Finite Horizon Equipment Replacement Analysis

Joseph C. Hartman
Alison Murphy
Lehigh University

Report No. 04T-002

Finite Horizon Equipment Replacement Analysis

Joseph C. Hartman and Alison Murphy

Industrial and Systems Engineering, Lehigh University, Harold S. Mohler Laboratory,

200 W. Packer Ave., Bethlehem, PA, 18015-1522, USA, jch6@lehigh.edu

Abstract

The optimal solution to the infinite horizon equipment replacement problem with stationary costs is to continually replace an asset at its economic life. The economic life is the age which minimizes equivalent annual capital and operating costs (EAC), including purchase, operating and maintenance costs less salvage values. The finite horizon problem requires a sequence of asset service lives such that the total service equals the horizon. We explore whether applying the infinite horizon solution to a finite horizon problem is justified, as we conjecture this may occur often in practice. With our presented integer knapsack approach, we can define a bound on the minimum number of times an asset is utilized at its economic life in a finite horizon problem. The bound can be derived from any feasible solution, although we provide closed form solutions for the case of convex EAC values. Finally, we illustrate that the bound is useful for reducing the computation of the associated dynamic program (for a positive interest rate) and often helps in solving the associated integer program (for an interest rate of zero). As the integer program's constraint defines a cyclic group problem, the bound may prove useful in other applications.

Subject classifications: Facilities/equipment planning: replacement. Dynamic programming: deterministic. Integer Programming.

1 Introduction

Capital equipment typically wears with age, defined by increasing operating and maintenance (O&M) costs and decreasing salvage values. If service is required for a number of periods, it may be economical to periodically replace the equipment. This paper examines the equipment replacement problem over a finite horizon under the assumption of stationary costs.

Specifically, the equipment replacement problem requires a sequence of keep or replace decisions for a single asset over a given horizon T, which may be infinite. Alternatively, the solution can be defined as a sequence of asset service lives, or ages, which an asset should be retained with the total service equaling T periods. An asset must be salvaged on or before its maximum service life N. Under the assumption of stationary costs, an asset is replaced with the purchase of a new, identical asset. For a finite horizon problem, the asset is salvaged at the end of the study period. The objective is to minimize discounted purchase and O&M costs less salvage values. We refer to this problem as REP.

Bellman (1955) provided the first solution to REP which did not require that assets be retained for the same length of time over a finite or infinite horizon. His dynamic programming formulation defines the state space as the age of an asset with the decision to keep or replace the asset in each period. If an asset can be retained for N periods, the maximum number of states in a period is N. For a T period problem, this translates to solving the problem in $O(2NT)$, or $O(NT)$, time.

Wagner (1975) presented an alternative dynamic programming formulation in which the time period is the state of the system and the decisions are to keep the asset for $1, 2, \ldots, N$ periods. This is similar to the Wagner-Whitin approach for economic lot sizing in which complete lots are produced (Wagner and Whitin 1958). With the maximum number of states in a given period being 1, a maximum of N decisions per state, and a total of T periods, this dynamic programming formulation can also be solved in $O(NT)$ time. The Wagner formulation has been extended by researchers to deal with realities such as technological change and multiple challengers (see, for example, Oakford, Lohmann and Salazar 1984, Bean, Lohmann and Smith 1985, 1994).

We present an alternative formulation for the finite horizon REP with stationary costs which can be described as an integer knapsack problem with an equality constraint or as a knapsack-partitioning problem (Johnson 1980b). For the case when the interest rate is zero, it can be modeled as a cyclic group problem (Gomory 1965) and solved with integer programming. However, for interest rates greater than zero, the nonlinearity of the objective function is more easily captured with dynamic programming. In our formulation, the space in the knapsack is time. Items to be placed in the knapsack are asset service lives with their size defined by the length of service. The size of the knapsack is defined by the length of the horizon.

Clearly, our motivation in this paper is not to present a method for computational purposes, as those previously defined in the literature are more than adequate. Additionally, we do not claim that this model is more adept at handling different modeling nuances, such as technological change, as Wagner's model has shown to be over time.

Rather, this paper is motivated by the relationship between the optimal solution to REP over a finite and an infinite horizon for the case of stationary costs. The optimal solution to the infinite horizon problem is to repeatedly replace an asset at its economic life. The economic life of an asset is the age which minimizes the equivalent annual costs (EAC) of owning and operating the asset (Thuesen and Fabrycky 1994). These costs include the purchase and O&M costs less salvage values. In general, O&M costs rise with age while salvage values decline. Thus, the optimal solution trades off the high cost of replacement (purchase less salvage) versus increasing O&M costs over time. The economic life of an asset is typically computed by calculating the EAC of retaining an asset for each of its possible service lives, ages 1 through N, as an integer solution is generally required. The minimum is then chosen from this set.

The optimal solution of replacing an asset at its economic life is only valid under the assumptions of an infinite horizon and stationary costs. However, many situations occur in practice where an asset is required for a finite length of service, especially if it is acquired to meet the needs of a given contract.

Engineering economy textbooks, such as Eschenbach (1995), Fleischer (1994), Park (1997) and

Thuesen and Fabrycky (1994), illustrate how to determine the economic life of an asset, but few address finding an optimal solution to the equipment replacement problem over a finite horizon, such as Park and Sharp-Bette (1990). While we have no literature to reference, we suspect this is because the finite horizon problem requires the solution of a dynamic program or a network flow formulation, which are techniques not learned by all engineers or financial managers. Thus, it is reasonable to assume that in practice, the infinite horizon solution may be applied to the finite horizon problem.

In this paper, we are interested in examining the validity of applying the infinite horizon solution (replacing an asset at its economic life) to the finite horizon problem. Assuming the solution to REP is defined by a sequence of asset service lives, we are specifically interested in developing a bound on the minimum number of times an asset is retained for its economic life in a given finite horizon solution. This bound is useful for two reasons: (1) Due to its familiarity and computational ease, practitioners may apply the infinite horizon solution to the finite horizon problem. This bound will provide some information on the validity of that decision. Furthermore, we explore conditions under which the bound will be at least one, including the case of convex EAC values. Determining whether at least one asset is retained for its economic life is critical as this is the decision to be implemented at time zero. (2) The bound may also be used to reduce the computational burden of solving the associated dynamic or integer program (for the case in which the interest rate is zero) for REP. Again, previous solutions to the equipment replacement problem are computationally efficient, so our motivation is more closely aligned with (1) above. However, the bound may prove useful in solving other integer knapsack or cyclic group problems.

The paper is divided into analyses assuming a positive interest rate and an interest rate equal to zero. Sections 2 through 6 examine the dynamic programming formulation, for the case of a positive interest rate, and its associated bounds on the minimum number of assets retained for their economic life, both analytically with convex costs (Section 5) and empirically (Section 6). The following sections analyze the case where the interest rate is zero with the use of integer programming. In addition to analytical (Section 8) and empirical (Section 9) tests of the bound, experiments on how the bound aids in solving the IP are presented in Section 10. We conclude in Section 11.

