Some other possible improvements in the present program are as follows:

1. Developing an improved automatic blocking strategy process;
2. Developing a technique to combine blocks and form trains automatically;
3. Developing a cost model to compare various strategies on a cost basis; and
4. Converting the whole system to time sharing with interacting blocking strategy and train editing capabilities.

The above is only a partial list, and several other features have been suggested during the course of the project. We hope that the present programs can eventually be augmented, by incorporating all the significant features, so that a highly efficient and useful tool will be available for railroad operators.

ACKNOWLEDGMENTS

The research work on which this paper is based was carried out as a part of a research project requested by the U.S. Railway Association, whose support and permission to present and publish the paper are gratefully acknowledged. However, the contents of this paper reflect our views, and we alone are responsible for the facts and accuracy of the information presented. The contents do not necessarily reflect the official views or the policy of the U.S. Railway Association.

REFERENCES

1. W. Siddiqee, D. A. D'Esopo, and P. L. Tuan. Blocking and Train Operations Planning. Stanford Research Institute, Menlo Park, CA, SRI Project 3759, Oct. 1976.
2. D. D'Esopo and W. Siddiqee. User's Manual for Network Analysis Computer Programs. Stanford Research Institute, Menlo Park, CA, SRI Project 3759, Nov. 1975.
3. P. L. Tuan and H. S. Proctor. A Railroad Classification Yard Simulation Model. Proc., 1975 Winter Computer Simulation Conference, Sacramento, CA.

Publication of this paper sponsored by Committee on State Role in Rail Transport.

Inventory Model of the Railroad Empty-Car Distribution Process

Craig E. Philip, Freight Car Utilization Program, Association of American Railroads
Joseph M. Sussman, Department of Civil Engineering, Massachusetts Institute of Technology

Techniques to improve freight-car fleet use are of considerable interest to the railroad industry. One potentially high improvement area is the disposition of empty cars within the network. This paper reports the first results of inventory control applied to one aspect of the process, namely the sizing of empty-car inventories at points in the network. First we evaluate existing techniques for distributing empty cars on a rail network. These techniques deal primarily with optimizing empty-car movements from areas of surplus to areas of deficit. To account for variations in supply and demand, we designed a discrete event simulation model that can determine optimum inventory level, for a single terminal area, as a function of (a) daily supply variations, (b) daily demand variations, and (c) cost of holding a car in a terminal awaiting loading compared to cost of having no car available to satisfy shipper demand. A first attempt to use the model to evaluate the performance of an actual railroad terminal area indicates that excessive inventories are maintained in surplus terminal areas. The applicability of the model to a real railroad operating situation is also demonstrated.

Empty-car distribution is an unavoidable problem for most railroads, because demand and supply are typically unbalanced in any given region. Thus, surpluses and deficits at terminal areas are inevitable, and some mechanism must be employed to move cars from points where they are not needed to points where they are.

Shippers feel the impact of the distribution mechanism directly. Car availability will largely be determined by the ability of the railroad to efficiently move cars from surplus to deficit areas.

This recurring need to manage and monitor car movement has come to dominate current empty-car distribution processes. The techniques used to allocate cars usually employ standard static optimization methods and thus rely on the hypothesis that levels of supply and demand will not vary significantly. Variations, however, do exist, and one of them is periodic shortages caused by railroads unreliably routing cars from surplus to deficit areas.

Some empty-car distribution practices have evolved to cope with this problem; individual terminal distributors, for example, often maintain an inventory of empty cars to protect against the uncertainties of supply and demand. Still, since distribution mechanisms seldom consider inventory levels, no strategy for determining appropriate inventory levels has yet been proposed, and costs to the railroad incurred by wasted car days or lost loads due to shortage can be directly related to these levels.

This report evaluates the theoretical implications and tests the methodology of one strategy for determining inventory level in a railroad operating environment. The proposed strategy grew naturally from our reexamination of the empty-car distribution process from the perspective of the local or terminal decision maker. Several theoretical solutions to the empty-car distribution problem, such as existing network models that determine flow rules, are contrasted with a theoretical construct of the need for empty-car inventories.

A discrete event simulation model of empty-car distribution determines the best target inventory level for a particular terminal area given the supply and demand characteristics of that terminal.

Finally, the results of sensitivity analyses of the impact of changes in railroad and shipper behavior on the optimum inventory level are presented, and the results of our first attempt to use the model to predict the best inventory level in a railroad terminal area, based on the actual flows into and out of the terminal, are given.

The results of this research effort have, to date, been encouraging. The model of freight-car distribution we tested accounts for the relationship between service reliability and freight-car utilization, and it may prove to be a useful tool when applied parallel to a traditional flow model to improve car distribution strategies. Much of what follows has been founded on the work of Philip (1).

THEORETICAL APPROACHES TO EMPTY-CAR DISTRIBUTION

Efficient empty-car distribution satisfies shipper demands at the lowest possible cost. There are two necessary and related approaches to empty-car distribution. The first, with its emphasis on empty-car movements to balance surplus and deficit areas, has been adopted in some form by most railroads. The second, which focuses on variable car supply to satisfy variable demand, has not been systematically analyzed. A theoretical base for such an analysis is the subject of what follows.

