Journal of Global Optimization manuscript No.
(will be inserted by the editor)

Solving Infinite Horizon Optimization Problems Through Analysis of a One-dimensional Global Optimization Problem

**Seksan Kiatsupaibul · Robert L. Smith ·
Zelda B. Zabinsky**

Received: date / Accepted: date

Abstract Infinite horizon optimization (IHO) problems present a number of challenges for their solution, most notably, the inclusion of an infinite data set. This hurdle is often circumvented by approximating its solution by solving increasingly longer finite horizon truncations of the original infinite horizon problem. In this paper, we adopt a novel transformation that reduces the infinite dimensional IHO problem into an equivalent one dimensional optimization problem, i.e., minimizing a Hölder continuous objective function with known parameters over a closed and bounded interval of the real line. We exploit the characteristics of the transformed problem in one dimension and introduce an algorithm with a graphical implementation for solving the underlying infinite dimensional optimization problem.

Keywords Infinite horizon optimization · Dynamic programming · Nonlinear programming · Hölder and Lipschitz continuous functions

1 Introduction

Decision making over an unbounded horizon presents several challenges to their formulation and solution. These include the necessity to forecast data over an infinite horizon, and thereby face a computational challenge of dealing with an infinite data set. One way to circumvent this issue is to simply assume that today's world is tomorrow's world, which reduces the infinite horizon problem to a one period problem via a dynamic programming recursion as in Markov decision process models. Alternatively, a finite data set that incorporates time dependent

S. Kiatsupaibul
Chulalongkorn Business School, Chulalongkorn University, Bangkok 10330, Thailand
E-mail: seksan@cbs.chula.ac.th

R. L. Smith
Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109, USA

Z. B. Zabinsky
Department of Industrial and Systems Engineering, University of Washington, Seattle, Washington 98195, USA

data associated with a long finite horizon still introduces a challenging forecasting and computational task, which realistically can only reliably deliver a limited finite horizon look ahead into the future. One can then attempt a planning horizon approach by solving a sequence of ever longer horizon problems in an attempt to approximate the next best decision that one must implement now [2,4,5,14]. These successive finite horizon problems are typically solved as dynamic programs [1,3,15]. However when the sequential decisions are strongly correlated (in cost or feasibility) the state space can rapidly grow to unmanageable sizes [9,11,16,18].

We adopt a novel approach in this paper to the task of finding the best initial decision for the infinite horizon planning problem by transforming it into a single one dimensional optimization problem. The resulting problem, although now finite dimensional, typically has multiple local optima and is thus a global optimization problem. However, we show the objective function is Hölder continuous and both of its bounding parameters are analytically given by easily determined parameters of the problem data. Moreover the feasible region is simply a closed and bounded interval of the real line. For an earlier attempt in this direction, see [10]. We exploit the characteristics of this transformed one dimensional problem to introduce an algorithm with a graphical technique that is guaranteed to find the optimal first decision to the infinite horizon problem whenever that decision is unique. Although hard to verify, this is almost always the case (see e.g., [13]).

Since our transformation to a one dimensional optimization problem is invertible, we have the opportunity to solve the original infinite horizon optimization problem by solving a one dimensional optimization problem whose solution can be transformed back to a solution of our original infinite dimensional optimization problem. This interesting and deep connection between infinite horizon optimization and the bounding parameters for a Hölder continuous function suggests the potential to more efficiently recover the next optimal decision within the infinite horizon optimization problem.

2 Mathematical Model

To study the infinite horizon problem, we consider, for purposes of exposition, the case of an infinite sequence of binary decisions {0,1}, although the theory is easily extended to arbitrary p-valued decisions. A familiar problem is the equipment replacement problem, where the binary decision is to keep or replace a piece of equipment over an infinite time horizon. A four-valued decision may include more options, such as: keep and do nothing, keep and provide maintenance, repair, or replace. Although this is an easily solved problem via dynamic programming when the cost and feasibility of choice of asset to replace is independent of past acquired assets, so that the age of the current asset represents the state variable, more complex binary decision problems can present formidable challenges to a dynamic programming formulation; in the worst case, the cost and feasibility of the current decision could depend on the entire history of previous decisions, and therefore, the state space becomes the nodes of an infinite binary tree.

In this case the number of states increases exponentially fast as one searches deeper into the tree and moreover a dynamic programming solution approach becomes pure enumeration of all decision sequences since all aspects of previous decisions become relevant in assessing the optimality of future decision alternatives. Our approach from this point of view may be seen as a kind of branch and bound approach for finding the best path in an infinite binary tree of decision sequences. Challenging aspects here include that one cannot fathom an incumbent path in finite time and strong bounds for pruning challenger paths are difficult to come by. We use the analytical characteristics of the underlying Hölder continuous function of the transformed one dimensional optimization problem to address both of these challenges.

To formalize this problem, we define the space of all decision strategies, $\mathcal{S}$, as an infinite product space of {0,1}, i.e.,

$$\mathcal{S} = \prod_{n=1}^{\infty} \{0, 1\},$$

and, for each strategy $s \in \mathcal{S}$, s is an infinite sequence of 0's and 1's, denoted,

$$s \equiv (s_n) \qquad \text{where } s_n \in \{0, 1\}, \text{ for all } n \in \mathbb{N}.$$

Throughout, $\mathbb{N}$ denotes the set of natural numbers, and $\mathbb{R}$ denotes the set of real numbers. We refer to $n \in \mathbb{N}$ as the period n, and refer to s_n as the decision for period n.

We create a metric space $(\mathcal{S}, d)$, where the metric d is defined as

$$d(s,t) = \sum_{n=1}^{\infty} \frac{|s_n - t_n|}{2^n}, \quad \text{for all } s, t \in \mathcal{S}.$$

This metric induces the product topology of component-wise convergence on $\mathcal{S}$. That is, the sequence $\{s^k\}$ converges to t if and only if the real number sequence $\{s^k_n\}$ converges to t_n, for all n, as k goes to infinity. By definition, we also have that $\{s^k\}$ converges to t if and only if $d(s^k, t)$ goes to zero as k goes to infinity. Under this topology, $\mathcal{S}$ is compact by the Tychonoff theorem [2]. Let the set of feasible strategies be denoted S, which is assumed to be a nonempty closed subset of $\mathcal{S}$. Note that S is compact since $\mathcal{S}$ is compact.

The *undiscounted cost function* for period n associated with strategy s is denoted by $c(s, n)$, where $c : S \times \mathbb{N} \to \mathbb{R}$. The cost $c(s, n)$ represents the undiscounted cost incurred in period n by following strategy s. The *total discounted cost* for strategy s is denoted $\tilde{c}(s)$, where $\tilde{c} : S \to \mathbb{R}$, and is given by

$$\tilde{c}(s) = \sum_{n=1}^{\infty} \frac{c(s,n)}{(1+r)^n}, \tag{1}$$

where $r > 0$ is the rate of interest.

We impose two assumptions on the undiscounted cost function.

