

Title of manuscript:

Replenishments in Price Fluctuating Situations:
Comparing the Performance of Optimal and Heuristic Approaches

Authors:

Keith A. Willoughby¹

Department of Finance and Management Science
Edwards School of Business
University of Saskatchewan
Saskatoon, SK Canada S7N 5A7
Phone: (306) 966-2128
Fax: (306) 966-2515
E-mail: willoughby@edwards.usask.ca

James J. Chrisman

Department of Management and Information Systems
College of Business and Industry
Mississippi State University
Mississippi State, MS 39762
Phone: (662) 325-1991
Fax: (662) 325-8651
E-mail: jchrisman@cobilan.msstate.edu

¹ Corresponding author

ABSTRACT

Determining effective replenishment strategies in stochastic environments can be a complex task. As a particular example of a price fluctuating setting, we analyze the replenishment of automobile fuel on a long journey. We develop four heuristic approaches to determine the timing of fuel purchases. Two of our approaches combine an assessment of current fuel prices and on-hand gasoline levels. A third heuristic exclusively examines prices, while the remaining approach solely considers on-hand fuel levels. We conduct a simulation experiment, and compare the performance of these heuristic policies to the optimal solution obtained via the Wagner-Whitin method.

Key words: price fluctuations, replenishments, heuristic methods.

1. INTRODUCTION

The analytical development of replenishment strategies often incorporates stochastic factors. Random demand and supply disruption represent two major sources of uncertainty for purchasing managers. Further, unit prices may fluctuate for products such as bulk commodities (cereals, sugar, fuel, etc.), stock market instruments and raw materials. In this case, a critical objective is to determine the timing of required replenishments. In other words, when is the price sufficiently low to warrant a purchase? Kingsman (1986, p. 358) declared that little attention has been devoted to “the problem of buying materials whose prices are continually fluctuating in a stochastic manner over time”. We would agree that Kingsman’s comment is still relevant today.

As a particular context for a price fluctuating situation, our paper considers the replenishment of automobile fuel for a driver on a particularly long journey. At each individual purchasing opportunity, the decision-maker has an important choice to make: given one’s on-hand level of fuel and the observed price, does one purchase gasoline, or

neglect this opportunity, hoping that lower prices might be observed at future communities?

We develop four heuristic approaches to determine the timing of fuel replenishments. Two methods (termed “Price Quantity” approaches) specifically analyze the cost and benefits of purchasing at a given price and current fuel level. When the benefits exceed the associated costs, a fuel purchase is justified. The methods differ with respect to the specific quantity to purchase; this shall be described in more detail later. Another technique (denoted by “Price String”) solely considers fuel prices. A purchase is made upon observing a string of decreasing prices, irrespective of on-hand fuel levels. The final approach (called “Threshold Level”) ignores prices. Replenishments are warranted when fuel levels drop at or below a particular threshold. We compare the performance of these heuristic approaches to the optimal solution obtained via use of the Wagner-Whitin method.

The format of this paper is as follows. The next section provides a review of the literature pertaining to the development of replenishment strategies in stochastic environments. The third section introduces our heuristic approaches for determining the timing of fuel replenishments. We also briefly describe the use of the Wagner-Whitin method to establish the optimal timing of purchases, provide an expression for total expected costs, and discuss the issue of dependent prices. Computational results are presented in the fourth section, while the concluding section features a summary of our paper and potential directions for further study.

2. LITERATURE REVIEW

The determination of the timing of replenishments is one of a class of problems known as “action-timing” problems. Howard (1971) offered an extensive analysis of these problems. In essence, a decision-maker must take a single action (this is different from our case in which an automobile driver must make a series of replenishment decisions). In Howard’s treatment, there exist a known number of opportunities at which the decision-maker may take this action, but the “value” of respective future opportunities is uncertain. Based upon this information, when should the action be taken?

Howard introduced several real-world decisions that may be modeled as action-timing problems. Interestingly, he discussed the purchase of automobile fuel for situations in which gasoline prices are random. As one might expect, foregoing the purchase of gasoline at a specific opportunity would prove attractive if the decision-maker were confident that subsequent gasoline prices were substantially lower than the one currently observed.