Traditional System Focus of Empty-Car Distribution

"The essence of car distribution and assignment is the process of providing destinations to empty cars and monitoring their movements towards those destinations" (2, case iv. 1). This definition embodies a rather appealing philosophical approach to the car distribution process when it is viewed from the system perspective. A car emptied at a point on the railroad where it is not needed for reloading must be moved to a potential reloading point. The process of deciding where to send which cars, then, becomes the essence of car distribution.

This process is further delineated in Figure 1, which highlights the three subsystems in all traditional car distribution systems:

1. Identification or prediction of empty-car supply,
2. Identification or prediction of the demand for empty cars, and
3. Allocation or control or both of car movements

Figure 1. Four-part traditional empty-car distribution process.

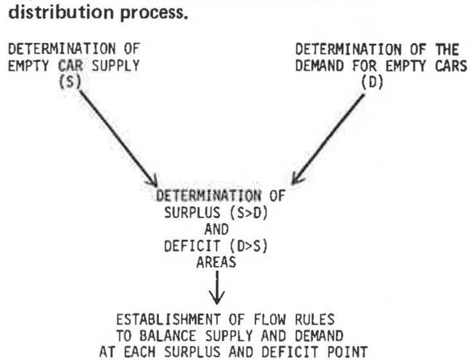

from surplus to deficit areas.

Demand and supply estimations define surplus and deficit areas, which themselves are only the inputs into the flow rule decision process; the quality of these flow rule decisions is necessarily limited by the accuracy of the demand and supply estimates. A recent report prepared for the Federal Railroad Administration (FRA) concluded that shipper demand varies a great deal (3, p. iii). Of even greater importance, it was found, is that demand level is not measured adequately by railroads and, with a few exceptions, is not even formally forecast.

Car supply itself is subject to at least as much variation as demand, because the receipt of unloaded cars from industry is the principal source of empty cars. In fact, empty-car supply is likely to vary even more than demand because of the variations introduced by unreliable movement of the cars by the railroads themselves.

Car Distribution as a Classical Transportation Problem

If certain simplifying assumptions are made, the problem of distributing empty cars from surplus to deficit points fits nicely into the form of what Dantzig (4, p. 299) and others have called the "classical transportation problem." This empty-car distribution problem as perceived in the classical sense is precisely one of determining a set of flow rules governing the movement of cars from surplus directly to deficit points; the objective function is to "minimize the cost of moving the cars into position [for loading] from the locations where they become available" (5, p. 147).

As one might expect, this has not been overlooked by theorists or practitioners. Models using either linear programing techniques or some other network optimization algorithm have been proposed repeatedly in the literature. One model for distributing wood rack-cars, was implemented with good results on the Louisville and Nashville Railroad Company, and the Missouri Pacific Lines periodically use a linear programing model to establish empty-car distribution guidelines. Dan Berman of the Southern Railway Company reports that a linear program is at the core of a system that manages the movement of the entire free-running fleet.

Shortcomings of the Traditional Approach

At first blush, the linear programing technique would seem to be an ideal solution to the problem of empty-car distribution, because it is offered as an optimum allocation of the empty-car fleet and thus minimizes the costs of allocation. Unfortunately, the solution is only optimum if the simplifying assumptions required to yield the solution are in fact correct, and in this case critical assumptions at odds with the realities of railroading have been made. For example, quality and uniformity of demand and supply forecasts are needed to define surplus areas, and the solution is optimum only when these forecasts are accurate and demand is stable.

A second, more subtle assumption has been made in forming the objective function. Here it is assumed that the only cost important to the distribution process is the penalty cost of moving cars from surplus to deficit areas. If the first assumption were true and the situation were in fact deterministic, then this second assumption might be plausible. The real costs associated with the stochastic nature of the process, however, should not be ignored in the solution strategy.

It is not the purpose of this study to indict existing car distribution practices because of the tremendous variations inherent in the levels of supply and demand. The system view of the problem, with its emphasis on flow rules and car movements, is an absolutely necessary component of any car distribution mechanism. Nevertheless, variations in supply and demand as well as forecasting difficulties need to be accounted for in any car distribution procedure.

Inventory Approach to Car Distribution: Terminal Perspective

The model described in this section evolved naturally from what has been called the traditional or system perspective on empty-car distribution. The typical proposed definition of the process resulted in network solutions to the problem. An alternative definition, however, suggested by Johnson (6) demands a new perspective and different solution strategies: "The main function of railway freight car distribution systems is to control the inventories of empty cars held to buffer the supply and demand at the loading points." This definition shifts the focus from the movement between areas to the surplus and deficit areas themselves.

Figure 2 (where L = loaded cars as supply, E_L = emptied cars, and D = demand) shows the terminal perspective appropriate for car distributors at both surplus and deficit areas. The process involves less network optimization and more inventory control. This conceptual framework suggests that the variable nature of empty-car supply and demand is related to the inventory maintained in terminal areas. A methodology to formally specify this relation is proposed in the next section.