2 The Positive Interest Rate Case: Dynamic Programming

We model REP as an integer knapsack problem with an equality constraint. As we do not place explicit bounds on the number of times an asset is utilized for a given length of time, this translates to the integer, unbounded knapsack problem. The problem can be solved as a traditional knapsack problem (inequality constraint) with dynamic programming if the solutions in which the capacity constraint does not hold at equality are ignored. The traditional approach (Gilmore and Gomory 1966) to solving this problem with dynamic programming defines the state of the system as the amount of filled capacity in the knapsack. The decisions entail placing items in the knapsack, which reduce the space remaining. The approach is known to be solvable in pseudo-polynomial time, $O(nK)$, where n is the number of items and K is the size of the knapsack. (The state space moves from 0 to K in each stage and can evaluate each item n in each stage.) A number of enhancements have been published to speed up the algorithm in practice (see Andonov, Poirriez and Rajopadhye 2000 and the references therein).

We present a different solution approach in order to solve REP. The alterations are motivated by our desire to bound the number of times an asset is utilized at its economic life in a finite horizon problem. They are also motivated by the fact that solving a knapsack problem does not provide the *sequence* of decisions required for REP. We first describe the method and then discuss how a sequence can be defined such that a solution to REP is produced.

In our dynamic programming formulation, each stage of the problem represents the opportunity to utilize an asset of service life n. As the maximum service life for an asset is N, there are N stages in our approach. At each stage, the decision is whether to purchase $0, 1, 2, \ldots, \lfloor T/n \rfloor$ assets of age n. We solve the dynamic program forwards with the state of the system defined as the total number of periods of service that have been accumulated through the current stage. The solution entails a number of assets and their respective service lives.

As noted earlier, the solution to REP requires a sequence of asset service lives. This is accomplished by ordering the stages of our dynamic program according to non-decreasing equivalent annual costs (EAC) of owning and operating an asset. Thus, the first stage examines the number of assets to be retained at their economic life, as this length of time corresponds to an asset operating at its minimum EAC. The second stage examines the number of assets with the next lowest EAC to be retained, etc. Note that in defining the stages in this manner, we define a sequence in which the knapsack is filled. This is critical for discounting purposes and ensures an optimal solution to REP. Consider the following theorem.

Theorem 1 *Assuming an interest rate $r > 0$ and given a set of assets and their respective service lives, the minimum net present value cost sequence is that which orders the assets according to non-decreasing EAC values.*

Proof. Given a sequence of b asset lives $n_1, n_2, \ldots, n_b$, the cash flow diagram for the sequence is defined in Figure 1, where c_{n_i} is the EAC of using an asset for n_i periods. In the figure, the c_{n_1} values occur in periods one through n_1 while the c_{n_2} values occur in periods $n_1 + 1$ through $n_1 + n_2$. The final asset's cash flows occur in periods $T - n_b + 1$ through T.

The net present value, p, of the cash flow sequence, assuming a positive, periodic interest rate r for discounting, can be calculated as:

$$p = \sum_{i=1}^{b} \sum_{t=1}^{n_i} \frac{c_{n_i}}{(1+r)^{t+\sum_{j=1}^{i-1} n_j}}$$

This merely discounts each cash flow to time zero as the sequence defines the number of periods in the future that a cash flow occurs.

This is a proof by contradiction and construction. Assume the minimum net present value cost sequence is defined such that the assets are not ordered according to non-decreasing EAC, such as

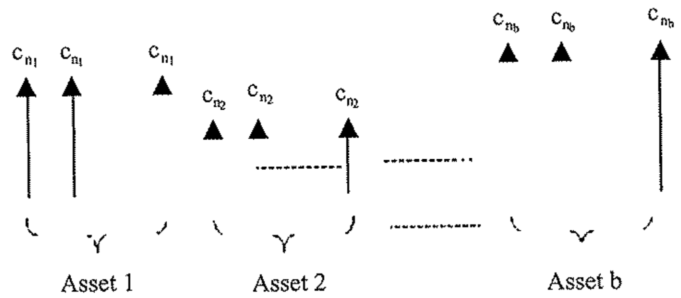

Figure 1: Cashflow diagram for a sequence of assets.

that depicted in Figure 1. (It is assumed that the drawing is to scale such that the length of the arrows represent the magnitude of the EAC values.)

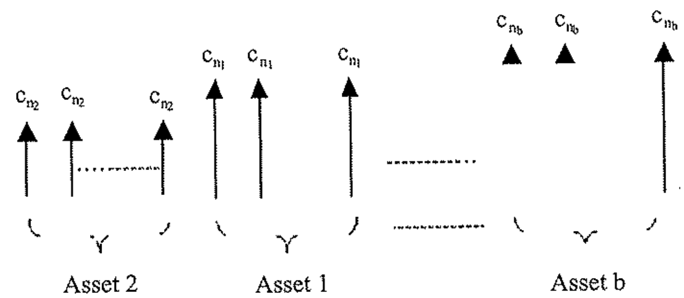

Figure 2: Cashflow diagram for re-sequenced assets.

Construct a new sequence of cash flows, as in Figure 2, such that the positions of the c_{n_2} cash flows and c_{n_1} cash flows are reversed. As $r > 0$, the new sequence will clearly have a lower net present value and by contradiction, the sequence not ordered by non-decreasing EAC cannot be optimal. Further, as any sequence not ordered according to non-decreasing EAC values can be replaced by a lower net present value cost sequence that is ordered accordingly (by switching their positions, as in our example), only a sequence ordered according to non-decreasing EAC values can be optimal. □

Theorem 1 defines the optimal sequence for utilizing assets over a finite horizon, given the length of time each asset is to be in service. This is critical to our dynamic programming model which analyzes asset service lives in each stage according to non-decreasing EAC values.

One final modification for the dynamic program, when compared to the traditional integer knapsack approach, is required due to discounting. If, for example, a solution entails keeping two assets for three periods each, then the net present value of the cost of using the assets are different as the second asset is purchased three periods later than the first asset. If the interest rate is positive, the purchase of two assets for three periods each is not equivalent, in terms of the net present value costs, to twice the cost of purchasing one asset for three periods. The cost of purchasing an additional asset of the same type must be discounted. Note that all assets in the sequence are discounted accordingly.

This can be easily handled with dynamic programming as the costs of purchasing $1, 2, \ldots, \lfloor T/n \rfloor$ assets are explicitly modeled. As each stage is defined by the length of time an asset is in service and the dynamic program tracks the number of periods of service accumulated thus far, this discounting can be handled straightforwardly. However, note that this would lead to a non-linear objective function if it were to be modeled as an integer program.

Given these changes to the traditional approach for the integer knapsack problem, our dynamic program, DP, can be formulated. Define the functional equation as follows:

$f_i(t)$ = minimum net present value of costs of owning and operating assets with service lives of $n_1^*, n_2^*, \ldots, n_i^*$ periods through time period t.

Recall that we sequence asset lives according to non-decreasing EAC values. This sequence must be mapped to the stage i such that the first stage $i = 1$ corresponds to the number of times an asset is retained for its economic life n_1^* and the second stage corresponds to the number of times an asset is retained for n_2^*, the length of time resulting in the second lowest EAC, etc. With $p(n_i^*)$ referring to the net present value of an asset's EAC values (at the time of occurrence) and α the single period discount factor, the recursion is defined as:

$$f_i(t) = \min_{m:\ t - mn_i^* \geq 0} \left\{ \alpha^{t - mn_i^*} \sum_{j=1}^{m} \alpha^{(j-1)n_i^*} p(n_i^*) + f_{i-1}(t - mn_i^*) \right\}, t = 1, 2, \ldots, T \tag{1}$$

The recursion is solved over stages $i = 1, 2, \ldots, N$, corresponding to the N possible service lives ordered according to non-decreasing EAC values $n_1^*, n_2^*, \ldots, n_N^*$. In each stage, the decision is to determine the number of copies m of an asset with service life n_i^* to incorporate into the solution. As we solve the DP forwards, the initial condition of $f_0(0) = 0$ is required. Also, if no copies of an asset are chosen in a given stage, it follows that $f_i(t) = f_{i-1}(t)$.