Other scenarios that may lend themselves to action-timing formulations include photography (a traveller on vacation with only a single frame remaining and no possibility of obtaining more film must make a critical decision as to when to take the picture), hiring (applicable for the case of a manager who must reach an irrevocable decision regarding the fate of each applicant before the next one is interviewed), and marriage (although Howard was reluctant to provide further details!).

Ahn and Kim (1998) developed a methodology for action-timing problems in which the decision-maker’s understanding of the environment is revised based on

previous observations. Further, Ross (1983) illustrated a stock option model, in which stock prices fluctuate and the decision-maker must decide both if and when to exercise the option. A decision rule is provided to indicate the conditions upon which the option ought to be exercised.

Some researchers have examined purchasing strategies for commodities (e.g. sugar, copper, cereals, and so forth). In these cases, prices fluctuate and the critical decision for a purchasing manager involves determining the timing and sizing of individual replenishments. Early contributors to this decision problem were Fabian, Fisher, Sasieni and Yardeni (1959). Basing their research on a real-world organization, they found that purchasing managers generally followed a myopic strategy (namely, purchasing sufficient materials to simply meet current monthly requirements). After forming one, two and three-month forecasts of future prices, the researchers developed decision rules to guide purchasing decisions. Specifically, if the unit price in a future period was less than the current price plus the holding cost to that period, then one's on-hand inventory ought to be sufficient to cover the requirements of those periods preceding the forecasted lower price. In this way, one would expect to have relatively little on-hand during the period in which unit prices were expected to be lower. Substantial purchases could then be made at these lower prices.

Kingsman (1969) analyzed the purchase of commodities subject to price fluctuations. He developed a simplified model in which a purchaser was entitled to a single buying opportunity per month, there were no restrictions on respective purchase quantities, and future prices followed the same probability density function. He then provided various decision models in which these simplifications were removed. His

decision rule consisted in a set of stock levels a buyer should target to achieve at each price break. Lower unit prices resulted in higher desired stock levels (that is, a purchaser would procure greater quantities at these lower prices).

Guimaraes and Kingsman (1990) developed a heuristic policy to guide the daily purchases of maize. Their heuristic suggested a price-trend buying policy. Specifically, buyers should purchase commodities only if the real trend in unit prices were positive. The more positive the trend, the larger would be the size of the resulting purchase. In this way, buyers could replenish their stocks before prices became prohibitively high.

Kingsman and Boussofiane (1989) examined specific ordering policies when prices solely rise (as in the case of inflation). The time between successive price increases followed a negative exponential probability distribution. They developed critical levels of inflation at which it became attractive (in terms of overall costs) to purchase now and thus carry stock in order to avoid a price increase in subsequent periods.

In an effort to explore purchasing behaviour of storable food products, Boizot, Robin and Visser (2001) created an inventory model that helped a household determine an optimal stock level of any product. The authors considered storage costs, as well as the fixed cost associated with any purchase decision. Prices were allowed to fluctuate due to price promotions.

Bonser and Wu (2001) studied a fuel procurement problem faced by electrical utilities. Given uncertain demand and market prices, purchasing managers had to determine the attractiveness of long-term contractual supply commitments (made at a set price) or purchasing fuel at current market prices. Edwards, Flaim and Howitt (1999)

developed a dynamic model of hydroelectric generation and purchases of replacement power to meet contractual demand obligations. Variability in peak and off-peak spot electricity prices contributed significantly to their suggested purchase strategy.

Finally, Pearn and Chien (1998) explored the traveling purchaser problem. In this case, a decision-maker has to buy various commodities in different markets, with dissimilar commodity prices. The objective is to minimize the sum of travel and purchase costs.

3. DEVELOPMENT OF QUANTITATIVE APPROACHES

We shall describe each of our quantitative approaches in turn; namely, the “Price Quantity” methods (described as “Price Quantity Fill” and “Price Quantity Look Ahead”), the “Price String” heuristic, the “Threshold Level” procedure, and the Wagner–Whitin optimal solution method.

3.1 PRICE QUANTITY FILL

Let us denote each community along the route by j ($j = 1, 2, \dots, n$), with n being the destination. At each community, the decision-maker observes a certain gasoline price, p_j (\$/litre). We have chosen a stationary uniform probability distribution to model gasoline prices. For now, we will assume that prices at individual communities are independently distributed. The values p_{min} and p_{max} represent the extreme values of the distribution. Such a distribution is plausible in those instances in which prices are subject to fluctuations, but a decision-maker is nonetheless confident about the highest and lowest values that could potentially be observed. We further note that, given this

probability distribution, the average fuel price, \bar{p} , is:

$$\bar{p} = \frac{P_{\min} + P_{\max}}{2} \quad (1)$$

We shall represent fuel usage (litres/kilometre) by U and inter-community distance (kilometres) by D . Without loss of generality, we shall assume a constant inter-community distance. As a result, the vehicle uses UD litres of gasoline between each successive pair of communities.