SIMULATION MODEL OF TERMINAL EMPTY-CAR CONTROL

To provide a new perspective on the empty-car distribution problem, we have sought a technique that would clarify the operation of a small part of the existing railroad environment and, if appropriate, would give us a simple tool for managing this environment more effectively. To this end, the elements of a discrete event simulation model that represents the empty-car inventory decisions of a surplus or deficit terminal area are described.

Figure 2. Terminal perspective on the empty-car distribution process.

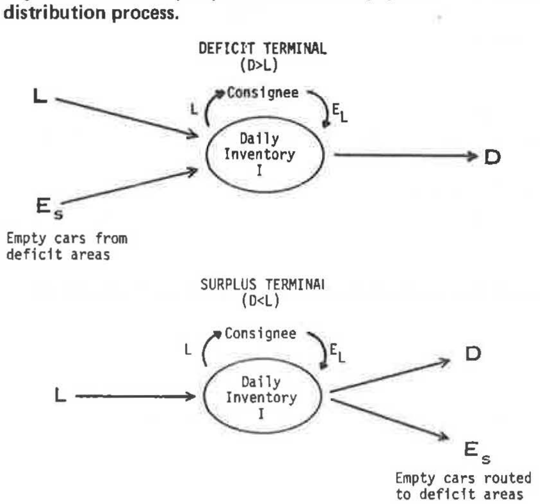

Basic Structure of the Model

As illustrated in Figure 2, most railroad terminal areas can be classified as being either "sources" (surplus) or "sinks" (deficit) for empty cars.

In a surplus terminal situation, empty-car supply normally exceeds demand. Each day consignor demand for empties will first be satisfied; then any empty cars remaining will be used to replenish the inventory (Figure 2). The model determines the number of cars, called the "target" inventory level, that should be in the inventory after replenishment. The following daily decision structure ranking is followed: (a) all daily demands are satisfied by the daily supply; (b) any empty cars not needed to satisfy the daily demands are used to bring inventory up to the target level; and (c), finally, any remaining empties are sent to a deficit area according to the flow rules.

In a deficit terminal area, demand for empty cars is generally greater than the number of loads terminated. Additional empties will be transshipped to the terminal area according to system flow rules, so that in the long run demand and supply will be in balance. Given this balance, the terminal decision maker cannot rely upon the daily flow of cars to establish or replenish his inventory, and additional cars will periodically be sent to the terminal to replenish the inventory. The model determines how large this inventory (initial inventory) should be at the beginning of each simulation period. Empty cars never flow out of the terminal area in this formulation, so a very simple decision structure is possible: All current demands are satisfied if possible, and any remaining empty cars are placed in the inventory.

For both surplus and deficit situations, the inventory level on a given day i will be determined by the day's new supply of and demand for empties, by the previous day's inventory, and by the prespecified target or initial inventory level (I°):

$$I_i = f(I_{i-1}, E_i^S, E_i^D, I^\circ) \qquad (1)$$

where

I_i = inventory level at the end of day i;
E_i^S = arrival of empty cars on day i;
E_i^D = demand for empty cars on day i;
I° = prespecified target or initial inventory level for the simulation.

I° is implicitly a part of the decision structure, because it effectively increases the supply of cars every day by an amount I°. Thus, for a day during the simulation period when supply exceeds demand, the result will be to increase the day's remaining inventory by I° cars. Likewise, on days when demand exceeds supply, the added cars will reduce the number of unsatisfied demands by an amount I°.

An optimum initial inventory level will be one that balances the costs of increasing the inventory level on surplus days against the costs of reducing the unsatisfied demand costs if shortages occur. The next section presents the method used to arrive at such a solution.

Determining Optimum Inventory Level

The terminal decision structures for both surplus and deficit situations have been specified. While they differ in several important ways, the daily inventory levels in both cases depend on the same set of four independent variables shown in Equation 1, of which only I° is specified by the decision maker. The other three depend on the external environment. For each I°, a different set of daily inventory levels I_t will emerge from the decision structure. The problem becomes one of selection from among different sets of daily inventory levels.

Terminal Cost Function

The principal function of an inventory of empty cars in a terminal is to diminish the impact of variations in

Figure 3. Using cost function to determine terminal cost.

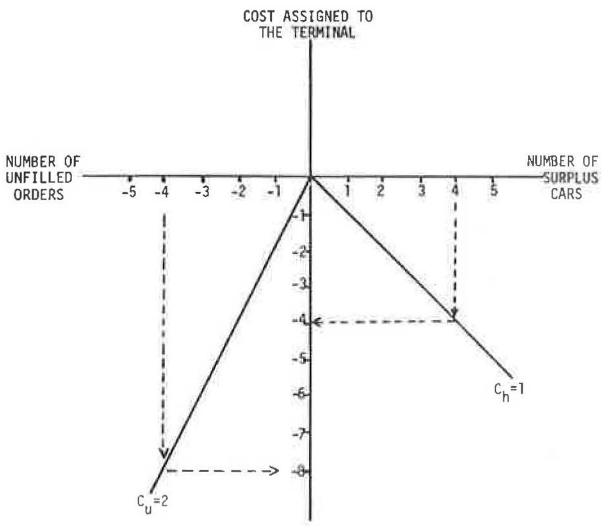

Figure 4. Example of the optimization process.

