and Miller, 1984; Ehrhardt and Mosier, 1984). In traditional analysis, service level or “part fill rate” is the probability that a stock level will cover demand. In contrast, “inventory optimization” defines total fill rate as a weighted average of the individual part fill rates, for example,

$$\text{Total Fill Rate} = \sum_j \left(E[D_j] \text{Prob}\{D_j \leq Q_j\} / \sum_k E[D_k] \right) \quad (1)$$

where D_j = demand rate (say, weekly) for stock-keeping unit (SKU) j ; Q_j = stock guidance (“order-up-to” quantity) for SKU j ; and $\text{Prob}\{D_j \leq Q_j\}$ = individual service level for SKU j .

Here the definition of total fill rate weights individual part fill rates by their expected demands (Hopp and Spearman, 1995). Other weighting schemes are certainly conceivable; for example, expected total cost of each SKU. The number of SKUs in a spare parts inventory may vary from a few hundred to several hundred thousand. Following (1), a simple inventory optimization formulation is

$$\text{Minimize : Total Inventory Cost} = \sum_j C_j Q_j$$

where T = the target fill rate for the entire inventory. Practical situations may expand this formulation to include additional constraints on fill rate targets for specific subsets of items, space constraints on certain item groups (for example, automobile body parts or windshield glass), and budgets on subsets of items.

$$\text{subject to } \frac{\sum_j (E[D_j] \text{Prob}\{D_j \leq Q_j\})}{\sum_k (E[D_k])} \geq T$$

The advantages of optimizing service parts inventory over total fill rate include:

- Controlling inventory to a single total fill rate goal, and thereby avoiding the need to specify a cost-effective service level for every SKU
- Relating a simple performance goal (total fill rate) to total inventory investment
- Considering trade-offs between service levels of individual part stocking levels

The principal disadvantage of the total fill rate approach is incurring low service levels on some SKUs. For typically large parts assortments, this disadvantage is acceptable in exchange for a manageable inventory control process. Published applications of the total fill rate approach to inventory optimization include Cohen, et al. (1990); Harris (1997); and Hopp, et al. (1997).

In real-world situations, slow-moving items can figure prominently in evaluating inventory effectiveness. For example, automotive dealers often stock from 5,000 to 20,000 spare parts. Stock order replenishment from a facing warehouse normally occurs weekly. Of all the dealer's parts, 90 percent or even more may be slow-moving parts, that is, those with 20 or more weeks of very low or zero demand in a year. To achieve a high total fill rate (for example, $\geq 90\%$), a dealer must stock a significant inventory of slow-moving parts, that is, fast-moving part inventories alone cannot achieve a dealer's high total fill rate target. Thus, obtaining a high service level on slow-moving parts prompts real curiosity, if not outright concern, about the right tail of the probability distribution of a slow-moving part's demand.

3 PURCHASER BEHAVIOR AND DISTRIBUTIONS OF DEMAND

Purchaser behaviors affect the appropriate choice of probability distribution. Miller (1995) discusses "macro models" of a real-world customer population that justify a theoretical distribution:

"Micro models deduce theoretical distributions directly from assumptions made about the 'behavior' of the underlying things that are being counted in the distributions...A macro model makes *hopefully plausible assumptions* about the overall population of the things being counted and *deduces the theoretical distribution that would result if the assumptions are true*. A macro model is not based on the 'behavior' of the individual things being counted." (emphasis added)

Factors of slow-moving item purchase behaviors include the following:

- On each purchase occasion, do buyers demand the item in quantities of just one? Or more than one?
- If the quantity per purchase occasion can be more than one, are there one or more modes (for example, shock absorbers)?
- Do some buyers in the population have a significantly different demand rate (for example, a national repair shop chain buying wholesale original equipment parts from a dealer versus a "shade-tree mechanic")?
- Can price promotions affect consumers' demand rates? For most slow-moving items, manufacturers or suppliers are not prepared to accommodate a demand burst, and are thus unlikely to offer promotions. Some style goods, such as manufacturer-branded appearance items, may present exceptions.
- Does seasonality create a nonstationary demand distribution? Demand for certain slow-moving repair parts will certainly exhibit seasonality (for example, for lawnmowers or snowmobiles), but many slow-moving items will present relatively stationary demand distributions.
- Does manufacturer or supplier phase-in/phase-out of SKUs create a nonstationary demand distribution?
- Do manufacturer or supplier engineering change orders create a nonstationary demand distribution?
- Does manufacturer vehicle or product launch or build-out create nonstationary demand distribution?