The following notation is required:

- A : fixed cost of replenishment (\$)
- C : capacity of vehicle's gasoline tank (litres)
- R : forced replenishment fraction
- RC : forced replenishment point (litres)
- I_j : on-hand fuel level upon entering community j , but before any replenishment decision (litres)
- Q_j : purchase quantity (litres) at community j (since in this case a purchase decision implies a “fill up”, the purchase quantity is equivalent to $C - I_j$)

By analyzing the benefits and costs of fuel purchases at individual communities, this heuristic approach determines when a decision-maker ought to take advantage of a purchasing opportunity. Fuel purchase is warranted if the following expression is satisfied:

$$(p_j - \bar{p})(C - I_j) + A < 0 \quad (2)$$

We propose that benefits accrue when one purchases sufficiently large quantities of fuel at relatively low prices. These “low prices” are determined by observing the difference between the observed p_j and the overall average price, \bar{p} . Further, one must recognize that the act of stopping to purchase fuel along the journey does require some time. Consequently, the “costs” of fuel replenishment involve the opportunity costs

associated with stopping to purchase gasoline (the parameter A). We must incorporate this feature into our replenishment strategy; else, a decision-maker, following this purchasing approach, would simply fill up with gasoline at any community that offered a lower price than \bar{p} (even if the concomitant purchase quantities were quite small). Our strategy suggests that a purchase decision ought to be made when there is either a relatively large difference in fuel prices or on-hand levels are becoming somewhat low (or both).

This heuristic approach needs to incorporate a forced replenishment point in order to guard against stockouts. Instances could conceivably arise in which, as on-hand fuel quantities are diminishing, a decision-maker encounters a series of communities featuring “higher than average” gasoline prices. Thus, (2) would not be satisfied and a purchasing decision would not be made (a stockout would then result). Consequently, we suggest that a decision-maker following this heuristic strategy be obligated to fill up whenever $I_j \leq RC$ (even if (2) is not satisfied at the particular community j).

We can represent this heuristic decision rule quite succinctly as follows:

For $j = 1, 2, \dots, n-1$

IF $(p_j - \bar{p})(C - I_j) + A < 0$

THEN $Q_j = C - I_j$

ELSE IF $I_j \leq RC$

THEN $Q_j = C - I_j$

ELSE proceed to community $j+1$

3.2 PRICE QUANTITY LOOK AHEAD

This heuristic approach is similar to the Price Quantity Fill method. As before, the decision-maker examines fuel prices at each community along with on-hand fuel levels. A purchase is warranted if (2) is satisfied.

However, in the event that (2) is not satisfied, and the decision-maker is obligated to make a forced replenishment decision (i.e. $I_j \leq RC$), we propose that the concomitant replenishment quantity not to be a “fill up”. Rather, this heuristic method suggests that one ought to purchase sufficient fuel to cover the requirements of traveling to the next two communities. Given that the vehicle uses UD litres of fuel between successive communities, a purchase of $2UD$ litres of fuel should be made. Such a relatively low fuel purchase has the drawback of necessitating a stop (and incurring the fixed replenishment cost), but it permits the decision-maker to reach a future community (one of the next two along with the route) with small on-hand fuel quantities. Recall that, for now, fuel prices are independently distributed, so there ought to be a very small chance of observing high prices at a long string of consecutive communities. It is very likely that the decision-maker will see a relatively low price somewhere in the immediate future. Arriving at this procurement point with dwindling on-hand quantities could provide cost savings by allowing the decision-maker to purchase substantial amounts at low prices.