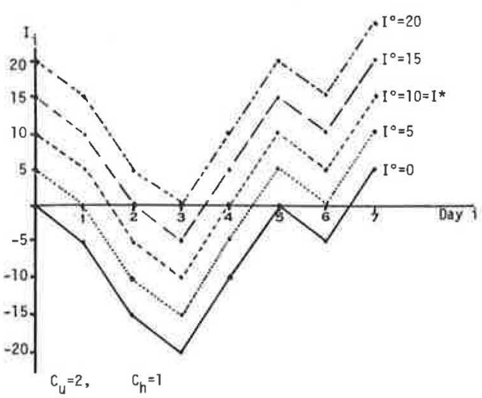

Figure 5. Generalized input structure.

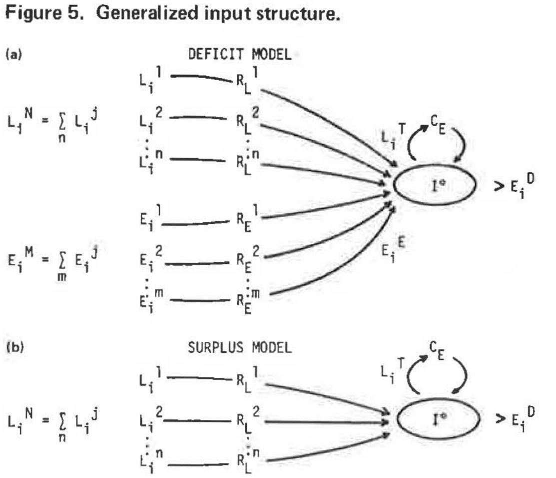

demand and supply levels on the area's ability to satisfy empty-car demand. Each car added to the inventory decreases the risk of a shortage, but increases empty-car inventory cost. Incurring some cost is an inevitable consequence of demand and supply variability, so the objective should be to minimize the expected total shortage plus inventory cost.

We can define a terminal cost function that accounts for this. For any inventory level (positive or negative), the cost function defines the cost to the system, and a simple piecewise linear function can be denoted as follows:

$$C(I_i) = \begin{cases} C_u(I_i) & \text{if } I_i < 0 \\ -C_h(I_i) & \text{if } I_i \geq 0 \end{cases} \tag{2}$$

Figure 3 illustrates the case where unsatisfied demand cost C_u equals 1 and holding cost C_h equals 2. If the inventory level is $I_t = 4$, then $C(4) = -1(4) = -4$; and if $T_t = -4$, then $C(-4) = +2(-4) = -8$. In a similar fashion, calculating the cost of a sequence of daily inventory levels is a simple matter of totaling individual costs for each day:

$$C^T = \sum_{i=q}^{Z} C(I_i) \tag{3}$$

where

q = first day of the simulation period;
Z = final day of the simulation period;
$C(I_t)$ = cost for the inventory level I_t; and
C^T = total cost for the period.

Recalling that the daily inventory level is itself a function of $I_t = f(I^{\circ})$, it is also possible to conclude that cost is a function of I°. Each value of I° implies the unique sequence of daily inventory levels that follow (Figure 4).

	Daily Cost				
Day	$I^{\circ}=0$	$I^{\circ}=5$	$I^{\circ}=10$	$I^{\circ}=15$	$I^{\circ}=20$
0	0	-5	-10	-15	-20
1	-10	0	-5	-10	-15
2	-30	-20	-10	0	-5
3	-40	-30	-20	-10	-15
4	-20	-10	0	-5	-10
5	0	-5	-10	-15	-20
6	-10	0	-5	-10	-15
7	-10	-10	-15	-20	-25
Total	-120	-80	-75	-85	-125

Each has a certain value of C^T associated with it. The remaining task is to find the I° value that defines the sequence of inventory levels with the lowest C^T.

Finding Minimum Cost Inventory Level

The nature of the decision and cost structures in the empty-car inventory problem defined here will ensure a well-behaved situation. As initial inventory increases, the number of demands not satisfied can diminish, but the daily inventory level can only increase. The piecewise linear cost structure equates fewer demands with lower cost and a larger inventory with increased cost over its entire range, so increasing I° may at first reduce C^T, but, if large enough, it will eventually increase C^T. Thus, the value of the C^T function will fall to the point where increased inventory cost exceeds decreased unfilled demand cost associated with an increase in I°. (The piecewise linear cost structure is not required to create the conditions described; however, any function whose slope is always positive for inventory levels less than zero and always negative for the inventory levels greater than zero will lead to the same optimality conditions.)

Given this functional relationship between C^T and I^o, it is possible to define a very simple search routine to determine the optimum inventory level (I^*). By successively testing $C^T(I^o)$ for increasing values of I^o, the optimum value of I^o will be found when $C^T(I^o)$ stops decreasing. This process is illustrated in a sample problem in Figure 4, in which the holding cost is 1 and the unsatisfied demand cost is 2; each day's individual cost for each I^o is recorded along with the C^T. For instance, the cost on day three for $I^o = 0$ is the inventory level (-20) multiplied by $C_u(2)$, which equals -40. As the results indicate, the optimum I^o is 10.