There are N stages, corresponding to the maximum service life of an asset and a maximum of $T + 1$ states in a stage, corresponding to the cumulative service acquired through stage i. For each state in a stage, there are a maximum of $\lfloor T/n \rfloor + 1$ possible decisions, where n represents the service life being examined in the stage. This maximizes at $T + 1$. This all corresponds to a worst case run time of $O(N(T + 1)^2)$.

Theorem 2 *DP solves the finite horizon REP with stationary costs.*

Proof. As defined, DP finds the minimum net present value cost sequence of assets to be used over the finite horizon T. According to Theorem 1, the stages are ordered such that the minimum net present value cost sequence is defined for any feasible replacement schedule. As the DP identifies the minimum net present value sequence over all feasible schedules, DP finds the optimal solution to REP. □

In the next section we pictorially examine this formulation and compare it to previous dynamic programming approaches. These differences provide our motivation for developing DP.

3 Networks for Dynamic Programming Approaches to REP

Figure 3 illustrates the representative networks for the dynamic programming approaches of Bellman (a) and Wagner (b) for the finite horizon REP. In Bellman's network, Figure 3(a), a node represents the age of the asset and is the state of the system, as labeled in the figure. Each arc represents either a keep or replace decision. Keeping the asset is denoted by an arc connecting nodes i and $i+1$ while replacing the asset connects nodes i and 1.

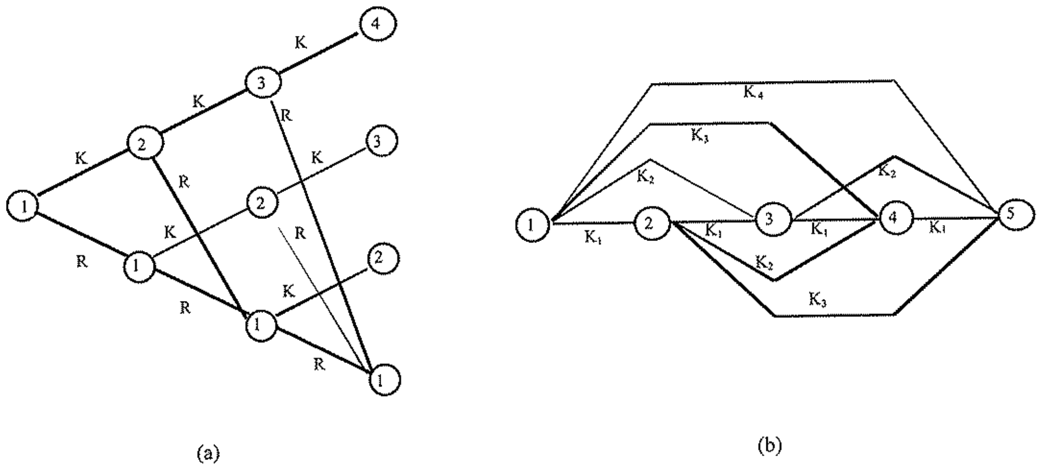

Figure 3: Dynamic programming networks for (a) Bellman and (b) Wagner approaches. Both are acyclic networks with flow from left to right.

In Wagner's network, Figure 3(b), a node represents the period and each arc represents the length of time to retain the asset. An arc connecting nodes t and $t+n$ represents retaining an asset for n periods. A stage represents one period of time in both the Bellman and Wagner formulations.

The network associated with DP is given in Figure 4. The nodes represent the cumulative service time that has been "placed" in the knapsack with the arcs representing the number of "copies" of an asset service life that is placed in the knapsack. In the first stage, either $0, 1, 2, \ldots, \lfloor T/i \rfloor$ copies of an asset with service life i can be placed in the knapsack. Each ensuing stage examines a new item (length of service life) to place in the knapsack, assuming the capacity constraint is not violated, resulting in a maximum of N stages.

These different networks illustrate the benefits of solving the associated dynamic programs. Wagner's network is the easiest to solve and can handle a variety of extensions, including technological change and multiple challengers. Bellman's network is slightly more difficult to solve and does not have the flexibility of Wagner's, but is easy to understand from a decision point of view as most managers determine whether to keep or replace an asset in a given period, not whether to

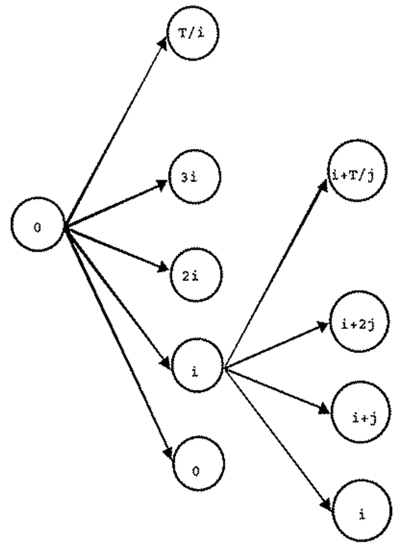

Figure 4: Dynamic programming networks for DP.

retain an asset for an ensuing two or three periods.

The benefit of the knapsack formulation is that one can readily determine the number of copies that the best asset will be used. Recall that in the optimal solution to REP with an infinite horizon under stationary costs, the asset is continuously replaced at its economic life. A natural question for the finite horizon problem is how many, if any, times do we replace an asset at its economic life? This is not easily answered using Bellman's or Wagner's formulation, unless the complete solution is found. However, this question can be asked in the first stage of the knapsack formulation approach. We specifically address this question in the next section.

4 Bounds on Assets Retained for their Economic Life

By definition, dynamic programs do not determine a feasible solution (and thus an upper bound to a minimization problem) until the final stage has been reached. This is in contrast to branch and bound approaches which generally solve relaxations to produce bounds in order to fathom branches. Morin and Marsten (1976), influenced by the applications of others, illustrated that simple bounding techniques can be used to fathom branches in dynamic programming networks. Their presentation was in the context of reducing the memory requirements of dynamic programs which can be extensive due to the curse of dimensionality (Bellman 1957) common to most applications.

Our motivation to implement simple bounding in DP is to determine a lower bound on the value of $m_{n_1^*}$, where n_1^* represents the economic life of an asset (service life where EAC is minimized) and m represents the number of copies used in the optimal solution to DP. Given a finite horizon T, a feasible solution to REP can be constructed easily by setting $m_{n_1^*} = \lfloor T/n_1^* \rfloor$, its upper bound, and taking $T - \lfloor T/n_1^* \rfloor$ copies of an asset retained for a single period. Another easily constructed solution is to combine the maximum number of assets retained at their economic lives with a single asset retained for $T - \lfloor T/n_1^* \rfloor$ periods. One could be more clever in constructing the solution to the remaining $T - \lfloor T/n_1^* \rfloor$ periods. For instance, a first-fit-decreasing (FFD) algorithm could be used. Regardless of how this is (easily) constructed, it represents a feasible solution with objective function value UB. (Note that the sequence of the assets is also important due to discounting.)