This policy is represented as follows:

For $j = 1, 2, \dots, n-1$

$$\text{IF } (p_j - \bar{p})(C - I_j) + A < 0$$

$$\text{THEN } Q_j = C - I_j$$

ELSE IF $I_j \leq RC$
THEN $Q_j = 2UD$
ELSE proceed to community $j+1$

3.3 PRICE STRING

In this particular case, the decision-maker only considers observed prices when making a purchase decision. Specifically, a replenishment is made if a string of two consecutive lower than average prices are seen. Further, the price at the current community must be lower than the price witnessed at the previous purchase opportunity. This would result in a pair of relatively low prices with a decreasing trend. Given our independent distribution of gas prices, it is rather likely that future communities may offer higher prices. Thus, purchasing now permits a decision-maker to exploit current price advantages.

We note that a potential drawback of the Price String heuristic is that one's purchasing decision may involve relatively small quantities. Even though prices are lower than average, the decision-maker would incur the fixed costs of replenishment each time a purchase is made.

To prevent running out of fuel, we incorporate the same rule as described in the previous two sections. That is, one is forced to make a purchase if $I_j \leq RC$, even if prices are disadvantageous. For sake of simplicity, we assume that a decision-maker fills up with fuel if a forced purchase is required.

The representation of this policy is as follows:

For $j = 1, 2, \dots, n-1$

```

IF  $p_j < p_{j-1} < \bar{p}$ 
THEN  $Q_j = C - I_j$ 
ELSE IF  $I_j \leq RC$ 
    THEN  $Q_j = C - I_j$ 
    ELSE proceed to community  $j+1$ 

```

3.4 THRESHOLD LEVEL

Our final heuristic approach disregards observed prices. A decision-maker purchases fuel when one's on-hand stock level gets "low enough". The threshold level policy stipulates that gasoline replenishments are made whenever $I_j \leq RC$. Therefore, this strategy is:

```

For  $j = 1, 2, \dots, n-1$ 
    IF  $I_j \leq RC$ 
    THEN  $Q_j = C - I_j$ 
    ELSE proceed to community  $j+1$ 

```

3.5 WAGNER-WHITIN METHOD

The Wagner-Whitin (1958) method is an application of dynamic programming, and may be used to provide an optimal solution. Given that one knew, a priori, the particular prices at the various communities, the Wagner-Whitin method can be used to determine the least cost timing of replenishments.

As an example of the steps involved in this algorithm, consider a driver who arrives at a particular community, say community 1, with no fuel on-hand. That is, a

replenishment is required at this point. (The interested reader is referred to Silver, Pyke and Peterson (1998) for a concise description of the algorithm). Let us suppose that $F(1)$ refers to the total cost of a replenishment in community 1, of sufficient size to cover the fuel requirements from community 1 to community 2. Trivially, since there will be only one possibility for $F(1)$, we know that it will be $F^*(1)$, the minimum cost solution for travel between communities 1 and 2.

Further, $F(2)$ refers to the total cost of a purchase in community 1, with enough fuel to satisfy needs from community 1 to community 3. For $F(2)$, there are 2 possible options to consider to meet requirements through to community 3:

- We could buy fuel in community 1 (enough to last us to community 2), then purchase enough fuel in community 2 to meet our needs to community 3.

- We could buy enough fuel in community 1 to last us until community 3.

Now, $F^*(2)$ represents the minimum of these two possibilities.

For $F(3)$, there are there are 3 possible options to consider to meet requirements through to community 4:

- We could cover ourselves to community 3 in the best possible fashion (that is, $F^*(2)$), then purchase sufficient fuel in community 3 to last us until community 4.

- We could cover ourselves to community 2 in the best possible fashion (that is, $F^*(1)$), then buy enough gasoline in community 2 to last us until community 4.

- We could have a single purchase in community 1 of sufficient quantity to satisfy our needs until community 4.

As above, $F^*(3)$ is the minimum of these three options.

The algorithm continues in this fashion, proceeding through the communities. A least cost replenishment strategy is then produced.

3.6 DETERMINATION OF TOTAL COSTS

We can derive an expression for determining the total costs of following any of the purchasing strategies (Price Quantity Fill, Price Quantity Look Ahead, Price String, Threshold Level, or Wagner-Whitin method). We begin by letting m represent the number of fuel purchases made along the journey (obviously, $m \leq n$). Moreover, any gasoline on-hand at the conclusion of the journey (ie. I_n) is "costed back" at the specific unit fuel price for which it was purchased. We will use the term p_k to represent the fuel price at the final community in which a purchase was made. It is subtracted from the total cost expression to signify that it represents something of a revenue for a decision-maker. Having this fuel on-hand means that it could satisfy usage on a subsequent trip. Furthermore, treating it as revenue does not penalize a decision policy that provides a fuel replenishment in a particular community close to the eventual destination, say, $n-1$.