Components of the Input Subprogram

The previous section's inputs were parameters defining the utility function (C_u and C_h), the daily empty-car supply (E_i^S), and the daily demand for empties (E_i^D). Much of the simulation model is given over to the process of determining these values, as is described in the following sections.

Cost Function Parameters

The cost function parameters are easily specified for our model's purposes because they are treated deterministically. However, they prove difficult to determine accurately in any particular inventory or railroad situation. Buffa (7) suggests that "though it is not difficult to develop a model for buffer stock based on the concept of balancing inventory and stockout cost, more often than not it is difficult, if not impossible, for management to isolate a realistic stockout cost."

In the railroad environment also, neither inventory cost nor cost of delayed or unfulfilled demand for cars is well defined. We therefore ran the model repeatedly using the same car supply and demand inputs but different cost ratios in order to reduce the importance of the cost specifications. It is the linear nature of the cost function that makes this possible, because the actual optimization routine is sensitive only to the cost ratio C_u/C_h and not to the absolute values of C_u and C_h themselves.

Daily Empty-Car Supply

One principal goal of this modeling effort was to determine the impact of rail network operations on the need for empty-car inventories in terminal areas. The input structure we established gives the model user a wide range of options with respect to the specifications of rail service and shipper behavior. This structure for surplus and deficit terminal areas is depicted in Figure 5.

For both surplus and deficit, the model assumes that loaded cars are destined for the terminal area in question. On each simulation day i, n groups of loaded cars (L_i^j) are generated. The number of cars in an individual L_i^j will vary from day to day according to the specified distribution, which can be different for each group. A Monte Carlo sampling procedure from each distribution was employed to determine the values of the L_i^js each day. This ensures that each L_i^j will have some known expected average value that keeps total daily number of loaded cars generated each day (L_i^N) constant over the simulation period.

At a deficit area, empty cars are also routed to the terminal. The same rules that apply to L_i^j and L_i^N also apply to E_i^M and E_i^n.

Trip-time distributions (R_E^j, R_L^j), which indicate how often a trip takes a particular number of days, are used to describe the railroad operating environment. The Monte Carlo procedure for appropriate trip-time distribution is used to determine the travel time for each group of cars, L_i^j or E_i^j. The trip-time distribution may be different for each group of cars, $R_E^1 = R_E^2$, for instance, but a particular group's trip time will be governed by one distribution during the entire simulation. Having selected the trip times, the arrival day for each group of cars at the terminal is determined; therefore arrival day equals departure day plus trip time.

Knowing the arrival day of each group of cars and the number of cars in each group, it is possible to determine the total number of loaded, L_i^T, and empty, E_i^E, cars arriving on day i. Empty cars immediately become part of the supply, but the loaded cars must first be unloaded and then returned to the terminal according to a "time to empty" distribution.

The two components of the daily new empty-car supply are now well defined. The E_i^L, added to the E_i^T, defines the total number of new empty cars each day (E_i^S):

$$E_i^S = E_i^T + E_i^L \tag{4}$$

From the terminal's perspective, a rather complicated network, characterized by numerous points of loaded and empty-car generation and equally numerous trip-time distributions, is reflected in the variation of a single input variable, E_i^S.

Daily Demand for Empty Cars

Daily empty-car demand is simpler to define than empty-car supply. The model user specifies a daily shipper demand distribution, which is sampled to determine a daily demand for cars (E_i^D), just as the loaded and empty group size distributions were sampled. The model in its present form also assumes that E_i^D and supply (E_i^S) are not correlated with one another. Use of a more complicated demand structure that assumes numerous independent sources of empty-car demand is possible but was not pursued in this study, because the behavior of a group of sources can be adequately represented for our purposes by a simple demand source.

Results of both a theoretical and an applied study using the model are presented in the next section.

EXPERIMENTAL RESULTS

Verification of our simulation model is a multifaceted problem, but, as it is based on theories of inventory control, its results should be consistent with theory. Also, as a model of rail terminal operations, it should be capable of evaluating current operation realistically. Since the core of the model is independent of the input subprograms, the model can be used to test both cases.

Theoretical Results Based on Hypothetical Input Data

This model can be used to show how rail operations, shipper behavior, and perceived operating costs affect optimum inventory level. Of the many possible relationships that can be analyzed, three of the most relevant concern the impact of

1. Improving trip-time reliability,
2. Lowering the cost of unfilled demand relative to the cost of holding a freight car, and

3. Decreasing variability in the number of loaded and empty freight cars generated.

The initial hypothesis, based on common sense and the classical theories of inventory control, is that each change should lower the optimum inventory level. Inputs that isolate these relationships in the deficit case are described and utilized in the following analysis.

To eliminate some sources of potential variation in the inputs, we created a simplified input structure that can be readily modified to isolate the three relationships outlined above. The demand for empty cars (E_1^o) is assumed to be constant each day and equal to 200 cars. Exactly 100 loaded cars in four equal groups of 25 and 100 empty cars in five equal groups of 20 cars are generated each day. The time required to empty each loaded car is always a day.

In addition, for a particular run, the trip-time distributions for all groups are modeled identically. This does not mean that the trip time for each group will be the same on a particular day, because it is independently selected from the underlying distribution. Finally, the ratio of the late load to inventory holding cost is assumed to be 2:1, and we used a simulation period of 7 d, based on a general railroad official's consensus that a weekly planning horizon for many car distribution decisions seems reasonable.