Examining the network of DP for the knapsack problem in Figure 4, we see that the first stage decisions are whether to place $0, 1, 2, \ldots, \lfloor T/n_1^* \rfloor$ copies of the economic life asset into the knapsack. Following each of these arcs and subsequent paths determines a feasible solution, or solutions, in which there are $0, 1, 2, \ldots, \lfloor T/n_1^* \rfloor$ copies of the minimum cost asset, respectively. If the optimal solution is unique, only one of these paths will define the minimum cost solution. A path can be fathomed (eliminated from consideration), as illustrated by Morin and Marsten (1976), if a lower bound produced on the branch exceeds any known upper bound. Given that the first stage decision is the number of copies of the minimum cost asset (assumed here), the best possible bound is to fill the remaining capacity of the knapsack with the next minimum cost asset. This asset has a life of length n_2^*.

Given m copies of the minimum cost asset, we can fit $(T - mn_1^*)/n_2^*$ copies of the second minimum cost asset. We are not concerned with integrality here as we only desire a lower bound to the problem. If the combination of the m copies of the economic life of the asset and the remaining space filled with the second best asset results in a lower bound that is greater than UB, then we can fathom the branch. The LB is increasing for decreasing m. Thus, if our LB exceeds the UB for a given value of m, the optimal solution must have greater than m copies of the minimum cost asset. This can be used to develop a lower bound on $m_{n_1^*}$, the number of copies of the economic life asset in the optimal solution. Specifically, we want to determine the minimum m such that the lower bound is less than our upper bound, or:

$$c_{n_1^*}(^{P/A,r,mn_1^*}) + \frac{1}{(1+r)^{mn_1^*}} c_{n_2^*}(^{P/A,r,T-mn_1^*}) \leq UB. \tag{2}$$

The term $(^{P/A,r,t}) = ((1+r)^t - 1)/(r(1+r)^t)$ is the present worth factor for an equal payment series (Thuesen and Fabrycky 1994) which merely reduces a periodic series of equal cash flows to a single net present value assuming discrete compounding at the interest rate r. The other asset's cash flows must be discounted to time zero, hence the $1/(1+r)^{mn_1^*}$ term.

Note that we require the minimum value of m such that Equation (2) is true. For any value less than the minimum value of m, the lower bound exceeds the upper bound and the branch can be fathomed. This says that there must be at least m copies of the economic life asset in the optimal solution.

Writing out the factor, we can rewrite (2) as:

$$\begin{aligned}
c_{n_1^*}\frac{(1+r)^{mn_1^*}-1}{r(1+r)^{mn_1^*}} + \frac{c_{n_2^*}}{(1+r)^{mn_1^*}}\frac{(1+r)^{T-mn_1^*}-1}{r(1+r)^{T-mn_1^*}} &\leq UB \\
\frac{c_{n_1^*}}{r} - \frac{c_{n_1^*}}{r(1+r)^{mn_1^*}} + \frac{c_{n_2^*}}{r(1+r)^{mn_1^*}} - \frac{c_{n_2^*}}{r(1+r)^T} &\leq UB \\
\frac{c_{n_2^*} - c_{n_1^*}}{r(1+r)^{mn_1^*}} &\leq UB + \frac{c_{n_2^*}}{r(1+r)^T} - \frac{c_{n_1^*}}{r}
\end{aligned}$$

$$(1+r)^{mn_1^*} \geq \frac{c_{n_2^*} - c_{n_1^*}}{r\left(UB + \frac{c_{n_2^*}}{r(1+r)^T} - \frac{c_{n_1^*}}{r}\right)}$$

Solving for m results in:

$$m \geq \frac{\ln\left[\frac{c_{n_2^*} - c_{n_1^*}}{r\left(UB + \frac{c_{n_2^*}}{r(1+r)^T} - \frac{c_{n_1^*}}{r}\right)}\right]}{n_1^*(\ln(1+r))} \tag{3}$$

We seek the minimum value of m such that Equation (3) is true, noting that m must be integer. Thus, the minimum number of times an asset is utilized at its economic life in a finite horizon equipment replacement problem with stationary costs is:

$$m_{n_1^*} = \max\left\{\left\lceil \frac{\ln\left[\frac{c_{n_2^*} - c_{n_1^*}}{r\left(UB + \frac{c_{n_2^*}}{r(1+r)^T} - \frac{c_{n_1^*}}{r}\right)}\right]}{n_1^*(\ln(1+r))} \right\rceil, 0\right\} \tag{4}$$

Once the minimum value of $m_{n_1^*}$ is determined, the arcs which explore values of $m_{n_1^*}$ less than the minimum can be removed. This is equivalent to shortening the horizon T as the minimum number of assets utilized at their economic life is bounded. While this computational advantage is of interest, we focus on the likelihood of $m_{n_1^*} \geq 1$ because it defines the decision to be implemented immediately. We analyze the case of convex costs in the following section and empirically explore the value of $m_{n_1^*}$ in the Section 6. We refer to $m_{n_1^*}$ as m in the following sections.

5 Convex Cost Analysis

As O&M costs generally rise (non-decreasing) and salvage values decline (non-increasing) with age, it is often assumed that the EAC of an asset are convex, or quasiconvex, with respect to service life. (Note that these conditions on O&M costs and salvage values are not sufficient for EAC values to be convex.) In the convex cost case, EAC costs are monotonically decreasing for asset service lives $1, 2, \ldots, n_1^*$ and monotonically increasing for asset service lives $n_1^*+1, n_1^*+2, \ldots, N$. If quasiconvex costs are assumed, the costs are non-increasing followed by non-decreasing.

In the convex cost case, construction of an upper bound which includes the maximum copies of an asset retained at its economic life is straightforward, as the best solution is to fill the remaining periods with an asset having a life equal to that number of periods. This is because any life greater than this age is infeasible and any life less than this age has a higher cost (convex cost assumption).

For simplicity, define $M = \lfloor T/n_1^* \rfloor$ and the EAC of retaining an asset for $T - Mn_1^*$ periods as $c_{n_3^*}$. Noting our UB, we can write Equation (2) as:

$$c_{n_1^*}(^{P/A,r,mn_1^*}) + \frac{1}{(1+r)^{mn_1^*}} c_{n_2^*}(^{P/A,r,T-mn_1^*}) \leq c_{n_1^*}(^{P/A,r,Mn_1^*}) + \frac{1}{(1+r)^{Mn_1^*}} c_{n_3^*}(^{P/A,r,T-Mn_1^*})$$

Solving for m as in the previous section results in:

$$m \geq \frac{\ln\left[\frac{c_{n_2^*}-c_{n_1^*}}{\frac{c_{n_3^*}-c_{n_1^*}}{(1+r)^{Mn_1^*}}-\frac{c_{n_3^*}-c_{n_2^*}}{(1+r)^T}}\right]}{n_1^*(\ln(1+r))} \tag{5}$$