Hence, the total costs, for any replenishment policy, are:

$$(A * m) + \sum_{j=1}^m (p_j * Q_j) - (I_n * p_k) \quad (3)$$

The first term represents the total fixed costs of replenishment, the second term provides the total fuel purchasing costs, while the last term costs back any on-hand fuel at the conclusion of the journey.

3.7 DEPENDENT FUEL PRICES

Recall our earlier assumption regarding the distribution of gasoline prices. Specifically, we assumed that prices at individual communities were independently distributed, drawn from a stationary uniform probability distribution. However, it is

likely that there is some dependence between community gasoline prices. This could occur since some regions may be known to have higher (or lower) prices than neighboring areas. For example, one could encounter a string of communities with relatively high prices (perhaps due to regional taxation issues).

We can incorporate dependence within the gasoline prices in the following fashion. Let us use the following expression (for $j \geq 2$):

$$p_j = \gamma p_{j-1} + (1-\gamma) \times (\text{Random price drawn from the } [p_{min}, p_{max}] \text{ distribution}) \quad (4)$$

where γ is a “dependence factor” between 0 and 1 (and, of course, p_1 would be a random price drawn from the uniform distribution). If $\gamma = 0$, then the gasoline prices are independently distributed. As γ increases, prices in successive communities become less and less independent. The price in a particular community has a larger share in determining the price in successive places. In the extreme case where $\gamma = 1$, then all prices are identical.

4. COMPUTATIONAL RESULTS

We shall now describe some computational results to compare the cost performance of our heuristic purchasing policies to the optimal solution obtained via the Wagner-Whitin method. Obviously, we know that the Wagner-Whitin approach will provide us with the least-cost solution assuming that we know the specific prices at all communities ahead of time. But, how much of a cost penalty is incurred by using the heuristic methods? Do any work better than the others? Under what conditions do the heuristic approaches offer superb, or deficient, performance?

We shall consider a decision-maker on a 3,000-kilometre journey. Further, we

will assume that one enters community 1 with no fuel on-hand. Thus, a purchase is required at this stop. We have the following “base case” parameter values (note that since we would never want a “stockout” of fuel, we use numerical problems where $RC \geq UD$).

<i>A</i> :	\$1.00
<i>C</i> :	60 litres
<i>D</i> :	20 kilometres (hence, they are 150 communities along this route)
<i>R</i> :	0.125
<i>RC</i>	7.5 litres
<i>U</i> :	0.1 litres per kilometre
<i>UD</i> :	2.0 litres
γ :	0 (prices are independently distributed)

Uncertainty in gasoline prices is reflected by the following variables:

p_{\min} :	\$0.40 (per litre)
p_{\max} :	\$0.50 (per litre)

As a result, \bar{p} is \$0.45 per litre. In one of our simulation experiments, we will use a normal distribution to represent fuel price uncertainty. In this particular case, the mean will be \$0.45 per liter, while the standard deviation is \$0.0167. We chose this specific value of the standard deviation so that the extreme points of the distribution (± 3 standard deviations from the mean) will be roughly equivalent to the minimum and maximum values provided by the uniform distribution.

A computer simulation model was developed using @RISK to generate 1,000 replications of this journey. As our performance measure, we computed the total average costs for each purchasing policy. We also varied different parameters, on a “one-at-a-time” basis, to establish the effect of a particular parameter setting on our results. Table 1 illustrates our simulation results (the percentage increase in total costs versus the optimal solution is shown in parentheses). The shaded cells correspond to the

least-cost approach of all the heuristic methods.

===== insert Table 1 about here =====

Table 1 shows that either of the Price Quantity heuristic policies leads to improved cost performance versus the Price String or Threshold Level decision rules. The best heuristic was always the Price Quantity Fill or Price Quantity Look Ahead strategy; in fact, each of these heuristic approaches was preferred in an equal number of cases (5). Generally, the Price String and Threshold Level policies featured cost penalties that were usually twice as much as those observed with the least-cost heuristic. Our claim is that a decision-maker could make better purchasing decisions by utilizing a decision rule that analyzes fuel price, current on-hand gasoline levels and the opportunity costs of stopping for fuel at each purchasing point. This is preferred to considering price levels in isolation, or merely replenishing when on-hand levels drop at or below a specific quantity. Of course, the Wagner-Whitin method outperforms all of our heuristic approaches, but this algorithm features increased levels of computational complexity. As previously outlined, this method (as we applied it in this particular case), assumes that a decision-maker knew specific fuel prices, a priori, at each of the communities.