Impact of Trip-Time Reliability

With these inputs, the only source of variability comes from the trip-time distribution, which allows study of the first relationship between trip-time reliability and the optimum inventory level by repeatedly running the model using different distributions.

The trip-time distribution defines both expected trip time (the mean) and the reliability of the trip. The variance is a conventional measure of a distribution's dispersion. The "n-day-%," developed by Martland (8), is also a measure of a distribution's central tendency and is derived as the maximum percentage of cars with trip times in a single n-day interval. Martland also proposes use of the measure "%-n-days-late," which is defined as the percentage of cars arriving n or more days later than the mean. As reliability improves (indicated by a smaller variance or a larger 2-day-%), predictability of trip time also increases.

Typical railroad trip-time distributions were selected from a compilation of actual trip times reported by a large shipper between seven origin-destination pairs. These distributions are listed in Table 1.

Two hundred 1-week simulations were performed with each of the trip-time distributions; the results are graphically presented in Figure 6, which shows that improvements in rail reliability (according to any of the three measures) are generally accompanied by smaller optimum inventory levels.

This imperfection reflects our inability to precisely define the variance of an actual distribution by using a single measure. Combining several of the measures, however, does provide a more adequate explanation of the results. The second distribution (2-day-% = 79) appears to have a higher inventory than can be explained by the 2-day-%, but if its %-2-days-late is also considered, it becomes clear that the higher inventory is caused by the extreme values of the distribution.

These seven trip-time distributions are used in the remainder of the analysis, and the seven runs of the model made with the inputs as specified for this first analysis will be referred to as the base case.

Impact of Cost Ratio

Simple changes in the base case permit testing of the second relationship for the impact of changes in the ratio of unfilled demand cost to holding cost. The base case is modified by first increasing the ratio from 2:1 to 3:1 and then decreasing it to 1:1. As before, 200 1-week simulations were run for each of the seven trip-time distributions, and in each case an optimum inventory level was determined. The results are illustrated in Figure 7.

When the penalty cost of not satisfying a shipper's demand increases relative to the empty-car holding cost (an increase in ratio), the optimum inventory required to minimize the terminal decision maker's cost increases regardless of the level of trip-time reliability. If the penalty is small, the decision maker will keep an inventory only if service is very unreliable.

These results are also consistent with those found in classical inventory theory. Safety stock is only justified when the cost of not maintaining it exceeds that of maintenance. When the cost is the same (per car in this case), a car supply must be very unreliable before the cost of the inventory justifies the reduction in the risk of stockout associated with it.

Impact of a Stochastic Car Generation Rate

The base case was designed so that trip times would vary while number of cars generated each day was constant. To consider the impact of variability in the number of loaded and empty cars generated, the base case is twice modified so that the number of cars generated is normally distributed, with means still equal to 25 and 20 for the five loaded and four empty moves respectively, and standard deviations of 25 and 50 percent of the respective means in the two cases tested.

Again, 200 1-week simulations were run for each trip-time distribution and standard deviation combination. Note also that an additional trip-time distribution, one with perfect reliability, was tested in order to isolate the impact of variation in the car generation rate on the inventory level.

The results are presented in Figure 8 and show that variability in the car generation rate does in fact increase the needed inventory for all levels of trip-time reliability.

It is difficult to make a direct comparison between the relative impacts of trip-time unreliability and generation variation, because there are no directly comparable measures of variance. The results do support the hypothesis, however, that each source of additional uncertainty increases the required inventory level.

Summary of Theoretical Results

To isolate and evaluate certain critical relationships between inputs to the terminal model and optimum inventory level, the model was exercised by using an appropriately designed set of inputs. Each of the tested sources of supply and demand variation increased the required inventory level both individually and when combined together. In addition, increases in the cost ratio, representing a larger stockout cost, also raised the optimum inventory. Each of these results is consistent with inventory control theory, the basis for the modeling structure. Although these results have no direct applicability to a particular railroad operating situation, we shall use the model to evaluate performance of an actual terminal area in the next section.

Table 1. Trip-time distributions for seven origin-destination pairs.

No. Trip Days	Percentage of Cars per Trip						
	Pair 1	Pair 2	Pair 3	Pair 4	Pair 5	Pair 6	Pair 7
1	0	0	0	0	0	0	0
2	0	2	6	0	0	0	0
3	0	2	45	0	17	0	0
4	4	64	33	11	41	0	4
5	12	15	14	26	19	0	15
6	57	10	2	43	8	15	27
7	27	1	—	16	15	28	13
8	1	2	—	0	—	30	12
9	—	1	—	0	—	11	9
10	—	0	—	0	—	2	6
11	—	2	—	2	—	7	4
12	—	0	—	0	—	5	4
13	—	1	—	2	—	0	5
14	—	—	—	—	—	2	1
Mean no. trip days	6.1	4.6	3.6	5.9	4.6	8.1	7.5
Percentage of cars 2 d late	84	79	78	69	61	59	42
Percentage of cars more than 2 d late	0	8	2	4	15	14	20
Variance	0.55	2.66	0.77	1.94	1.65	2.24	5.58

Figure 6. Impact of trip-time reliability on optimum inventory level.