In order for us to guarantee that at least one asset is retained at its economic life, then m must be greater than or equal to 1. We are interested in this case as it is the decision to be implemented immediately and replacement decisions are generally revisited in the future. This translates to Equation (5) being positive (due to rounding up to the nearest integer) such that:

$$\frac{\ln\left[\frac{c_{n_2^*}-c_{n_1^*}}{\frac{c_{n_3^*}-c_{n_1^*}}{(1+r)^{Mn_1^*}}-\frac{c_{n_3^*}-c_{n_2^*}}{(1+r)^T}}\right]}{n_1^*(\ln(1+r))} > 0$$

which means that:

$$c_{n_2^*} - c_{n_1^*} > \frac{c_{n_3^*} - c_{n_1^*}}{(1+r)^{Mn_1^*}} - \frac{c_{n_3^*} - c_{n_2^*}}{(1+r)^T} \tag{6}$$

For the quasiconvex cost case, we know that $0 \leq c_{n_1^*} \leq c_{n_2^*} \leq c_{n_3^*}$. From this, we can examine the two extremes of the value $c_{n_2^*}$. For the case where $c_{n_2^*} \rightarrow c_{n_1^*}$, Equation (6) becomes:

$$0 > \frac{c_{n_3^*} - c_{n_2^*}}{(1+r)^{Mn_1^*}} - \frac{c_{n_3^*} - c_{n_2^*}}{(1+r)^T}$$

As $T \geq Mn_1^*$, the right hand side will never be less than zero if $r > 0$. So in this situation, we cannot guarantee that there is even one asset retained at its economic life in this solution. Given our UB assumption, this should not come as a great surprise, as $c_{n_2^*} \rightarrow c_{n_1^*}$, the lower bound is as low as possible given that there is an integer number of assets retained at their economic life.

The other extreme is the case where $c_{n_2^*} \rightarrow c_{n_3^*}$. In this situation, Equation (6) becomes:

$$c_{n_3^*} - c_{n_1^*} > \frac{c_{n_3^*} - c_{n_1^*}}{(1+r)^{Mn_1^*}}$$

which is true for any positive r value. Thus, in this case, we are guaranteed to have at least one asset retained at its economic life in the optimal solution to the finite horizon problem. This would be re-assuring to a decision-maker with costs of this structure.

6 Empirical Testing of Bound

We performed a number of experiments in order to determine whether the bound produced in (4) can be "useful." That is, will the value of m generally be at least 1. Clearly, the value of m is dependent on the quality of UB. For the following experiments, we compute a number of upper bounds. For all of the upper bounds, we assume that the maximum number of assets retained at their economic lives (age n_1^*) are used in the solution. The remaining required periods of service, $T - \lfloor T/n_1^* \rfloor$, are completed as follows:

UB1: $T - \lfloor T/n_1^* \rfloor$ assets utilized for one period each.

UB2: asset of service life $T - \lfloor T/n_1^* \rfloor$ utilized once.

UB3: two asset combinations, including $T - \lfloor T/n_1^* \rfloor - 1$ service life with an asset for one period, $T - \lfloor T/n_1^* \rfloor - 2$ service life with an asset for two periods, etc.

UB4: FFD solution for remaining $T - \lfloor T/n_1^* \rfloor$ periods.

The following two sections empirically analyze the cases where costs are random and convex. For the random cost case, all of the above upper bounds were calculated while UB2 was utilized for the convex cost case, as it dominates the others.

6.1 Varying T and Random Costs

In our first experiment, we examine the impact of both costs and the horizon length on the minimum value of m. For horizons $T = 20$ through 100, we generated random EAC values according to $U[1, 100]$, $U[1, 1000]$ and $U[1, 10000]$ for an asset with a maximum age of $N = 20$. For each combination of costs and horizon, we generated 100 instances of costs and an interest rate (U[0.05,.0.20]). For each instance, we computed each UB and m as defined in (4).

Table 1: Experimental results for $i > 0$ with varying T and random EAC values.

Cost	T	n_1^*	m	$T - mn_1^*$	$\lfloor T/n_1^* \rfloor$	$m \geq 1$	mn_1^*/T	$m/\lfloor T/n_1^* \rfloor$
100	20	9.32	3.4	5.89	3.75	84.00%	70.55%	80.83%
1000	20	10.19	2.94	7.23	3.58	76.00%	63.85%	71.83%
10000	20	9.75	3.42	7.54	3.99	76.00%	62.30%	71.67%
100	40	10.42	6.61	5.44	6.9	96.00%	86.40%	92.31%
1000	40	11.65	5.6	7.31	6.12	94.00%	81.73%	89.03%
10000	40	10.33	6.66	5.07	6.86	98.00%	87.33%	94.52%
100	60	10.53	10.38	6.27	10.67	99.00%	89.55%	93.79%
1000	60	10.81	10.49	5.52	10.75	98.00%	90.80%	94.48%
10000	60	10.7	8.54	5.54	8.74	98.00%	90.77%	95.40%
100	80	10.68	12.08	6.98	12.6	100.00%	91.28%	95.29%
1000	80	10.78	10.14	8.77	10.78	99.00%	89.04%	92.95%
10000	80	11.43	8.8	12.18	9.92	98.00%	84.78%	88.80%
100	100	11.01	16.56	5.61	16.87	100.00%	94.39%	97.47%
1000	100	10.13	20.1	7.31	20.63	100.00%	92.69%	95.47%
10000	100	11.18	15.17	9.58	15.72	99.00%	90.42%	93.71%

In Table 1, the first two columns (Cost and T) define the experiment. From the 100 trials, the average economic life n_1^*, average minimum m, and average resulting new horizon $T - mn_1^*$ are defined. The value of $\lfloor T/n_1^* \rfloor$ defines the maximum number times an asset can be retained at its economic life.

The final three columns summarize the results. The seventh column defines the percentage of the 100 problems in which $m \geq 1$. Thus, this is the percentage of problems in which at least one asset is retained at its economic life. As seen by the column, this value is over 70% for the small horizon problems and approaches 100% for the long horizon problems. This leads to average reductions in the horizon of over 60% to over 90%. The final column is an average of the ratio of m over $\lfloor T/n_1^* \rfloor$ for each trial, or the percentage of possible economic lives guaranteed to be in the optimal solution.

This data reveals a few facts. First, the number of copies m of the economic life asset grows with respect to the horizon. This is expected as a larger horizon approximates an infinite horizon. Second, unexpectedly, the spread of costs has little impact on m. It was expected that higher values of m would result with a larger spread of costs as there would not be a concentration of possible near-optimal solutions.

Given the results of this experiment, it appears that a manager can feel fairly confident in using the first asset in a sequence of replacements at its economic life. We examine the convex cost case next.

6.2 Varying Economic Life and Convex Costs

Our second experiment assumes that EAC values are convex in the age of the asset. For this experiment, we defined the economic life of the asset and then generated a set of convex costs under the assumption that costs between ages were differentiated according to $U[1, 1000]$. Note that we altered the spread of the costs, as in the previous experiment, but did not see any significant variation in results. For each economic life, we generated a random horizon $U[20, 100]$ and solved for the value of m as before. The interest rate was generated as previously. For each possible economic life between 2 and 20, inclusive, we generated 100 trials.

Table 2 summarizes our results in the same format as the previous experiment. As expected, an increase in the economic life of an asset led to a decrease in the possibility that it would be in the optimal solution to the finite horizon problem. However, the percentage of time that $m \geq 1$ is over 79% for all situations, which is much greater than the random cost case.