1. For each $s \in S$ and $n \in \mathbb{N}$, $c(s, n)$ is deterministic and depends only on s_1 to s_n; it does not depend upon s_k where $k > n$. Consequently, for any $s, t \in S$, if $s_k = t_k$, for all $k = 1$ to n, then

$$c(s, n) = c(t, n).$$

2. For each $s \in S$ and $n \in \mathbb{N}$, $c(s,n)$ grows at most geometrically fast in n at a rate of growth γ less than r, i.e.,

$$|c(s,n)| \leq B(1+\gamma)^n, \quad \text{for all } s \in S, \text{ for all } n \in \mathbb{N}, \tag{2}$$

where $B > 0$, and $-1 < \gamma < r$.

The two assumptions imply that the total discounted cost $\tilde{c}(s)$ is finite for all $s \in S$. They also imply that the total discounted cost $\tilde{c}(s)$ is continuous over S. The assumptions leading to continuity of $\tilde{c}(s)$, and the compactness of S allow us to state an optimization problem with the existence of an optimal solution.

The class of infinite horizon optimization (IHO) problems we consider is,

Program 1

$$\min \tilde{c}(s)$$

$$\text{s.t.} \quad s \in S.$$

We denote an optimal strategy by s^*, and the nonempty set of all optimal solutions to Program 1 by S^*.

We show how to transform Program 1 into the following one dimensional global optimization problem,

Program 2

$$\min \tilde{f}(y)$$

$$\text{s.t.} \quad y \in [a,b]$$

where a and b are real numbers with $a < b$ and the objective function $\tilde{f}$ is a Hölder function, i.e.,

$$|\tilde{f}(y) - \tilde{f}(z)| \leq M|y-z|^\alpha, \quad \text{for all } y, z \in [a,b] \subseteq \mathbb{R}, \tag{3}$$

where M is a positive real number and α is a real number with $0 < \alpha \leq 1$.

Our transformation provides closed form expressions for M and α in terms of known parameters (specifically, r, B and γ from (1) and (2)). The transformation has an inverse, allowing the set of optimal solutions of the transformed problem to be mapped back to optimal solutions of the original problem. There exist several solution methods in the global optimization literature for Program 2, especially for the subclass of Lipschitz functions where $\alpha = 1$ (see [6], [8], [19], [7], and [17]).

Program 2 can be a tractable global optimization problem, and if not solvable exactly, at least good approximate solution techniques are known. An approximation may be enough, since in practical terms, only the first period decision s_1^* of an optimal strategy is needed now within a rolling horizon framework. Towards this end, we introduce a new algorithm, operating on the one dimensional real number domain, to solve for the optimal first decision of the original infinite dimensional problem. Furthermore, since the transformed problem is a one dimensional global optimization problem on a closed interval, it also allows for graphical representations of the problem and the algorithm, enabling a decision maker to visualize the problem and even solve it graphically. This connection between IHO and global optimization may lead to future research between the two communities.

3 Construction of a Global Optimization Problem Equivalent to the Infinite Horizon Optimization Problem

Starting with the infinite horizon optimization problem as formulated in Program 1, a natural mapping from a feasible strategy s in S to a real number in the interval $[a, b]$, for real numbers $a < b$, is by binary expansion. For example, the strategy $(0, 1, 1, 1, \ldots)$ can be represented in base 2 by $0.0111\ldots_2$, which equals one-half. Formally, we can define a function $\tilde{x} : S \to [0, 1]$ by

$$\tilde{x}(s) = \sum_{n=1}^{\infty} \frac{s_n}{2^n}, \quad \text{for all } s = (s_n) \in \mathcal{S}. \tag{4}$$

Note that $\tilde{x}$ is an onto function, but it is *not* one-to-one, because $(0, 1, 1, 1, \ldots)$ and $(1, 0, 0, 0, \ldots)$ map to the same point, one-half.

We desire a mapping that is one-to-one, so that there exists an inverse mapping. For that reason, we use a mapping x with a base-3 expansion. Since $S \subset \mathcal{S}$, we construct a function x that maps $\mathcal{S}$ into the interval $[0, 1/2]$,

$$x(s) = \sum_{n=1}^{\infty} \frac{s_n}{3^n}, \quad \text{for all } s = (s_n) \in \mathcal{S}. \tag{5}$$

This mapping x, as we next discuss, is now a one-to-one mapping.

Let the image of x over $\mathcal{S}$ be denoted by X. Defined in this way, X is the set of all real numbers in $[0, 1/2]$ that, when represented in the base-3 expansion, have digits that are 0 or 1. We have no 2's in our base-3 expansion because $s_n \in \{0, 1\}$, and thus we do not have to address the ambiguity between $0.1000\ldots_3$ and $0.0222\ldots_3$ (both base-3 expansions of one-third). For example, because our base-3 expansion never ends in a string of 2's, when we refer to the base-3 expansion of one-third, we are referring to $0.1000\ldots_3$. This convention assures us that x is *one-to-one*, and hence invertible. As we saw, had we adopted a binary representation, we would lose this one-to-one property.

The inverse mapping $x^{-1}(y)$ for $y \in X$ is easy to interpret; the kth digit in the 3-adic expansion of y represents the kth period decision in $x^{-1}(y)$. For example, the inverse mapping of $y = 0.1$ is the strategy $s_1 = 1$ and $s_n = 0$ for $n = 2, 3, \ldots$. Even though the mapping x is easy to work with, its range X, the image of x over $\mathcal{S}$, is a special type of set, known in mathematical analysis as a Cantor set, and is difficult to work with. In the following, we extend the range of X to achieve a transformation that is easier to work with.

It is well-known from mathematical analysis that X, as a Cantor set, is compact with Lebesgue measure zero. Also, $x : \mathcal{S} \to X$ is a continuous mapping. To see this, observe that, for s and t in $\mathcal{S}$,

$$|x(s) - x(t)| = \left| \sum_{n=1}^{\infty} \frac{s_n - t_n}{3^n} \right| \le \sum_{n=1}^{\infty} \frac{|s_n - t_n|}{3^n} \tag{6}$$

where the inequality follows from the triangle inequality, and hence

$$\le \sum_{n=1}^{\infty} \frac{|s_n - t_n|}{2^n} = d(s, t).$$

Therefore, if a sequence $\{s^k\}$ converges to t in $\mathcal{S}$, then $\{x(s^k)\}$ converges to $x(t)$ in X as k goes to infinity, implying the continuity of x.

Consider now the set of feasible strategies S and denote the image of x over S by Y, called the *set of feasible solutions* or simply the *feasible region*. The feasible region Y is compact, because S is compact and the image of a compact set under a continuous map is compact. Also, Y has Lebesgue measure zero, because $Y \subset X$ and X has Lebesgue measure zero.