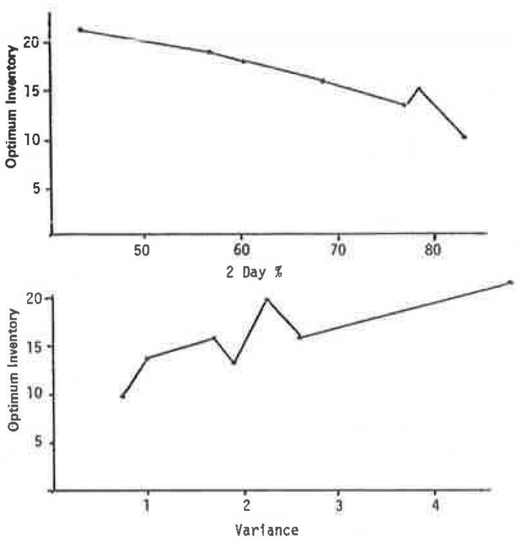

Figure 7. Impact of changes in cost ratio.

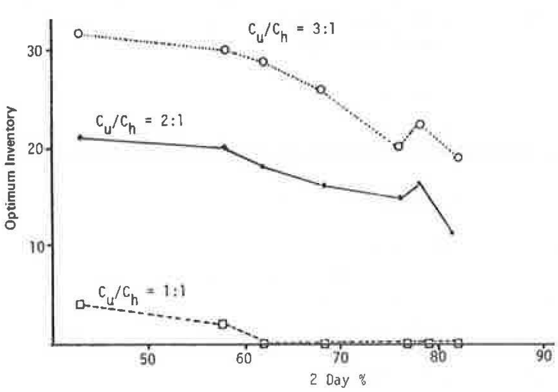

Applying the Model to a Real Terminal Area

The theoretical results suggest that an appropriately specified inventory model structure can be used to determine the empty-car inventory that should be maintained in each terminal or area of a railroad. If the model can be verified, then the inventory information, along with data on total available empties and empty re-

Figure 8. Impact of supply variability on optimum inventory level.

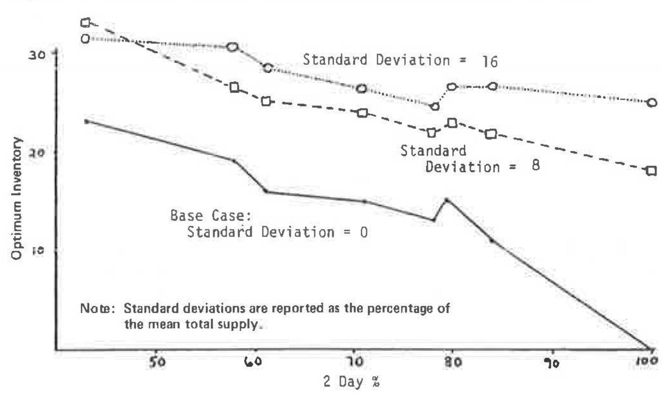

positioning costs, might become part of a method to better allocate empty cars.

Currently, no railroad attempts to quantitatively determine an optimum empty-car inventory level. Several large class I railroads were contacted during the course of this study: we found one both willing and able to provide daily data on empty-car movements for each car type and major terminal area. In concert with this particular railroad, we attempted to apply the model.

Our railroad has a traditional but sophisticated car distribution process that relies on flow rules to balance system surpluses and shortages. The number of cars maintained in a terminal area's inventory is determined by the terminal manager himself. The system decision maker can define a route for all surplus cars, but the local decision maker is largely responsible for determining how many cars are surplus.

Available Data

For each terminal area and car type, daily disaggregated data (each car is identified) are available showing (a) number of empty cars on hand and (b) number of empty cars out to industry. While these two categories do not in themselves define the inputs required by the model, the disaggregate nature of the data makes it possible to calculate empty inventory, empties loaded, empty arrivals, and empty cars routed to another terminal. One key variable is obviously missing—the number of empty cars demanded but not provided—and without this information it is impossible to balance the inventory cost against the cost associated with lost or delayed loads.

Table 2. Model results for the case study.

Ratio of Unfilled Demand to Holding Cost	Target Inventory		Average Daily Inventory		Average Daily Flow to Deficit Areas		Total Late Loads	
	Actual	Generated[a]	Actual	Generated[a]	Actual	Generated[a]	Actual	Generated[a]
Existing situation	—	—	49	—	9	—	0	—
1:1	0	0	0	0	9	9	71	33
2:1	0	0	0	0	9	9	71	33
5:1	4	2	3	1	9	9	47	22
10:1	16	5	13	4	9	9	10	11
100:1	22	14	19	12	9	9	0	0

[a] Results generated by using the stochastic input data.

Figure 9. Analysis of mean and standard deviation for area.