A narrower price range, or greater dependence between fuel prices, leads to better average performance for the Threshold Level heuristic. For the cases in which the price distribution was limited between \$0.43 and \$0.47 per litre (note that the average was still \$0.45), the average cost difference between the best heuristic (Price Quantity Fill) and Threshold Level was just over \$1.00. When prices in neighboring communities became extremely dependent ($\gamma = 0.80$), the difference in average costs was only \$0.15. Further, we note that in each of these cases, the cost performance of these two heuristic policies

closely approached that of the optimal solution. The Price Quantity Look Ahead rule featured extremely poor performance in the case of extreme dependence in fuel prices (cost penalty of 35.60%). By following this policy, a decision maker would only purchase small quantities of gasoline in the event that a forced replenishment was required. Should a string of relatively high prices be encountered (which is plausible with the significant dependence in the price distribution), one would be caught in a situation of continually purchasing small quantities of fuel. The fixed costs of purchasing, incurred each time a stop is made, would combine to produce the large costs observed in this case.

Results were very similar for different levels of R , the forced replenishment fraction. By the way, the value of $R = 0.04$ suggests that a decision-maker would wait until fuel volume was at $1/25^{\text{th}}$ of tank capacity before purchasing fuel. This yields a forced replenishment point of 2.4 litres, barely above the 2 litres required for travel between successive communities. Changing the level of A , the fixed cost of making a purchase, leads to greater average costs across the five replenishment policies. In this case, the Price Quantity approaches continue to maintain healthy advantages over the other two heuristic policies.

5. SUMMARY AND DIRECTIONS FOR FURTHER STUDY

We have examined heuristic approaches to guide the timing of fuel replenishments when unit prices are subject to fluctuations. We developed Price Quantity heuristic procedures that relied on a consideration of costs and benefits at each respective purchasing opportunity. The purchase of fuel is warranted only in those cases in which benefits exceed (opportunity) costs. Through a set of simulation experiments, we have shown that considering the interaction of fuel prices, on-hand gasoline levels and opportunity costs leads to improved cost performance over following policies that limit one's attention to either price or quantity levels. Although we developed a model in the context of purchasing automobile fuel, these heuristics could conceivably be extended to other scenarios (e.g. commodities, raw materials) in which price uncertainties exist. We feel these policies mimic the manner in which managers consider relevant factors in making real-world purchasing decisions. To wit, the Price Quantity heuristic permits a decision-maker to consider price and on-hand inventory. Should prices be higher than average, and a purchase decision must be made, one procures a set quantity. The Price Quantity Look Ahead rule suggested that in the event of a forced purchase quantity, one only buy enough to meet needs for the immediate future, trusting that lower prices will be forthcoming.

There are a few extensions that one could introduce into our modeling effort. Recall that our heuristic approaches solely considered the timing of fuel purchases. We were not concerned with the quantity one ought to purchase, simply with when this decision should be made. We either filled up with gasoline, or purchased sufficient fuel

to cover requirements for the next two communities (should a forced replenishment be needed in the Price Quantity Look Ahead heuristic). Ideally, one could conduct a more elaborate, sophisticated analysis into the specific quantities one ought to purchase at each procurement point.

Our model examined deterministic usage of automobile fuel. In reality, factors such as wind and vehicle speed may cause gasoline usage to fluctuate. Stochastic usage would mean that decision-makers would need to retain a specific amount of fuel as “safety stock”. The actual level of safety stock would be selected so as to provide a certain likelihood of no stockouts during the journey.

Moreover, we have used a constant inter-community distance (D). We could modify our model to reflect the case in which the distance between communities was not constant (i.e. D_{ij} would give the distance between communities i and j). Changing our model to reflect this situation would require an adjustment in the forced replenishment point, RC . In order to permit zero chance of stockouts (for the case of deterministic usage), we would need to set this replenishment point in such a way that it covered total gasoline usage between the most-widely separated successive communities. In other words, RC would equal $\max_{i,j} (D_{ij} * U)$.