Since the function x defined in (5) is a one-to-one mapping, the inverse function x^{-1} exists. Define an objective function $f : Y \to \mathbb{R}$ by

$$f(y) = \tilde{c}(x^{-1}(y)), \quad \text{for all } y \in Y \tag{7}$$

where $\tilde{c}$ is the total discounted cost of Program 1. This is the first step in our transformation of Program 1, and the following theorem states that the objective function (7) preserves the set of optimal strategies.

Theorem 1 *The optimal set of strategies S^* to Program 1 is equal to an inverse mapping of the set of optimal solutions Y^* to the following problem;* $\min f(y)$ *s.t.* $y \in Y$.

Proof The proof is straightforward because of the construction of the one-to-one and onto mapping x and its inverse function x^{-1}. If $\tilde{c}(s) < \tilde{c}(t)$ for $s, t \in S$, then $f(x(s)) < f(x(t))$ because $f(x(s)) = \tilde{c}(x^{-1}(x(s))) = \tilde{c}(s)$. Thus if s^* is optimal to Program 1, $y = x(s^*) \in Y^*$. Similarly, if $f(y) < f(z)$ for $y, z \in Y$, then $\tilde{c}(x^{-1}(y)) < \tilde{c}(x^{-1}(z))$. Consequently, if $y^* \in Y^*$, then $s = x^{-1}(y^*)$ is optimal to Program 1. Hence, $Y^* = x(S^*)$ and $S^* = x^{-1}(Y^*)$. □

The following theorem establishes the Hölder property of the objective function f over Y.

Theorem 2 *The objective function $f(y) = \tilde{c}(x^{-1}(y))$ is a Hölder function on Y, i.e.,*

$$|f(y) - f(z)| \leq M|y - z|^{\alpha}, \quad \textit{for all } y, z \in Y, \tag{8}$$

where $0 < \alpha \leq 1$, and M is a positive constant. In particular, we may set $M = 4B/(1-\beta)$ and $\alpha = \min\{\log_3(1/\beta), 1\}$ where $\beta = (1+\gamma)/(1+r)$, $0 < \beta < 1$.

Proof Consider $y, z \in Y$, and let $s = x^{-1}(y)$, $t = x^{-1}(z)$. Let k be the first period where the decision differs between s and t. That is, if $k = 0$, then $s_n = t_n$ for all $n = 1, 2, \ldots$; if $k = 1$, then $s_1 \neq t_1$; and for any $k = 2, 3, \ldots$, then $s_n = t_n$ for $n = 1, \ldots, k-1$, and $s_k \neq t_k$.

If $k = 0$, then $y = z$, because $y = \sum_{n=1}^{\infty} s_n/3^n = \sum_{n=1}^{\infty} t_n/3^n = z$, and hence $|f(y) - f(z)| = 0$. Thus (8) holds for any positive M and $0 < \alpha \leq 1$.

If $k \in \{1, 2, \ldots\}$ then

$$|f(y) - f(z)| = |\tilde{c}(s) - \tilde{c}(t)|$$

and because $s_n = t_n$ for $n = 1, \ldots, k-1$, we know by Assumption 1 that $c(s, n) = c(t, n)$ for $n = 1, \ldots, k-1$, thus we have,

$$
\begin{aligned}
&= \left| \sum_{n=k}^{\infty} c(s,n)/(1+r)^n - \sum_{n=k}^{\infty} c(t,n)/(1+r)^n \right| \\
&\leq \left| \sum_{n=k}^{\infty} c(s,n)/(1+r)^n \right| + \left| \sum_{n=k}^{\infty} c(t,n)/(1+r)^n \right| \\
&\leq \sum_{n=k}^{\infty} \left| c(s,n)/(1+r)^n \right| + \sum_{n=k}^{\infty} \left| c(t,n)/(1+r)^n \right|
\end{aligned}
$$

by the triangle inequality, and

$$
\leq 2 \sum_{n=k}^{\infty} B\left((1+\gamma)/(1+r)\right)^n
$$

by Assumption 2, and because $0 < \beta = (1+\gamma)/(1+r) < 1$, we have

$$
\begin{aligned}
&= 2B \sum_{n=k}^{\infty} \beta^n \\
&= 2B\beta^k/(1-\beta).
\end{aligned}
$$

Now, for $k = 1, 2, \ldots,$

$$
|y - z| \geq \frac{1}{2 \cdot 3^k}. \tag{9}
$$

To see this,

$$
\begin{aligned}
|y - z| &= |x(s) - x(t)| \\
&= \left| \sum_{n=k}^{\infty} \frac{s_n}{3^n} - \sum_{n=k}^{\infty} \frac{t_n}{3^n} \right| \\
&\geq \left| \frac{s_k - t_k}{3^k} \right| - \left| \sum_{n=k+1}^{\infty} \frac{s_n - t_n}{3^n} \right| \\
&\geq \frac{1}{3^k} - \sum_{n=k+1}^{\infty} \left| \frac{s_n - t_n}{3^n} \right| \\
&\geq \frac{1}{2 \cdot 3^k}
\end{aligned}
$$

and thus (9) holds. Consequently, we have $|y - z|^\alpha \geq \left(\frac{1}{2 \cdot 3^k}\right)^\alpha$ for all $\alpha > 0$. Therefore, $(2^\alpha)(3^\alpha)^k |y - z|^\alpha > 1$, and thus

$$
|f(y) - f(z)| \leq 2B \left(\frac{\beta^k}{1-\beta}\right) (2^\alpha)(3^\alpha)^k |y - z|^\alpha = \frac{(2)(2^\alpha)B}{1-\beta} (3^\alpha \beta)^k |y - z|^\alpha.
$$

Let α be such that

$$
0 < \alpha \leq \min\{\log_3(1/\beta), 1\}.
$$

Since $\alpha \leq \log_3(1/\beta)$, we have $3^\alpha \beta \leq 1$. Since $\alpha \leq 1$, we have $(2)(2^\alpha) \leq 4$. Hence,

$$
|f(y) - f(z)| \leq \frac{4B}{1-\beta} |y - z|^\alpha
$$

and we conclude that f is a Hölder function, i.e.,

$$|f(y) - f(z)| \leq M|y - z|^{\alpha}$$

with $0 < \alpha \leq \min\{\log_3(1/\beta), 1\}$ and $M = 4B/(1-\beta)$. □

Up to this point, we have transformed the original infinite horizon optimization problem in an infinite dimensional space into a global optimization problem with a Hölder objective function in a one dimensional space. However, our task is not complete because the feasible region Y is a compact subset of a Cantor set, and not the entire closed interval $[a, b]$ as we intended.

To complete the task, we extend the objective function f to the whole interval $[a, b] = [0, 1/2]$ in such a way that the extended function $\tilde{f}$ preserves the same Hölder condition of f, i.e., $\tilde{f}$ satisfies (3) with the same α and M as those of f.