7 The Zero Interest Rate Case: Integer Programming

Consider the case of the interest rate being zero. This assumption is justified when interest rates are low and/or the time horizon is short. In this situation, we are not concerned with the time value of money and the sequencing of assets is not required in REP because any sequence of the same assets results in the same net present value of costs. Thus, to solve REP, we only need to find the combination of assets that results in minimum costs while ignoring their sequence. This eliminates the need to discount and thus eliminates the non-linearity of the problem. We can model

Table 2: Experimental results for $i > 0$ with varying n_1^* and convex EAC values.

T	n_1^*	m	$T - mn_1^*$	$\lfloor T/n_1^* \rfloor$	$m \geq 1$	mn_1^*/T	$m/\lfloor T/n_1^* \rfloor$
59.44	2	29.28	0.88	29.48	100.00%	98.08%	99.00%
60.11	3	19.14	2.69	19.69	100.00%	95.28%	97.43%
60.33	4	13.93	4.61	14.7	100.00%	91.78%	94.62%
60.52	5	11.22	4.42	11.72	100.00%	91.38%	94.79%
59.01	6	8.44	8.37	9.42	98.00%	85.35%	90.05%
55.99	7	6.96	7.27	7.57	99.00%	85.19%	90.95%
55.86	8	5.49	11.94	6.52	95.00%	75.52%	82.12%
61.63	9	5.76	9.79	6.39	97.00%	80.07%	87.24%
59.87	10	4.6	13.87	5.52	93.00%	73.83%	80.23%
57.27	11	3.83	15.14	4.79	92.00%	70.89%	77.84%
62.49	12	3.97	14.85	4.73	95.00%	72.11%	81.21%
58.66	13	3.35	15.11	4.1	91.00%	68.58%	75.61%
59.06	14	2.84	19.3	3.68	87.00%	60.53%	71.30%
59.41	15	2.6	20.41	3.45	87.00%	60.61%	69.95%
58.07	16	2.37	20.15	3.14	84.00%	59.64%	71.20%
61.58	17	2.22	23.84	3.16	79.00%	55.66%	63.75%
58.6	18	2.07	21.34	2.79	79.00%	57.76%	68.53%
60.33	19	2.04	21.57	2.74	86.00%	60.60%	71.62%
57.17	20	1.86	19.97	2.39	80.00%	59.71%	71.83%

this situation with integer programming.

Define T as the length of the horizon and the variable x_i as the number of times an asset is kept in service for i consecutive periods within T. Further define c_i as the EAC for retaining the asset for i periods, with N being the maximum number of periods that an asset can be retained. With these definitions, the formulation IP follows:

$$\min \sum_{i=1}^{N} ic_i x_i$$

subject to:

$$\sum_{i=1}^{N} ix_i = T \qquad (7)$$

$$x_i \in \{0, 1, \ldots, N\} \qquad (8)$$

The formulation above goes by a number of names in the literature, including the equality knapsack problem (Lee 1997) and knapsack-partitioning problem (Johnson 1980b). The objective function minimizes the costs of using a number of assets for their respective service lives i with the total required service of T defined in Constraint (7). The coefficients in Constraint (7) define the master-partitioning problem (Johnson 1980b) or the cyclic group problem (Gomory 1965, Johnson 1980a). As illustrated in Johnson (1980a), the group problem is interesting because as the value of the right hand side is increased, the solution is cyclic. This cyclic nature is of interest here as it relates to the infinite horizon solution.

As $T \to \infty$, the solution is obvious in that the variable x_i with the lowest annual cost $ic_i/i = c_i$ is set to infinity. In the equipment replacement literature, this corresponds directly with the economic life of an asset. In an infinite horizon problem, the solution is to repeatedly replace an asset at its economic life, or the age which minimizes annual equivalent costs, c_i. We have previously defined $\min_i c_i = c_{n_1^*}$.

In the finite horizon problem, the linear programming relaxation of IP mimics the economic life solution in that the variable $x_{n_1^*}$, with minimum value of $c_{n_1^*}$, is set to its maximum value. Clearly, we can add the following upper bounds:

$$x_i \leq \left\lfloor \frac{T}{i} \right\rfloor$$

such that the linear programming relaxation will set $x_{n_1^*}$ to $\lfloor T/n_1^* \rfloor$. The remaining amount $T - \lfloor T/n_1^* \rfloor$ is taken by $x_{n_2^*}$ with the next lowest value c_i value, $c_{n_2^*}$. This continues until the knapsack is "filled" and can result in a non-integer solution (although only one variable will be non-integer).

In the context of this paper, we are interested in the lower bound on $x_{n_1^*}$. Following our reasoning from Section 4, we can define an UB to IP with any feasible solution. Again, we construct a feasible solution by setting $x_{n_1^*} = \lfloor T/n_1^* \rfloor$, its upper bound, and taking $T - \lfloor T/n_1^* \rfloor$ copies of x_1, an asset retained for a single period, or one copy of an asset retained for $T - \lfloor T/n_1^* \rfloor$ periods. Other UBs can be constructed as before.

We define the lower bound here as we did with DP, in that $0, 1, \ldots, \lfloor T/n_1^* \rfloor$ copies of $x_{n_1^*}$ are utilized, with the remaining periods "filled" with the second best asset. Determining the minimum number of copies of $x_{n_1^*}$ is the minimum m such that:

$$mn_1^* c_{n_1^*} + (T - mn_1^*) c_{n_2^*} \leq UB. \tag{9}$$

Solving for m results in:

$$m \geq \frac{UB - Tc_{n_2^*}}{n_1^*(c_{n*_1} - c_{n_2^*})}$$

As there may be non-integer and negative values, we formally define m as:

$$m = \max\left\{ \left\lceil \frac{UB - Tc_{n_2^*}}{n_1^*(c_{n*_1} - c_{n_2^*})} \right\rceil, 0 \right\} \tag{10}$$

As the denominator in (10) is negative, by definition, the value of m is improved with the quality of UB. Furthermore, large values of T and small values of n_1^* improve the bound. We specifically analyze m in the following section under the assumption of convex costs.

8 Convex Cost Analysis

Assuming convex costs , the lowest upper bound solution that can be created when the maximum copies of the asset retained at its economic life is to purchase an asset with service life $T - Mn_1^*$, where $M = \lfloor T/n_1^* \rfloor$. Noting this, we can re-examine Equation (9) such that:

$$mn_1^* c_{n_1^*} + (T - mn_1^*) c_{n_2^*} \leq Mn_1^* c_{n_1^*} + (T - Mn_1^*) c_{n_3^*}$$

Solving for m results in:

$$m \geq \frac{Mn_1^*(c_{n_1^*} - c_{n_3^*}) + T(c_{n_3^*} - c_{n_2^*})}{n_1^*(c_{n_1^*} - c_{n_2^*})} \tag{11}$$

As any value of m can be rounded up to the next integer, we know that in order to guarantee that at least one asset is retained at its economic life, then:

$$\frac{Mn_1^*(c_{n_1^*} - c_{n_3^*}) + T(c_{n_3^*} - c_{n_2^*})}{n_1^*(c_{n_1^*} - c_{n_2^*})} > 0$$

such that:

$$\frac{T}{mn_1^*} < \frac{c_{n_1^*} - c_{n_3^*}}{c_{n_2^*} - c_{n_3^*}} \tag{12}$$

With quasiconvex costs, we know that $0 \leq c_{n_1^*} \leq c_{n_2^*} \leq c_{n_3^*}$. For the case where $c_{n_2^*} \to c_{n_1^*}$, then Equation (12) becomes:

$$\frac{T}{mn_1^*} < 1$$

As $T \geq Mn_1^*$, the left hand side will always be greater than or equal to one. So, as the case with discounting, we cannot guarantee that there is even one asset retained at its economic life in the optimal solution.