Finally, to obtain an enhanced understanding of how the different parameters affect cost performance, we could run simulations by varying two or more parameters at a time.

REFERENCES

- Ahn, J.H. and Kim, J.J., "Action-timing Problem with Sequential Bayesian Belief Revision Process," *European Journal of Operational Research*, 105, (1998) 118-129.
- Boizot, C., Robin, J-M., and Visser, M., "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, 111, 470 (2001), 391-419.
- Bonser, J.S. and Wu, S.D., "Procurement Planning to Maintain both Short-term Adaptiveness and Long-term Perspective," *Management Science*, 47, 16 (2001), 769-786.
- Edwards, B.K., Flaim, S.J., and Howitt, R.E. Howitt. "Optimal Provision of Hydroelectric Power under Environmental and Regulatory Constraints," *Land Economics*, 75, 2 (1999), 267-283.
- Fabian, T., Fisher, J.L, Sasieni, M.W. and Yardeni, A., "Purchasing Raw Material on a Fluctuating Market," *Operations Research*, 7, (1959), 107-122.
- Guimaraes, R.C. and Kingsman, B.G., "A Day-to-Day Buying Policy for Commodities – A Study of Purchasing Maize," *Journal of the Operational Research Society*, 41, (1990), 25-37.
- Howard, R.A., **Dynamic Probabilistic Systems, Volume II: Semi-Markov and Decision Processes**, John Wiley & Sons, New York (1971).
- Kingsman, B.G., "Commodity Purchasing," *Operational Research Quarterly*, 20, (1969), 59-79.
- Kingsman, B.G., "Purchasing Raw Materials with Uncertain Fluctuating Prices," *European Journal of Operational Research*, 25, (1986), 358-372.
- Kingsman, B.G. and Boussofiane, A., "Ordering and Stockholding under Price Inflation when Prices Increase in Successive Discrete Jumps," *Engineering Costs and Production Economics*, 17, (1989), 395-407.
- Pearn, W.L. and Chien, R.C. "Improved Solutions for the Traveling Purchaser Problem," *Computers and Operations Research*, 25, 11 (1998), 879-885.
- Ross, S.M., **Introduction to Stochastic Dynamic Programming**, Academic Press, New York (1983).
- Silver, E.A., Pyke, D.F. and Peterson, R., **Inventory Management and Production Planning and Scheduling**, John Wiley & Sons, New York (1998).
- Wagner, H. and Whitin, T.M., "Dynamic Versions of the Economic Lot Size

Model,” *Management Science*, 5, 1 (1958), 89-96.

Table 1
Simulation Model Results

Scenario	Heuristic Purchasing Policies				Optimal Solution
	Price Quantity Fill	Price Quantity Look Ahead	Price String	Threshold Level	
Base Case	134.80 (5.38%)	134.42 (5.08%)	142.62 (11.49%)	140.86 (10.12%)	127.92
A = \$0.50	133.44 (7.14%)	132.36 (6.27%)	135.53 (8.82%)	138.14 (10.91%)	124.55
A = \$1.50	137.19 (4.48%)	138.83 (5.73%)	150.14 (14.34%)	143.87 (9.57%)	131.31
$p_j \sim U[0.43, 0.47]$	139.92 (3.16%)	146.11 (7.72%)	142.30 (4.91%)	140.97 (3.93%)	135.64
$p_j \sim U[0.37, 0.53]$	133.00 (10.69%)	131.75 (9.65%)	140.24 (16.71%)	141.08 (17.41%)	120.16
R = 0.04	134.77 (5.33%)	134.18 (4.87%)	142.13 (11.08%)	140.88 (10.11%)	127.95
R = 0.25	135.48 (5.88%)	135.40 (5.81%)	143.94 (12.49%)	142.09 (11.04%)	127.96
$\gamma = 0.40$	137.49 (3.93%)	140.90 (6.51%)	156.00 (17.92%)	140.98 (6.57%)	132.29
$\gamma = 0.80$	140.86 (2.53%)	186.30 (35.60%)	165.66 (20.58%)	141.01 (2.63%)	137.39
$p_j \sim N(0.45, 0.0167)$	137.59 (3.45%)	142.11 (6.85%)	142.26 (6.96%)	140.96 (8.83%)	133.00