In fact, there are a number of possible extensions that preserve the Hölder condition, see, for example, [12]. For the ease of implementation on a computer with a simple error bound analysis, as shown in Section 4, we choose a piecewise linear extension, defined as follows,

$$\tilde{f}(y) = \begin{cases} f(y) & \text{if } y \in Y, \\ f(y_1) + \frac{y-y_1}{y_2-y_1}(f(y_2) - f(y_1)) & \text{if } y \in [0, 1/2] - Y, \end{cases} \tag{10}$$

where $y_1 = \text{argmin}_{u \in Y}\{|y - u| : u < y\}$ and $y_2 = \text{argmin}_{u \in Y}\{|y - u| : u > y\}$. The points y_1 and y_2 are the nearest adjacent points in Y to y, or the *adjacent original feasible points*. These adjacent original feasible points, y_1 and y_2, exist because Y is compact.

Theorem 3 *The extended function $\tilde{f}$ is a Hölder extension of f over $[0, 1/2]$ preserving the same Hölder condition as in Theorem 2.*

Proof See the Appendix. □

We have now constructed the mathematical program that we want, Program 2, where the objective function is the Hölder extension in (10) and the feasible region is $[0, 1/2]$.

It remains to show that we can recover the set of optimal solutions of Program 1 from that of the transformed problem, Program 2. Let S^* be the optimal set of strategies to Program 1. Recall from, Theorem 1, that

$$x(S^*) = Y^* = \text{argmin}_{y \in Y} f(y) = \text{argmin}_{y \in Y} \tilde{f}(y),$$

where the last equality holds because $\tilde{f} = f$ on Y. Let

$$\tilde{Y}^* = \text{argmin}_{y \in [0,1/2]} \tilde{f}(y)$$

be the set of optimal solutions to Program 2. We have $Y^* \subseteq \tilde{Y}^*$ because the function value of any point in the extended domain is greater than or equal to the function values of its adjacent original feasible points. However, there is a possibility that there exists $\tilde{y}^* \in [0, 1/2]$ and $\tilde{y}^* \notin Y$ that is optimal to Program 2. This occurs only when its adjacent original feasible points, y_1 and y_2 in Y, are both optimal. Therefore, whenever we find an optimal solution to the Program 2, we can always recover an original optimal solution by moving to one of the adjacent original feasible points and taking the inverse function.

Example 1 Consider a stationary cost equipment replacement problem with a one-year-old equipment at the beginning of period one, assuming no maximum physical life. Model this problem by the binary sequences of buy-keep decisions, where 1 represents a "buy/replace" decision and 0 represents a "keep" decision. A stationary cost has the property that the cost of n periods of a strategy s, $c(s,n)$, is the same cost if the strategy starts at a different time. To illustrate, consider strategy $s = (s_1, s_2, s_3, \ldots)$ and strategy $t = (t_1, s_1, s_2, s_3, \ldots)$. If the cost structure is stationary, then

$$\tilde{c}(t) = \frac{c(t,1)}{(1+r)^1} + \sum_{n=2}^{\infty} \frac{c(s,n-1)}{(1+r)^n}. \tag{11}$$

Now suppose for this example that the optimal solution of Program 2, denoted by y^*, is unique, and the first period decision of y^* is to keep the equipment. Then we can write $y^* = (0, y_2^*, y_3^*, \ldots)$, and using the transformation,

$$y^* \in (0.000..._3, 0.011..._3) = (0, 1/6).$$

In contrast, consider a strategy z, where the first period decision is to buy a piece of equipment, that is, $z \in (0.100..._3, 0.111..._3) = (1/3, 1/2)$. Now, given that the first period decision of z is to buy, and by the stationarity of the cost structure, the optimal subsequent decisions are y^*, because starting at period 2 of z is equivalent to starting at period 1 of y^* with a one-year-old piece of equipment. Therefore, if $z_1 = 1$, then the optimal decision is $(1, 0, y_2^*, y_3*, \ldots) \in (0.100..._3, 0.1011..._3)$, and if $z_1 = 1, z_2 = 1$, then the optimal decision is $(1, 1, 0, y_2^*, y_3*, \ldots) \in (0.1100..._3, 0.11011..._3)$.

More generally, the unique optimal solution of $\tilde{f}$ over the open interval

$$\left(\sum_{i=1}^{n} (1/3)^i, \sum_{i=1}^{n} 1/3^i + \sum_{i=n+2}^{\infty} 1/3^i \right)$$

for $n = 1, 2, ...$ is given by $z^* = \sum_{i=1}^{n} (1/3)^i + \sum_{i=n+1}^{\infty} y_i^*/3^i$, representing a sequence of buy decisions for the first n periods, followed by y^*.

An interesting feature of this example is that the function $\tilde{f}$ of this problem has at least a countably infinite number of strict local optimal solutions. However, in practice, we are only interested in the first period strategy which corresponds to accuracy within the first "decimal" (actually, tri-adic) point. This suggests that the analytical expressions for the parameters in the Hölder extension allows for an efficient algorithmic approach that can eliminate intervals with an infinite number of local optima.

A motivation for this transformation from an infinite dimensional problem to a one dimensional Hölder function is that there exist numerical procedures, based on branch-and-bound methods, to optimize Hölder and Lipschitz functions [6–8, 17, 19]. With the extension, we can go one step further to devise a graphical algorithm to visualize and, simultaneously, solve it. In fact, an approximate solution that is accurate to the first decimal place would suffice as it would provide the optimal first period decision, which is the decision we need to make now. In Section 4, we introduce a graphical method to solve the problem for the optimal first period decision.

4 A Graphical Algorithm for Finding the Optimal First Decision and a Numerical Illustration

A simple algorithm, Algorithm 1, that is based on branch and bound is now constructed to guarantee discovery of the optimal first period decision, when it is unique, to the general IHO problem. It should be noted that the subintervals $[0, 1/6] = [0.0\dot{0}_3, 0.0\dot{1}_3]$ and $[1/3, 1/2] = [0.1\dot{0}_3, 0.1\dot{1}_3]$ correspond to the first decision equal to 0 and 1, respectively. This is due to the definition of the transformation in the previous section. In the algorithm, we form a set of sample points Q_i for iteration i, compute the upper bound and the lower bound of the minimum corresponding to each of the two subintervals based on Q_i, then refine Q_i until the upper bound of one subinterval is lower than the lower bound of the other subinterval, announcing the discovery of the first optimal decision. We formalize the algorithm below.