For the case where $c_{n_2^*} \to c_{n_3^*}$, Equation (12) becomes:

$$\frac{T}{mn_1^*} < \lim_{c_{n_2^*} \to c_{n_3^*}} \frac{c_{n_1^*} - c_{n_3^*}}{c_{n_2^*} - c_{n_3^*}} = \infty,$$

In this case, we can guarantee than at least one asset will be retained at its economic life in the optimal solution to the finite horizon problem. Thus, if the economic life of an asset is "unique" in that $c_{n_2^*}$ does not approach $c_{n_1^*}$, then we should be confident in our initial decision to purchase an asset and retain it for its economic life.

9 Empirical Testing of Bound

We repeat the experiments presented in Section 6 assuming the interest rate is zero. The value of m is calculated using (10) for the random cost case and the integer value of (11) for the convex cost case. Upper bounds were determined similarly as in the previous experiments.

9.1 Varying T and Random Costs

Our first experiment with the interest rate being zero follows that of the first experiment in Section 6. The results are given in Table 3. One would expect the results to be worse for the case of no interest, as discounting lessens the effect of the costs of assets utilized at the end of the horizon. For our upper bound, this entails the final $T - Mn_1^*$ periods.

Table 3: Experimental results for $r = 0$ with varying T and random EAC values.

Cost	T	n_1^*	m	$T - mn_1^*$	$\lfloor T/n_1^* \rfloor$	$m \geq 1$	mn_1^*/T	$m/\lfloor T/n_1^* \rfloor$
100	20	10.13	2.92	7.17	3.4	75.00%	64.15%	73.33%
1000	20	10.44	3.22	7.58	3.73	71.00%	62.10%	69.17%
10000	20	10.2	2.85	9	3.41	67.00%	55.00%	64.00%
100	40	10.34	6.58	10.65	7.49	88.00%	73.38%	79.55%
1000	40	10.6	4.94	11.57	5.85	85.00%	71.08%	78.24%
10000	40	9.52	6.66	12.48	7.64	81.00%	68.80%	74.62%
100	60	10.76	9.37	10.63	10.01	92.00%	82.28%	86.58%
1000	60	11.38	7.99	11.62	8.67	91.00%	80.63%	84.71%
10000	60	8.9	13.68	10.35	14.38	91.00%	82.75%	84.93%
100	80	10.5	11.84	15.71	13.12	91.00%	80.36%	82.83%
1000	80	10.64	13.05	18.12	14.74	87.00%	77.35%	80.83%
10000	80	10.08	14.27	16.01	15.52	88.00%	79.99%	82.27%
100	100	10.8	13	17.96	14.49	95.00%	82.04%	85.05%
1000	100	11.64	10.38	22.23	12.38	91.00%	77.77%	80.09%
10000	100	10.01	19.05	16.89	20.46	95.00%	83.11%	85.76%

While this is true, as seen in the Table 3, the difference is not as dramatic as expected as the percentage of time that $m \geq 1$ lies in the range of 67% to 95% for each the problem data sets. As with the discounted cost case, it would appear that the decision to keep an asset for its economic life (at least once) still appears valid. Additionally, 64% was the lowest percentage of economic life assets utilized (among possible) in the optimal solution. This should give further reassurances to the decision maker. The other results, including an increase in the value of m with the horizon time T followed as in the previous experiment.

9.2 Varying Economic Life and Convex Costs

For the convex cost case, the results are summarized in Table 4. As we can see with these results, there is a drastic deterioration in the percentage of time that we can guarantee an asset is retained at its economic life if the economic life is closer to the horizon. This is dramatically worse than the case where the interest rate is positive.

Table 4: Experimental results for $r = 0$ with varying n_1^* and convex EAC values.

T	n_1^*	m	$T - mn_1^*$	$\lfloor T/n_1^* \rfloor$	$m \geq 1$	mn_1^*/T	$m/\lfloor T/n_1^* \rfloor$
54.83	2	26.65	1.53	27.18	99.00%	96.91%	97.78%
59.99	3	18.94	3.17	19.68	99.00%	93.46%	95.21%
61.11	4	13.84	5.75	14.88	97.00%	89.78%	92.64%
59.56	5	10.07	9.21	11.56	93.00%	82.89%	85.59%
56.98	6	7.58	11.5	9.12	90.00%	77.75%	81.24%
58.96	7	6.31	14.79	8.01	86.00%	70.88%	74.71%
54.92	8	4.43	19.48	6.43	84.00%	62.76%	67.61%
60.49	9	4.67	18.46	6.35	82.00%	64.22%	68.16%
55.25	10	3.13	23.95	5.06	77.00%	54.13%	59.39%
55.3	11	2.58	26.92	4.56	70.00%	45.62%	52.44%
60.71	12	2.42	31.67	4.6	62.00%	42.72%	47.93%
57.99	13	1.82	34.33	4.03	59.00%	37.63%	41.45%
58.02	14	1.64	35.06	3.72	48.00%	33.81%	37.00%
57.54	15	1.21	39.39	3.37	40.00%	25.07%	27.23%
59.04	16	1.07	41.92	3.18	37.00%	24.72%	28.83%
62.68	17	0.89	47.55	3.18	35.00%	20.74%	24.50%
57.27	18	0.71	44.49	2.71	34.00%	20.18%	24.28%
59.46	19	0.43	51.29	2.65	22.00%	12.66%	15.38%
56.88	20	0.78	41.28	2.33	39.00%	22.86%	27.83%

10 Computational Testing of IP with Bound

We noted in the introduction that the goal of this paper was not to present a new computational approach to the problem, as there are already viable approaches available. However, there may be some computational benefit to computing the value of m and reducing the search space. We have already described how the value of m can fathom branches in the DP network. Here, we analyze the ability to reduce the computational effort of solving the IP.

We performed two sets of experiments. For the case of random costs, we generated costs for asset ages one through 20 assuming $U[1, 50]$. Note that as we are testing the benefits for IP, the costs are not discounted. For each of 100 instances, we computed the value of m using our upper bounds as described earlier. If the value of m was greater than zero, we calculated the new horizon T as $T - mn_1^*$. We generated an integer program for the original T value and the new value of T and solved it using LINDO 6.1 (1985) on a Silicon Graphics Octane Workstation. For each instance, we recorded the number of simplex iterations and branches evaluated in the branch and bound tree, as reported by LINDO.

We repeated the 100 trial test under the assumption of convex costs. The economic life was randomly generated and then convex costs, with random intervals of $U[1, 50]$, were generated. The results of both experiments are given in Table 5.

For the random cost data instances, the average number of simplex iterations required to solve a problem was 29.18 with 6.29 branches. After computing m, the average number of iterations was reduced to 17.49 while the average number of branches decreased to 2.51. The first section of Table 5 shows a breakdown of the 100 instances. Of the 100 random data instances, the number of simplex iterations was reduced in 76 of the problems, while it increased for 9 and remained the same for 15. When examining branch reduction, 40 of the problems were reduced while none got worse.