Numerous authors have concluded that the negative binomial distribution is generally most suitable for modeling the probability of demand for slow-moving items. The Santa Clara University Retail Workbench yielded many data sets fit well with the negative binomial (Agrawal and Smith, 1996). Jacobs and Wagner (1989) found the negative binomial distribution effective for modeling the demand distribution of 21 U. S. Air Force data sets. While Murthi, et al. (1993), discovered conditions under which alternative mixture models were superior, those authors found the negative binomial distribution surprisingly robust with respect to alternative forms of heterogeneity among consumers' per-purchase demand rates.

A more or less plausible macro model that results in the negative binomial is as follows:

- Consumers in the population may have different and unknown positive demand rates λ .
- The distribution of λ should be plausible in that, for different values of its distribution parameters, the distribution of λ has reasonable shapes, namely J- and bell-shaped forms. The gamma distribution is suitable.
- The number of items demanded on a purchase occasion is Poisson distributed with a mean λ that is thus gamma-distributed. Since the negative binomial distribution is a mixture of the Poisson and gamma

distributions, this macro model of slow-moving item demand yields a negative binomial distribution.

Miller (1995) poses similarly elegant, straightforward, plausible assumptions about demand for a slow-moving item to yield any of beta-binomial or three unnamed (or, at least, unrecognizable) distributions apparently in the hypergeometric family. He comments:

“The major point of these various examples is this.

Given any particular area of interest it is quite easy to generate a profusion of models, each of which will lead to a possible theoretical distribution that may fit observed distributions from the area. Each of these theoretical distributions will have any number of possible underlying models, not just the particular model that we used in deducing the distribution. We have, then, an embarrassment of riches if we take this approach to understanding frequency distribution.”

In the context of inventory optimization, a retailer or distributor would face a daunting market analysis burden to create appropriate macro models for each of an enormous number of SKUs. As the following section illustrates, the available SKU demand data probably will not encode complete answers to the purchase behavior issues enumerated above.

4 AUTOMATED FITTING OF PROBABILITY DISTRIBUTIONS OF DEMAND

In an application such as inventory optimization, the large number of SKUs typically involved would require an automated facility to identify and fit appropriate probability distributions. To illustrate the challenges of fitting these kinds of data, I used ExpertFit (Averill Law and Associates) to fit slow-moving demand data simulated from known distributions. Table 1 illustrates samples from four distributions (Poisson, binomial, negative binomial, and beta-binomial) that various academics and practitioners have recommended for slow-moving items. These samples are representative of 52 weeks of demand records. Note that each item displays one or more weeks of zero demand, and each simulated demand has a small sample mean (approximately 2) and low maximum. Table 2 presents the true parameters and sample statistics for each simulated item demand. The simulation assumes that each item demand is completely stationary and thus lacks both seasonality or trend.

Table 1: Simulated Weekly Slow-Moving Demand Data

Week	1	2	3	4	5	6	7	8	9	10
Item A	2	0	3	2	2	2	1	0	0	1
Item B	0	1	1	3	2	1	3	2	1	2
Item C	1	6	1	3	4	1	0	0	2	1
Item D	0	1	14	0	1	0	2	1	0	5
Week	11	12	13	14	15	16	17	18	19	20
Item A	0	3	3	1	1	3	1	3	0	1
Item B	1	1	2	0	2	1	2	1	2	0
Item C	5	5	3	3	3	6	1	1	1	2
Item D	3	3	3	0	3	2	7	1	7	1
Week	21	22	23	24	25	26	27	28	29	30
Item A	2	1	1	5	4	1	1	1	2	1
Item B	3	2	2	1	2	3	3	2	1	2
Item C	1	2	3	2	5	0	0	2	3	8
Item D	4	3	0	0	0	7	1	0	0	0
Week	31	32	33	34	35	36	37	38	39	40
Item A	3	6	0	2	0	1	1	4	1	2
Item B	3	1	3	2	1	2	0	3	2	3
Item C	1	1	0	6	1	1	2	3	4	3
Item D	3	4	0	2	0	11	1	1	1	1
Week	41	42	43	44	45	46	47	48	49	50
Item A	3	6	1	2	1	2	3	4	3	7
Item B	1	1	2	1	1	2	3	3	2	1
Item C	1	0	0	1	5	1	1	3	0	5
Item D	1	2	5	0	2	14	0	4	5	10
Week	51	52								
Item A	3	1								
Item B	1	2								
Item C	0	0								
Item D	1	0								