To implement the algorithm on a computer, we need to approximate the cost function by truncating the infinite series corresponding to a horizon T. In the course of Algorithm 1, we opt to extend the truncated horizon T_i at each iteration i, reducing the truncation error as i increases. For iteration i, let the truncated cost be denoted by $\tilde{c}(s, T_i)$ where

$$\tilde{c}(s, T_i) = \sum_{n=1}^{T_i} \frac{c(s, n)}{(1 + r)^n}.$$

The truncation error $|\tilde{c}(s) - \tilde{c}(s, T_i)|$ is then

$$|\tilde{c}(s) - \tilde{c}(s, T_i)| = \left| \sum_{n=T_i+1}^{\infty} \frac{c(s, n)}{(1 + r)^n} \right| \leq \frac{B\beta^{T_i+1}}{1 - \beta}. \tag{12}$$

Denote the truncation error bound by ε_i where

$$\varepsilon_i = \frac{B\beta^{T_i+1}}{1 - \beta}.$$

We next transform the truncated cost function over the infinite sequence domain into the truncated objective function over the real numbers. For $s \in S$, let $y = x(s)$. Define the truncated objective function, with respect to truncation horizon T_i,

$$f^{T_i}(y) = \tilde{c}(x^{-1}(y), T_i).$$

With $f(y) = \tilde{c}(x^{-1}(y))$ being defined as before and by (12), we then have

$$|f(y) - f^{T_i}(y)| \leq \varepsilon_i.$$

It should be noted that $f^{T_i}(y)$ is defined only on the Cantor set. We then extend $f^{T_i}(y)$ to $\tilde{f}^{T_i}(y)$ over the whole interval $[0, 1/2]$ by the piecewise linear extension, $\tilde{f}$, discussed in Section 3. By the property of the piecewise linear extension, for all $y \in [0, 1/2]$,

$$|\tilde{f}(y) - \tilde{f}^{T_i}(y)| \leq \varepsilon_i. \tag{13}$$

Now we construct the upper bound and the lower bound of the minimal objective function value based on the set of sample points Q_i, taking into account the truncation error.

Proposition 1 *Given a set of sample points Q_i, a truncation horizon T_i and its corresponding truncation error bound ε_i, let*

$$\bar{f}_i([a,b]) = \min_{y \in Q_i \cap [a,b]} \tilde{f}^{T_i}(y) + \varepsilon_i, \tag{14}$$

and

$$\underline{f}_i([a,b]) = \min_{y \in Q_i \cap [a,b]} \tilde{f}^{T_i}(y) - \varepsilon_i - M(\Delta_i/2)^\alpha, \tag{15}$$

where Δ_i is the sample spacing size of Q_i. Then $\bar{f}_i([a,b])$ and $\underline{f}_i([a,b])$ are upper and lower bounds of the minimum value of $\tilde{f}$ over the interval $[a,b]$, respectively.

Proof First, we show that (14) provides an upper bound of the minimum. Because $Q_i \cap [a,b] \subset [a,b]$ and (13), we have

$$\min_{y \in [a,b]} \tilde{f}(y) \leq \min_{y \in Q_i \cap [a,b]} \tilde{f}(y) \leq \min_{y \in Q_i \cap [a,b]} \tilde{f}^{T_i}(y) + \varepsilon_i.$$

Now consider the lower bound of the minimum given in (15). Since $\tilde{f}(y)$ is continuous on a closed interval, it attains the minimum at a point $y^* \in [a,b]$. Let q^* be the point in Q_i that is closest to y^*, so

$$|y^* - q^*| \leq \Delta_i/2.$$

By Theorem 2 and Theorem 3,

$$\tilde{f}(y^*) \geq \tilde{f}(q^*) - M(\Delta_i/2)^\alpha.$$

By (13) and because $q^* \in Q_i \cap [a,b]$, we have,

$$\tilde{f}(y^*) \geq \tilde{f}^{T_i}(q^*) - \varepsilon_i - M(\Delta_i/2)^\alpha \geq \min_{y \in Q_i \cap [a,b]} \tilde{f}^{T_i}(y) - \varepsilon_i - M(\Delta_i/2)^\alpha. \square$$

Algorithm 1 is a simple demonstration of how we can employ the transformation to solve the problem. Proposition 2 formally states that the algorithm indeed solves Program 1 for a unique optimal first decision in finite time.

Algorithm 1 Solving an Infinite Horizon Problem by Global Optimization

Require: A sequence of spacing $\{\Delta_i\}_{i=0}^\infty$ such that $\Delta_i \downarrow 0$, and a sequence of truncation horizon $\{T_i\}_{i=0}^\infty$ such that $T_i \uparrow \infty$. Set $i = 0$ and form an initial set of sample points Q_0 with spacing no greater than Δ_0.

1: **while** $\bar{f}_i([0,1/6]) \geq \underline{f}_i([1/3,1/2])$ and $\bar{f}_i([1/3,1/2]) \geq \underline{f}_i([0,1/6])$ **do**

2: Form a new set of sample points Q_{i+1}, a refinement of Q_i, with spacing no greater than Δ_{i+1}. Extend the truncation horizon to T_{i+1}.

3: Set $i = i + 1$.

4: **end while**

5: If $\bar{f}_i([0,1/6]) < \underline{f}_i([1/3,1/2])$, the optimal first decision is 0. On the other hand, if $\bar{f}_i([1/3,1/2]) < \underline{f}_i([0,1/6])$, the optimal first decision is 1.

Proposition 2 *If Algorithm 1 terminates, it delivers the first optimal decision. Conversely, if there is a unique optimal first period decision, Algorithm 1 terminates with the optimal first decision in finite time.*

Proof Suppose Algorithm 1 terminates. If $\bar{f}_i([0,1/6]) < \underline{f}_i([1/3,1/2])$, by construction, the upper bound of the cost incurred by the optimal strategies in the infinite horizon problem with 0 as their first decision is lower than the lower bound of the cost incurred by the optimal strategies with 1 as their first decision. Therefore, the optimal first decision is 0. On the other hand, if $\bar{f}_i([1/3,1/2]) < \underline{f}_i([0,1/6])$ the upper bound of the cost incurred by the optimal strategies with 1 as their first decision is lower than the lower bound of the cost incurred by the optimal strategies with 0 as their first decision. Hence, the optimal first decision is 1.

Suppose there is a unique optimal first period decision. Let

$$\delta = \mid \tilde{f}^*([0,1/6]) - \tilde{f}^*([1/3,1/2]) \mid$$

be the gain from adopting an optimal first period decision over the other one, where $\tilde{f}^*([a,b]) = \min_{y\in[a,b]} \tilde{f}(y)$. Since the optimal first period decision is unique, we have $\delta > 0$.

Without loss of generality, assume that the optimal first period decision is 1. For each i, let

$$\varepsilon_i = \frac{B\beta^{T_i+1}}{1-\beta}.$$

Since $T_i \uparrow \infty$, $\varepsilon_i \downarrow 0$. Therefore, there exists m such that

$$0 < \varepsilon_m < \delta/3.$$

Let $\kappa = \delta - 3\varepsilon_m > 0$. Since $\Delta_i \downarrow 0$, there exists $n > m$ such that the Hölder bound error from the sample points to the remaining points is smaller than $\kappa/2$, i.e.,

$$M(\Delta_n/2)^\alpha < \kappa/2.$$

Let $y_1^* = \mathrm{argmin}_{y\in[1/3,1/2]} f(y)$ and $u \in Q_n \cap [1/3,1/2]$ be the sample point that is closest to y_1^*. Hence,

$$\begin{aligned}
\bar{f}_n([1/3,1/2]) &= \min_{y\in Q_n\cap[1/3,1/2]} \tilde{f}^{T_n}(y) + \varepsilon_n \\
&\le \tilde{f}^{T_n}(u) + \varepsilon_n \\
&< f(y_1^*) + \kappa/2 + \varepsilon_n \\
&= \tilde{f}^*([1/3,1/2]) + \kappa/2 + \varepsilon_n. \qquad (16)
\end{aligned}$$