For the convex cost data, the average number of iterations was 68.91 and 15.35 branches. These figures reduced to 62.68 and 12.96, respectively, after computing m. For these 100 problems, 53 instances showed a reduction in the number of iterations while 23 saw a reduction in the number of branches. Nine instances saw an increase in the number of iterations and three had an increase in the number of branches evaluated.

The remaining data in Table 5 provides the reduction (for improved instances) or increase (for worse cases) in the number of iterations and branches, respectively, as a percentage of those in the original problem (original value of T). For each of these cases, the average value of m was computed in addition to the average value of $m/\lfloor T/n_1^* \rfloor$, as this is the percentage of the maximum number of economic lives that are included in the optimal solution. There is clearly a correlation between this percentage and whether there was an improvement in the computation. Thus, while we were interested in determining whether $m \geq 1$ for decision-making purposes, achieving computational benefits requires that $m \rightarrow \lfloor T/n_1^* \rfloor$. This would probably require the use of more sophisticated upper bounding procedures than the ones utilized in this paper.

Unfortunately, the bound did not prove as useful as desired for the convex cost case, which is

Table 5: Results (simplex iterations and branches in the branch and bound tree) from solving 100 instances of IP with m assuming random and convex costs.

		Random		
Iterations	Instances	Avg. Reduction	Avg. m	$m/\lfloor T/n_1^* \rfloor$
Improved	76	71.9%	7.26	93.8%
Worse	9	(56.0%)	7.11	65.9%
Same	15	–	4.33	28.4%
Branches				
Improved	40	47.7%	6.3	90.5%
Worse	0	–	–	–
Same	60	–	7.15	75.5%
		Convex		
Iterations	Instances	Avg. Reduction	Avg. m	$m/\lfloor T/n_1^* \rfloor$
Improved	53	77.8%	13.88	89.9%
Worse	9	(30.4%)	5.55	59.2%
Same	38	–	0.73	0.05%
Branches				
Improved	23	81.7%	10.26	74.8%
Worse	3	(58.4%)	3	45.6%
Same	74	–	7.69	49.0%

clearly the more challenging case to be solved with integer programming. For the convex cost case, 30 instances required more than 10 branches to be evaluated when compared to only 15 cases for the random data. The bound reduced the number of branches to evaluate in 11 of the 30 convex problems and 11 of the 15 random cost problems.

While the statistics are not overly impressive, more than half of the problem instances show computational improvements, for both the random and convex cost cases, when the bound on m was utilized. Furthermore, the amount of improvement in these cases is substantial. It is noted that the computation of m is not included in these results. However, for the convex cost case, the value is defined in closed form while computing the upper bounds for the random case is quite straightforward.

11 Conclusions and Directions for Future Research

This paper has presented a new dynamic programming approach to the finite horizon equipment replacement problem with stationary costs. Although the problem has been solved previously with formulations by Bellman and Wagner, this model was introduced in order to examine the relationship between the infinite horizon solution (to replace an asset continuously at its economic life) and the finite horizon solution. The presented integer knapsack approach, with a non-linear objective function to capture discounting and the sequencing of assets, was used to bound the minimum number of times an asset would be utilized at its economic life in the finite horizon solution.

This bound is useful for a number of reasons. First, determining the bound provides some insight as to whether an infinite horizon solution can be applied in a finite horizon setting. Second, if the bound is shown to be at least one, then the optimal time zero decision is known as the decision to keep an asset at its economic life is implemented immediately. Third, the bound may be used to reduce the computational burden of solving the associated dynamic program. Although the bound can be constructed using any upper bound to the problem, we provide a closed form solution under the assumption of convex costs and analyze it, providing conditions where we can guarantee that at least one asset is retained at its economic life in the optimal solution of the finite horizon problem.

For the case where the interest rate is zero, we illustrate that the model can be solved as an integer program with the constraint defined as that of a group problem. Here, the bound carries the same insight as with the dynamic program and is shown to improve the computation of the IP in a number of generated instances. Again, the bound can be produced from any feasible solution and we provide and analyze a closed form solution for the bound under the assumption of convex costs.

Previous work of Wagner (1975) has been shown to be extremely efficient in handling various, important modeling nuances in replacement analysis, such as multiple challengers and technological change. It is unclear as to whether this formulation can handle these nuances as sequencing asset service lives would not be trivial. Regardless, it is believed that the work here may be useful for other types of integer knapsack models which exhibit non-linear returns, such as models with economies of scale where additional items of the same type cost less than the previous item.

12 Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. DMI-9984891.

References

ANDONOV, R., V. POIRRIEZ AND S. RAJOPADHYE. 2000. Unbounded knapsack problem: Dynamic programming revisited. *European Journal of Operational Research*, 123,394-407.

BEAN, J.C., J.R. LOHMANN AND R.L. SMITH. 1985. A dynamic infinite horizon replacement economy decision model. *The Engineering Economist*, 30,99-120.

BEAN, J.C., J.R. LOHMANN AND R.L. SMITH. 1994. Equipment replacement under technological change. *Naval Research Logistics*, 41,117-128.

BELLMAN, R.E. 1955. Equipment replacement policy. *Journal of the Society for the Industrial Applications of Mathematics*, 3,133-136.

BELLMAN, R.E. 1957. *Dynamic Programming.* Princeton University Press, Princeton, NJ.

ESCHENBACH, T.G. 1995. *Engineering Economy, Applying Theory to Practice.* Irwin, Chicago, IL.

FLEISCHER, G.A. 1994. *Introduction to Engineering Economy.* Prentice-Hall, Inc., Boston, MA.

GILMORE, P.C. AND R.E. GOMORY. 1966. The theory and computation of knapsack functions. *Operations Research*, 14,1045-1074.

GOMORY, R.E. 1965. On the relation between integer and noninteger solutions to linear programs. *Proceedings of the National Academy of Science*, 53,260-265.

Lindo Systems Inc. 1985 Lindo (Linear, Interactive, Discrete Optimizer), Version 6.1, Chicago, IL.

JOHNSON, E.L. 1980. *Integer Programming: Facets, Subadditivity and Duality for Group and Semi-group Problems.* Society for Idustrial and Applied Mathematics, Philadelphia, PA.

JOHNSON, E.L. 1980. Subadditive lifting methods for partitioning and knapsack problems. *Journal of Algorithms*, 1,75-96.

E.K. Lee. On facets of knapsack equality polytopes. 1997. *Journal of Optimization Theory and Applications*, 1,223-239.

MORIN, T.L. AND R.E. MARSTEN. 1976. Branch-and-bound strategies for dynamic programming. *Operations Research*, 24,611-627.

OAKFORD, R.V., J.R. LOHMANN AND A. SALAZAR. 1984 A dynamic replacement economy decision model. *IIE Transactions*, 16,65-72.

PARK, C.S. 1997. *Contemporary Engineering Economics.* Addison-Wesley, Menlo Park, CA, Second edition.

PARK, C.S. AND G.P. SHARP-BETTE. 1990. *Advanced Engineering Economics.* John Wiley and Sons, New York.

THUESEN, G.J. AND W.J. FABRYCKY. 1994. *Engineering Economy.* Prentice-Hall, Inc., Upper Saddle River, NJ, Eighth edition.

WAGNER, H.M. 1975. *Principles of Operations Research.* Prentice-Hall, Inc., Englewood Cliffs, NJ.

WAGNER, H.M. AND T.M. WHITIN. 1958. Dynamic version of the economic lot size model. *Management Science*, 5,89-96.