The characteristics of these simulated demand data would certainly argue strongly for modeling with a discrete probability distribution that is truncated at the left. Results of the ExpertFit analysis appear in Table 3, and are encouraging for the following reasons:

- Correct detection of both binomial and negative binomial distributions
- χ^2 goodness-of-fit rejections of distributions whose Lexis ratios (variance-to-mean) are contrary to the binomial and negative binomial
- Correct resolution of geometric versus negative binomial distribution for Item C (the geometric distribution is a special case of the negative binomial distribution, and the true negative binomial distribution was relatively “close” to a geometric distribution)
- Rejection of the Poisson distribution for modeling the Item D beta-binomial data

Since ExpertFit does not include the beta-binomial (or any other three-parameter discrete distributions), its failure on Item D is hardly surprising or unreasonable.

Often retail business systems contain little or no capacity to store and analyze SKU demand history, so that having 52 weeks of data could be unlikely. To support financial reporting and planning, some systems will produce and archive monthly demand data by SKU. Table 4 provides 13 monthly (4-week) aggregations of the data from Table 1. Fitting these data with ExpertFit yields the results shown in Table 5. In three of the four simulated

demands, ExpertFit recommends the discrete uniform distribution; in the case of the Item D beta-binomial data, the tool favors the negative binomial distribution (whose Lexis ratio exceeds one, in contrast to the underlying beta-binomial distribution). As expected, demand history aggregation that would be typical of retail business systems confounds state-of-the-art analytical efforts to identify underlying probability distributions from the demand pattern.

Table 2: Distributions Used to Generate the Simulated Weekly Slow-Moving Demand Data

	Item A	Item B	Item C	Item D
Distribution	Poisson	Binomial	Negative Binomial	Beta-Binomial (Hypergeometric)
True Mean	2	2	2	2
True Variance	2	0.66667	3.3333	0.71429
Lexis Ratio (Var/Mean)	1	0.33333	1.6667	0.35714
Sample Mean	2	1.7115	2.1923	2.5192
Sample Variance	2.6275	0.7975	3.9231	11.588
Sample Max	7	3	8	14
Sample Lexis Ratio	1.3137	0.466	1.7895	4.6000

Table 3: Results of ExpertFit Analyses of Simulated Weekly Slow-Moving Demand Data

	Item A	Item B	Item C	Item D
1 st Choice	Negative Binomial	Binomial	Negative Binomial	Geometric
2 nd Choice	Poisson	<i>Poisson</i>	<i>Geometric</i>	Negative Binomial
3 rd Choice	<i>Geometric</i>	<i>Discrete Uniform</i>	<i>Poisson</i>	<i>Poisson</i>
4 th Choice	<i>Discrete Uniform</i>	<i>Geometric</i>	<i>Discrete Uniform</i>	<i>Discrete Uniform</i>
True Distribution	Poisson	Binomial	Negative Binomial	Beta-Binomial

Note: Italics indicate distributions rejected in equal-width χ^2 test

Table 4: Simulated Aggregated Monthly Slow-Moving Demand Data

Month	Item A	Item B	Item C	Item D
1	7	5	11	15
2	5	8	5	4
3	4	5	13	11
4	8	5	15	8
5	5	5	5	10
6	9	8	8	7
7	7	10	7	8
8	12	7	13	7
9	3	8	8	13
10	8	8	12	4
11	12	5	2	8
12	10	9	10	20
13	14	6	5	16

Table 5: Results of ExpertFit Analyses of Simulated Monthly Slow-Moving Demand Data

	Item A	Item B	Item C	Item D
1 st Choice	Discrete Uniform	Discrete Uniform	Discrete Uniform	Negative Binomial
2 nd Choice	Negative Binomial	Binomial	Negative Binomial	<i>Poisson</i>
3 rd Choice	Poisson	Poisson	Poisson	Discrete Uniform
4 th Choice	<i>Geometric</i>	<i>Geometric</i>	<i>Geometric</i>	Geometric
5 th Choice	<i>Logarithmic series</i>	<i>Logarithmic series</i>	<i>Logarithmic series</i>	<i>Logarithmic series</i>
True Distribution	Poisson	Binomial	Negative Binomial	Beta-Binomial

Note: Italics indicate distributions rejected in equal-width χ^2 test

The failure to capture information about lost sales of slow-moving items creates an additional estimation challenge. In the case of a service-parts operation, many business systems provide a lost sales recording capability, but poorly trained or overly busy staff may fail to record the lost sales information. In a self-service environment, such as a department store, many lost sales go unobserved. Failure to record lost sales creates a right-censored population; Agrawal and Smith (1996) offer an estimation method for negative binomially distributed demand. Nahmias (1994) provides a method for normally distributed demand with unobserved lost sales.