Let $y_0^* = \mathrm{argmin}_{y\in[0,1/6]} f(y)$ and $y_0 = \mathrm{argmin}_{y\in Q_n\cap[0,1/6]} \tilde{f}^{T_n}(y)$. Hence,

$$\begin{aligned}
\underline{f}_n([0,1/6]) &= \min_{y\in Q_n\cap[0,1/6]} \tilde{f}^{T_n}(y) - \varepsilon_n - M(\Delta_n/2)^\alpha \\
&= \tilde{f}^{T_n}(y_0) - \varepsilon_n - M(\Delta_n/2)^\alpha \\
&\ge (\tilde{f}(y_0) - \varepsilon_n) - \varepsilon_n - M(\Delta_n/2)^\alpha \\
&\ge \tilde{f}^*([0,1/6]) - 2\varepsilon_n - M(\Delta_n/2)^\alpha \\
&> \tilde{f}^*([0,1/6]) - 2\varepsilon_n - \kappa/2. \qquad (17)
\end{aligned}$$

From (16) and (17) and because $\varepsilon_n \le \varepsilon_m$,

$$\begin{aligned}
\underline{f}_n([0,1/6]) - \bar{f}_n([1/3,1/2]) &> \tilde{f}^*([0,1/6]) - \tilde{f}^*([1/3,1/2]) - 3\varepsilon_n - \kappa \\
&= \delta - 3\varepsilon_n - \kappa \\
&\ge \delta - 3\varepsilon_m - \kappa \\
&= 0. \qquad (18)
\end{aligned}$$

Observe that (18) is the terminating condition of Algorithm 1. Therefore, the algorithm terminates at step n with 1 as the optimal first decision. $\square$

To illustrate our algorithm and graphical method, we construct a challenging infinite horizon binary decision tree problem by setting the discounted cost at period n given decisions up to period n as a randomly distributed cost between zero and β^n to emulate a highly unstructured problem with exploding state space. Though satisfying the assumptions required in Section 2, this is a very challenging dynamic programming problem since the cost at period n in general depends on the entire history of decisions prior and up to n, rendering any efficiencies via dynamic programming futile.

The numerical results use the values of $r = 0.1$ for the rate of interest, $\gamma = -0.4$ for the rate of growth, and $B = 1$, resulting in $\beta = 0.5455$, $\alpha = 0.5517$ and $M = 8.8$. The algorithm applied to this problem instance is executed on the R statistical software platform. We set the schedules of Δ_i and T_i as follows,

$$\Delta_i = \frac{1}{10 + 10i}, \quad T_i = \max\{20 + i,\ 61\}.$$

We set the bound of T_i to 61 because the corresponding truncation error bound ε_i will be equal to $2.22e - 16$, which is the machine precision on the R platform.

Figure 1 shows the plots of the transformed objective function together with $\underline{f}([0, 1/6])$ and $\bar{f}([1/3, 1/2])$ at $i = 1, 10, 100$. $\bar{f}([1/3, 1/2])$ is plotted by the dotted horizontal line, while $\underline{f}([0, 1/6])$ is plotted by the solid horizontal line. The algorithm terminates at $i = 10$ when the termination condition is reached, i.e., when $\bar{f}([1/3, 1/2]) < \underline{f}([0, 1/6])$. From the graph, we can conclude that, in this specific problem instance, the optimal solution is located in the interval $[1/3, 1/2]$. This is equivalent to stating that the optimal first period decision to the corresponding infinite horizon problem is 1.

The advantage of this method is that it enables us to graph the cost function of an infinite dimensional problem and use visual inspection to study the problem, as illustrated in Figure 1. From a visual inspection, it can easily be seen that the minimum of the objective function $\tilde{f}$ in the interval $[1/3, 1/2] = [0.1\dot{0}_3, 0.1\dot{1}_3]$ is lower than that in the interval $[0, 1/6] = [0.0\dot{0}_3, 0.0\dot{1}_3]$ without resorting to the lower bound and the upper bound lines. With an ability to graph the objective function, a decision maker can identify the optimal first period decision from visually inspecting the rough shape of the graph early in the course of the algorithm, as shown this example. This leads to a simple approximation procedure. It should be noted that the graphical representation requires both a continuous mapping and continuous extension, which is the main idea of this article.

Another enticing feature of this procedure is that it enables us to graphically drill down to investigate not only the optimal first period decision, but also the second, third and beyond. This can be best illustrated by Figure 2 where we can visually solve for the optimal first, second and third period decisions to this complex infinite dimensional decision problem.

5 Discussion

The development in this paper provides the mapping to a one dimensional problem and algorithmic approach for an infinite horizon optimization problem with binary

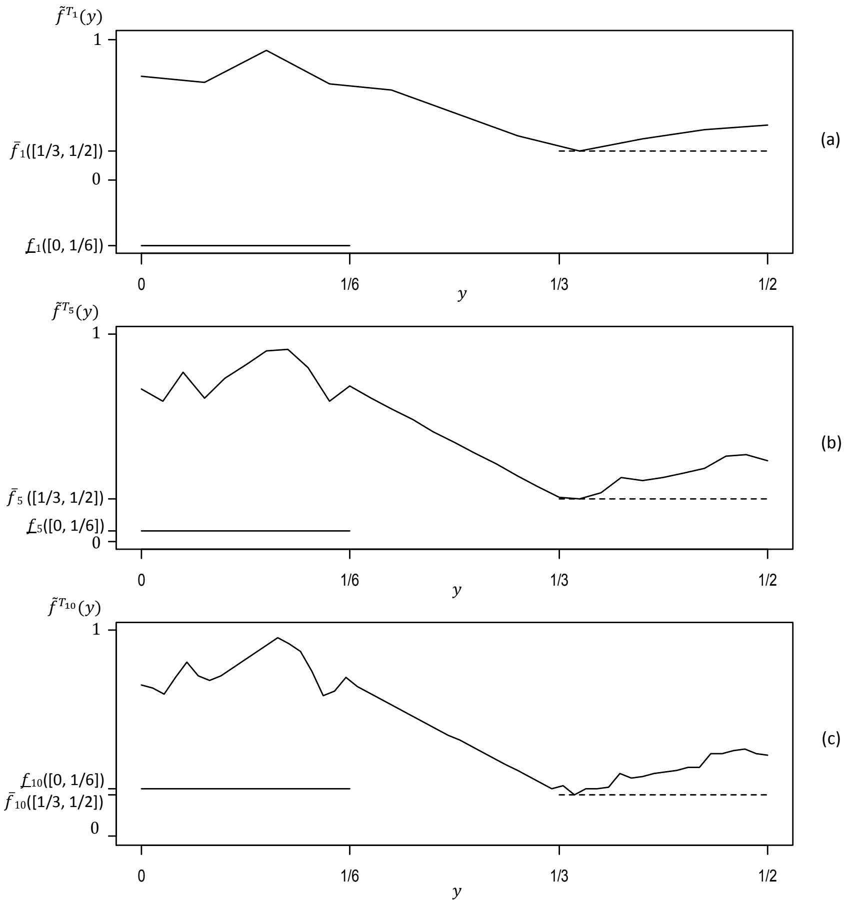

Fig. 1 The development of the graph of the objective function pertaining to the problem instance at different i^{th} iteration in the course of Algorithm 1. (a) $i = 1$, (b) $i = 5$ and (c) $i = 10$. Each plot also shows the corresponding $\bar{f}_i([1/3, 1/2])$ and $\underline{f}_i([0, 1/6])$, in the dotted horizontal line and the solid horizontal line, respectively. The algorithm terminates at $i = 10$ when the solid line is first located higher than the dotted line.