While this problem is perhaps unavoidable in any deficit area, the characteristics of this railroad's decision-making structure permit and perhaps encourage a different problem in surplus areas. The surplus terminal manager is not penalized for maintaining an excessive inventory of empty cars (in the model's terminology, his holding cost, C_h, is very low), leading him, as the theoretical results suggest, to maintain a large inventory. If surplus areas do maintain excessive inventories, then a precise measure of demand is available in the "number of empties loaded," since we assume that loads are seldom lost and empty cars are always available. A surplus area was therefore selected for analysis.

Characteristics of the Selected Surplus Area

Data listing the empty cars on hand and out to industry at the terminal each day for the selected car type were collected for a 1-month period in 1976.

It was possible to create a complete history of the empty-car decisions at the terminal from these data. On the average, 14 cars arrived each day; 5 cars were damaged; 9 cars were routed to the appropriate deficit area. An average inventory of 50 cars was maintained over the period, and at no point did the inventory level drop below 20 cars, which supports the assumption that loads were never lost. These characteristics are schematically summarized in Figure 9.

Using the Model to Evaluate the Terminal's Performance

The historical data, provided by the railroad, record variations in supply and demand as they actually occurred. Thus, instead of repeatedly simulating the situation using hypothetical inputs, the model was run by using the actual supply and demand data.

With these inputs, the optimum empty-car inventory was calculated for different ratios of delayed load to holding costs. The results, reported in Table 2, reveal that the average inventory required to avoid any lost or late loads was only 19 cars; average daily flow was 14 cars; average daily demand for empties was 5 cars; and the maximum inventory was 91 cars.

For the purposes of comparison, the stochastic model form was also run. Both supply and demand were assumed to be normally distributed with means and standard errors equal to those found in the actual demand and supply data. These results were also reported in Table 2 and are similar to those determined by using the actual data.

For the 1-month period examined, both sets of results suggest that the inventory level maintained during the period is oversized regardless of the cost hypothesis used. The average inventory of cars needed to ensure that no load be lost is only 19 cars; the average of 30 additional cars sitting in the inventory were not needed and were of no value to the terminal operator. Of course, even a superficial, qualitative examination of the daily supply and demand characteristics indicates that an inventory of 50 cars is excessive. The model is useful to the extent that it can quantify the degree of excess inventory, given the unique characteristics of the flows into and out of the area.

The cost of maintaining this excess inventory will depend on a number of factors, including car type, time of year, and age of the fleet. The particular car type analyzed was in short supply when the data for this study were collected, but during times of car shortage inventory cost will approach the opportunity cost associated with loads lost or delayed elsewhere on the railroad.

SUMMARY AND CONCLUSIONS

This paper reports the first results of applying an inventory control theory to one aspect of the empty-car decision-making process. The empty-car distribution process and its importance to overall railroad performance were first reviewed. Then the theory of inventory control and its applicability to empty-car distribution were outlined and a simulation model based on this theory presented. Theoretical results were discussed, and, finally, the model applied to an actual railroad situation.

The theoretical results are consistent with classical inventory control theory. A positive correlation has been discovered between variability in the supply or demand and the number of cars needed in the inventory.

The model has also been successfully utilized to evaluate the performance of an actual railroad terminal area, and the results in this case are consistent with railroad thinking. The model data required that a surplus area be analyzed; data needed to evaluate a deficit area were not available. During the 1-month period while data were being collected, the inventory level in that surplus area was found to be oversized. This conclusion corroborated opinions voiced by several individuals at the railroad. Application of the model to deficit areas should prove useful in the long run, but data on lost or late loadings are not presently available.

It should be recognized that these results are preliminary in nature. While the single case study performed does seem to verify the theoretical results, more terminal areas will be investigated. The model itself should be looked upon as part of a more comprehensive set of models for use as a tool for managers seeking to balance empty-car inventories over an entire network.

ACKNOWLEDGMENTS

This research, based on the work of C. E. Philip, was performed under the sponsorship of the Car Utilization Program of the Association of American Railroads, supported in part by the Federal Railroad Administration of the U.S. Department of Transportation.

REFERENCES

1. C. E. Philip. A Discrete Event Simulation Model of Car Distribution Decisions in a Railroad Terminal Area. Department of Civil Engineering, M.I.T., Cambridge, MA, masters thesis, Feb. 1977.
2. Manual of Car Utilization Practices and Procedures. AAR Freight Car Utilization Research/Demonstration Program, Washington, DC, June 8, 1976.
3. W. K. Minger and others. Freight Car Demand Information and Forecasting Research Project. Phase I—Final Report, AAR Research and Test Department, Washington, DC, Mar. 1975.
4. G. B. Dantzig. Linear Programming and Extensions. Princeton Univ. Press, New York, 1963.
5. W. W. White and A. M. Bomberault. An Algorithm for Empty Car Allocation. IBM Systems Journal, Vol. 8, 1969.
6. J. W. Johnson. Freight Car Distribution: The Nature of the Problem. CORS Journal, Vol. 1, No. 1, Dec. 1963.
7. E. Buffa. Production-Inventory Systems. Richard D. Irwin, Homewood, IL, 1968.
8. C. D. Martland. Railroad Trip Time Reliability: Evaluation of Performance Measures and Analysis of Trip Time Data. Studies in Railroad Operations and Economics, June 1972.

Publication of this paper sponsored by Committee on State Role in Rail Transport.