5 COST OF MISIDENTIFICATION OF SKU DEMAND DISTRIBUTION

The cost of misidentifying SKU demand distributions takes one of the following forms:

- If the distribution's right tail overestimates the cumulative probability, the inventory optimization could predict an unrealistically high total fill rate, and thus underinvest in stock. Some comments in various articles by Harris (for example, 1997) may indicate such results.
- If the distribution's right tail underestimates the cumulative probability, the inventory optimization could predict a lower total fill rate than should be expected, and thus cause overinvestment in inventory.

The exact consequence to inventory optimization of misidentifying SKU demand distributions is unpredictable; the results above are correct for setting the service level of an individual SKU as, for example, in Jacobs and Wagner (1989) or Silver (1991). However, a multi-SKU inventory optimization will "compensate" somewhat for misidentification by reallocating investment relative to the costs and expected demands of all the SKUs.

Table 6 illustrates the effect of assuming all SKU demands are Poisson-distributed when they are, in fact, negative-binomially distributed. Abundant empirical studies notwithstanding, many authors (for example, Murthi, et al. (1993)), as well as some commercial demand forecasting software, suggest the Poisson as a suitable demand distribution for slow-moving items. The costs and negative binomial parameters for the 20 SKUs in Table 6 are arbitrary, and the inventory optimization used the Poisson probabilities to compute expected fill rate. The Solver™ add-in for Microsoft Excel™ computed the solution in the column labeled "Stock Level."

Table 6 shows that, in general, the Poisson overestimates the true negative binomial cumulative probabilities of demand, and thus overstates the expected fill rate by approximately 5 percent for this particular contrived example. This expected fill rate misestimation from demand distribution misidentification is severe enough to warrant concern about the rigor of probability modeling an inventory optimizer would enforce.

The numerical error revealed in this small sample is of only qualitative value. Inventory optimization involves solving a nonlinear program over a set of integer variables (the stock levels). While the Excel Solver is generally effective and accurate for linear and nonlinear optimization over continuous variables, its algorithm is somewhat questionable over discrete variables (for example, receiving a bad starting solution to the problem in Table 6, Solver will fail to find an optimal solution; what constitutes a "bad starting solution" is not obvious). Experiments with several starting solutions yielded the solution in Table 6, so that it is probably an optimal solution, or one very close to a true optimum.

Table 6: Inventory Optimization Assuming Poisson-Distributed Demands When "True" Demand Distributions Are Negative Binomial

Item	Unit Cost	"True" Negative Binomial Distribution (NBD)					Stock Level	Cum Probability of Demand	
		s	p	Mean	Variance	Lexis Ratio		Poisson	NBD
1	\$1.27	3	0.2083	11.402	54.741	4.801	19	0.986803	0.866377
2	\$6.73	4	0.3280	8.195	24.985	3.049	17	0.997943	0.9488226
3	\$2.10	7	0.6308	4.097	6.496	1.585	19	1	0.9999617
4	\$15.44	2	0.5948	1.363	2.291	1.681	13	1	0.99997
5	\$7.91	1	0.9759	0.025	0.025	1.025	16	1	1
6	\$17.05	4	0.8950	0.469	0.524	1.117	12	1	1
7	\$13.37	3	0.2434	9.324	38.304	4.108	14	0.94707	0.8208325
8	\$4.56	1	0.8142	0.228	0.280	1.228	18	1	1
9	\$7.73	5	0.8794	0.686	0.780	1.137	16	1	1
10	\$12.30	2	0.3000	4.667	15.556	3.333	14	0.999891	0.9738884
11	\$14.23	4	0.4320	5.259	12.173	2.315	13	0.998881	0.9738874
12	\$1.49	6	0.3761	9.955	26.473	2.659	19	0.99671	0.9506677
13	\$12.26	1	0.1114	7.976	71.589	8.976	14	0.983149	0.8299737
14	\$16.45	3	0.6862	1.372	2.000	1.457	14	1	0.9999981
15	\$14.62	2	0.2178	7.181	32.962	4.590	13	0.984419	0.8701254
16	\$15.38	5	0.4828	5.355	11.091	2.071	13	0.998678	0.9780117
17	\$3.96	3	0.9947	0.016	0.016	1.005	18	1	1
18	\$13.36	4	0.6400	2.250	3.516	1.563	14	1	0.9999474
19	\$18.88	8	0.3722	13.491	36.242	2.686	12	0.410273	0.4818636
20	\$5.54	2	0.8501	0.353	0.415	1.176	17	1	1
Target Fill Rate								90%	
Expected Fill Rate								90.5%	85.3%