decisions, however the approach is easily extended to p-valued decisions. A p-valued decision may include, for example, different maintenance schedules (keep and do nothing, keep and maintain, or keep and repair), or different replacement schedules (trade-in when the machine is only a few years old, sell as a used machine, or pay to have a worn-out machine removed as scrap).

To generalize to p-valued decisions, the strategy sequence would be modified to $s_n \in \{0, 1, \ldots, p-1\}$, and the mapping in (5) would be modified as

$$x(s) = \sum_{n=1}^{\infty} \frac{s_n}{(p+1)^n}, \quad \text{for all } s = (s_n) \in \mathcal{S}. \tag{19}$$

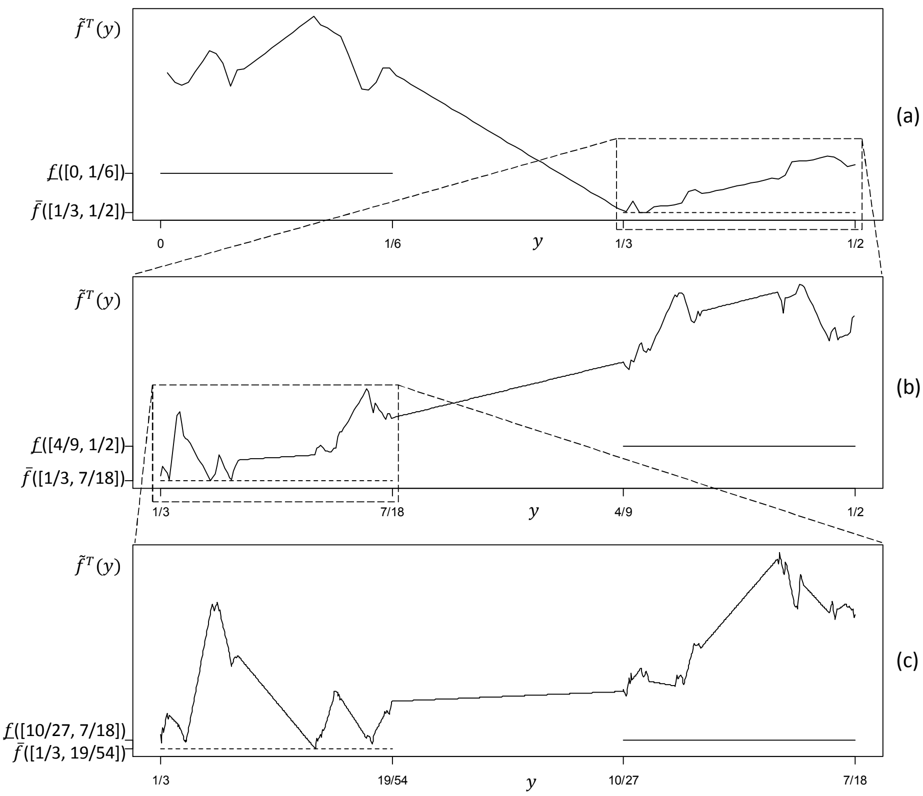

Fig. 2 The magnified objective functions of the global optimization problem. (a) The graph of the objective function over the complete feasible set. It visually suggests the optimal first period decision is 1. (b) The graph of the objective function over the feasible set corresponding to the optimal first decision equal to 1. It visually suggests the second optimal decision is 0. (c) The graph of the objective function over the feasible set corresponding to the optimal first period and second period decisions equal to 1 and 0, respectively. It visually suggests the optimal third period decision is 0.

The analysis extends in a straight-forward manner, with a slight modification in Theorem 2 by letting $\alpha = \min\{\log_{p+1}(1/\beta), 1\}$ with the constant term $M = (2pB)/(1-\beta)$.

6 Conclusion

In this article, we have constructed an equivalence between infinite horizon optimization problems and global optimization problems. We have found a hidden structure between the two optimization problems, leading to a bound of the objective function in the transformed problem that is easy to compute, enabling one to construct a practical algorithm to solve the infinite horizon problem. Specifically, we provide a simple algorithm based on the branch-and-bound algorithm to solve the problem. The branch-and-bound operation on the transformed one dimensional Hölder function implicitly prunes the infinite binary tree in the original state space, narrowing down the choices so that we can finally pin down the optimal first period decision. This relationship also offers the opportunity to explore graphical procedures to approximate the global optimization solution to sequentially find the next optimal decision for the infinite horizon problem.

Appendix

Proof of Theorem 3.

The proof that the piecewise linear extension preserves the Hölder condition can be carried out in a straightforward manner as follows. We are to show that for any pair $x, y \in [0, 1/2]$,

$$|\tilde{f}(x) - \tilde{f}(y)| \leq M|x - y|^{\alpha}. \tag{20}$$

Case 1: if both $x, y \in Y$, (20) follows readily, since $\tilde{f} = f$ by construction of $\tilde{f}$, and f satisfies the Hölder condition by Theorem 2,

$$|\tilde{f}(x) - \tilde{f}(y)| = |f(x) - f(y)| \leq M|x - y|^{\alpha}, \quad \text{for all } x, y \in Y. \tag{21}$$

Case 2: if exactly one of $x, y \notin Y$, WLOG, let $x \in Y$, but $y \notin Y$. Suppose $y > x$. It is easy to check that the function g defined by

$$g(z) = \tilde{f}(x) + M(z - x)^{\alpha} = f(x) + M(z - x)^{\alpha}$$

is a concave function over $z \in [x, \infty)$. Therefore, the hypograph of g on $[x, 1/2]$ is a convex set. Denote this convex set by K. Since $y \notin Y$, there exist $y_1 = \text{argmin}_{u \in Y}\{|y - u| : x \leq u < y\}$, and $y_2 = \text{argmin}_{u \in Y}\{|y - u| : u > y\}$. Since $y_1, y_2 \in Y$, by (21), $(y_1, \tilde{f}(y_1)), (y_2, \tilde{f}(y_2)) \in K$. Since K is convex, $(y, \tilde{f}(y))$, which is the convex combination of $(y_1, \tilde{f}(y_1))$ and $(y_2, \tilde{f}(y_2))$, is also in K. Hence,

$$\tilde{f}(y) \leq g(y) = \tilde{f}(x) + M(y - x)^{\alpha}. \tag{22}$$

Similarly, it is easy to check that the function h defined by

$$h(z) = \tilde{f}(x) - M(z - x)^{\alpha} = f(x) - M(z - x)^{\alpha}$$

is a convex function over $z \in [x, \infty)$. Therefore, the epigraph of h on $[x, 1/2]$ is a convex set. With this fact, applying the same argument, it follows that

$$\tilde{f}(y) \geq h(y) = \tilde{f}(x) - M(y - x)^{\alpha}. \tag{23}$$

By (22) and (23),

$$|\tilde{f}(y) - \tilde{f}(x)| \leq M(y - x)^{\alpha}.$$

Similarly, when $y < x$,

$$|\tilde{f}(y) - \tilde{f}(x)| \leq M(x - y)^{\alpha}.$$

Hence,

$$|\tilde{f}(x) - \tilde{f}(y)| \leq M|x - y|^{\alpha}, \quad \text{for all } x \in Y, \text{ for all } y \notin Y. \tag{24}$$

Case 3: if both $x, y \notin Y$, suppose $y > x$. The function g defined by

$$g(z) = \tilde{f}(x) + M(z - x)^{\alpha}$$

is a concave function over $z \in [x, \infty)$. Therefore, the hypograph of g on $[x, 1/2]$ is a convex set. Denote this convex set by K. Since $y \notin Y$, there exist

$$y_1 = \begin{cases} \operatorname{argmin}_{u \in Y}\{|y-u| : x < u < y\} & \text{if } \{u : x < u < y, u \in Y\} \neq \emptyset, \\ x & \text{if } \{u : x < u < y, u \in Y\} = \emptyset, \end{cases}$$

and $y_2 = \operatorname{argmin}_{u \in Y}\{|y-u| : u > y\}$. Since $y_2 \in Y$ and $y_1 \in Y$ or $y_1 = x$, by (24), we have $(y_1, \tilde{f}(y_1)), (y_2, \tilde{f}(y_2)) \in K$. Since K is convex, $(y, \tilde{f}(y))$, which is the convex combination of $(y_1, \tilde{f}(y_1))$ and $(y_2, \tilde{f}(y_2))$, is also in K. Hence,

$$\tilde{f}(y) \leq g(y) = \tilde{f}(x) + M(y-x)^{\alpha}. \tag{25}$$

Similarly, the function h defined by

$$h(z) = \tilde{f}(x) - M(z-x)^{\alpha}$$

is a convex function over $z \in [x, \infty)$. Therefore, the epigraph of h on $[x, 1/2]$ is a convex set. With this fact, applying the same argument, it follows that

$$\tilde{f}(y) \geq h(y) = \tilde{f}(x) - M(y-x)^{\alpha}. \tag{26}$$

By (25) and (26),

$$|\tilde{f}(y) - \tilde{f}(x)| \leq M(y-x)^{\alpha}.$$

Similarly, when $y < x$,

$$|\tilde{f}(y) - \tilde{f}(x)| \leq M(x-y)^{\alpha}.$$

Hence,

$$|\tilde{f}(x) - \tilde{f}(y)| \leq M|x-y|^{\alpha}, \quad \text{for all } x, y \notin Y. \tag{27}$$

By (21), (24), and (27), the theorem is proved. □

Acknowledgements This work was supported in part by the National Science Foundation under Grant CMMI-1333260.

References

1. Alden, J.M., Smith, R.L.: Rolling horizon procedures in nonhomogeneous markov decision processes. Operations Research **40**, S183–S194 (1992)
2. Bean, J.C., Smith, R.L.: Conditions for the existence of planning horizons. Mathematics of Operations Research **9**, 391–401 (1984)
3. Bean, J.C., Smith, R.L.: Conditions for the discovery of solution horizons. Mathematical Programming **59**(2), 215–229 (1993)
4. Chand, S., Hus, V., Sethi, S.: Forecast, solution and rolling horizons in operations management problem: A classified bibliography. M&SOM **4**(1), 25–43 (2002)
5. Ghate, A.: Infinite horizon problems. In: J.J. Cochran, L.A. Cox Jr., P. Keskinocak, J.P. Kharoufeh, J.C. Smith (eds.) Wiley Encyclopedia of Operations Research and Management Science. John Wiley & Sons (2011)
6. Gourdin, E., Jaumard, B., Ellaia, R.: Global optimization of Hölder functions. Journal of Global Optimization **8**, 323–348 (1996)
7. Hansen, P., Jaumard, B., Lu, S.H.: Global optimization of univariate lipschitz functions: I. Survey and properties, and II. New algorithms and computational comparison. Mathematical Programming **55**(1), 251–272, 273–292 (1992)
8. Horst, R., Pardalos, P.M. (eds.): Handbook of Global Optimization. Kluwer Academic Publishers, London (1995)

9. Huang, K., Ahmed, S.: A stochastic programming approach for planning horizons of infinite horizon capacity planning problems. European Journal of Operational Research **200**(1), 74–84 (2010)
10. Kiatsupaibul, S.: Markov Chain Monte Carlo Methods for Global Optimization. Ph.D. thesis, University of Michigan (2000)
11. Lian, Z., Liu, L., Zhu, S.X.: Rolling-horizon replenishment: Policies and performance analysis. Naval Research Logistics **57**(6), 489–502 (2010)
12. McShane, E.J.: Extension of range of function. Bulletin of American Mathematical Society **40**, 837–842 (1934)
13. Ryan, S.M., Bean, J.C., Smith, R.L.: A tie-breaking algorithm for discrete infinite horizon optimization. Operations Research **40**, S117–S126 (1992)
14. Schochetman, I.E., Smith, R.L.: Infinite horizon optimization. Mathematics of Operations Research **14**, 559–574 (1989)
15. Schochetman, I.E., Smith, R.L.: Finite dimensional approximation in infinite dimensional mathematical programming. Mathematical Programming **34**(3), 307–333 (1992)
16. Sharky, T.C.: Infinite linear programs. In: J.J. Cochran, L.A. Cox Jr., P. Keskinocak, J.P. Kharoufeh, J.C. Smith (eds.) Wiley Encyclopedia of Operations Research and Management Science. John Wiley & Sons (2011)
17. Shubert, B.O.: A sequential method seeking the global maximum of a function. SIAM Journal on Numerical Analysis **9**, 379–388 (1972)
18. Wu, P., Hartman, J.C.: Case study: solving a rental fleet sizing model with a large time-space network. The Engineering Economist **55**, 71–104 (2010)
19. Zhang, B., Wood, G.R., Baritompa, W.P.: Multidimensional bisection: The performance and the context. Journal of Global Optimization **3**(3), 337–358 (1993)