6 CONCLUSIONS

Inventory managers faced with high service-level requirements and many SKUs appreciate the simplicity of inventory optimization, as well as the explicit control it offers over total investment and total fill rate. In the author's experience, practitioners are generally unfamiliar with the choices and implications of various discrete probability distributions. The issues and examples presented in this paper suggest that a) identification of distributions from typically available demand history can be extremely difficult, and b) misapplication of a demand distribution will yield unsatisfactory inventory optimization results. The facets of slow-moving item purchase behaviors suggest a number of fertile research directions with solid practical value to establish an inventory optimization capability with predictable performance.

ACKNOWLEDGMENTS

Thanks to Mark Grabau of Andersen Consulting for his helpful suggestions on this paper. Thanks also to Greg Clay for his previous collaboration on simulation of automotive spare parts inventory management.

ExpertFit is a trademark of Averill M. Law and Associates, P. O. Box 40996, Tucson, Arizona, USA. Solver and Excel are trademarks of the Microsoft Corporation, Redmond, Washington, USA.

REFERENCES

- Agrawal, Narendra, and Stephen A. Smith, 1996, "Estimating Negative Binomial Demand for Retail Inventory Management with Unobservable Lost Sales," *Naval Research Logistics*, vol. 43, pp. 839-861.
- Azoury, Katy S., and Bruce L. Miller, 1984, "A Comparison of the Optimal Ordering Levels of Bayesian and Non-Bayesian Inventory Models," *Management Science*, vol. 30, no. 8, p. 993-1003.

- Cohen, Morris, Pasumarti V. Kamesam, Paul Kleindorfer, Hau Lee and Armen Tekerian, 1990, "Optimizer: IBM's Multi-Echelon Inventory System for Managing Service Logistics," *Interfaces*, vol. 20, no. 1 (Jan-Feb), p. 65-82.
- Ehrhardt, Richard, and Charles Mosier, 1984, "A Revision of the Power Approximation for Computing (s,S) Policies," *Management Science*, vol. 30, no. 5, p. 618-622.
- Harris, Terry, 1997, "Optimized Inventory Management," *Production and Inventory Management Journal*, vol. 38, no. 1 (1st quarter), p. 22-25.
- Hopp, Wallace J., and Mark L. Spearman, 1995, *Factory Physics: Foundations of Manufacturing Management*, Richard D. Irwin, Chicago.
- Hopp, Wallace J., Mark L. Spearman and Rachel Q. Zhang, 1997, "Easily Implementable Inventory Control Policies," *Operations Research*, vol. 45, no. 3, p. 327-340.
- Jacobs, Raymond A., and Harvey M. Wagner, 1989, "Reducing Inventory Systems Costs by Using Robust Demand Estimators," *Management Science*, vol. 35, no. 7, p. 771-787.
- Miller, David W., 1995, *Fitting Frequency Distributions to Data: Philosophy and Practice*, Book Resource, New York.
- Murthi, B. P. S., Kannan Srinivasan and Pandu R. Tadikamalla, 1993, "Robustness of the Negative Binomial Distribution to Model Purchase Incidence: A Simulation Study," *Communications in Statistics—Simulation and Computation*, vol. 22, no. 2, pp. 399-420.
- Nahmias, Steven, 1994, "Demand Estimation in Lost Sales Inventory Systems," *Naval Research Logistics*, vol. 41, pp. 739-757.
- Silver, Edward A., 1991, "A Graphical Implementation Aid for Inventory Management of Slow-Moving Items," *Journal of the Operational Research Society*, vol. 42, no. 7, pp. 605-608.

AUTHOR BIOGRAPHY

FRANK GRANGE is an experienced manager in Andersen Consulting's Centre for Process Simulation and Modelling. He has more than 20 years of experience applying simulation and optimization modeling to applications in business engineering, information technology, telecommunications, manufacturing, defense/aerospace, petroleum, and mining. He is currently an Adjunct Professor at the Colorado School of Mines, where he received a B. S. degree in Chemical and Petroleum Refining Engineering in 1972, and an M. S. and Ph. D. in Mineral Economics in 1974 and 1977, respectively.