

University of Groningen

On the Optimal Policy for the Single-product Inventory Problem with Set-up Cost and a Restricted Production Capacity

Foreest, N. D. van; Wijngaard, J.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:
2010

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Foreest, N. D. V., & Wijngaard, J. (2010). *On the Optimal Policy for the Single-product Inventory Problem with Set-up Cost and a Restricted Production Capacity*. (SOM Research Reports; Vol. 10005). University of Groningen, SOM research school.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

On Optimal Policies for Production-Inventory Systems with Compound Poisson Demand and Setup Costs

Nicky D. Van Foreest and Jacob Wijngaard
Faculty of Economics and Business
University of Groningen

September 25, 2012

Abstract

In this paper, we consider a single-item, one-machine production-inventory system with compound Poisson demand. The production facility may be in production or idle. While in production, the production rate is constant and positive, and is zero while idle. System costs consist of switching costs and quasi-convex inventory and backlogging costs. We provide conditions when (s, S) -policies are optimal under the long-run average expected cost criterion. These conditions are met in particular when the inventory costs are convex. The developed method in the proof is easy to apply to more general cases. Moreover, the method allows us to compute optimal policies very efficiently.

1 Introduction.

The Economic Production Quantity (EPQ) model is one of the classical production-inventory models. In this continuous-time model, demand is assumed to be deterministic and constant, and production alternates between ‘on’ and ‘off’; if ‘on’, it produces at fixed rate, if ‘off’, it produces nothing. A fixed set-up cost K is incurred every time production is switched on, and linear holding and backlogging costs are accrued for each unit of stock per unit time. In many logistic environments, however, the assumption of constant and deterministic demand is not appropriate. In this paper, we consider a stochastic version of the EPQ model in which demand is compound Poisson rather than deterministic, and the inventory cost function h is quasi-convex rather than linear. The paper has two objectives. The first is theoretical and directed at structure results, namely, to identify conditions on the cost structure and the demand distribution such that (s, S) -policies are optimal under the long-run average expected cost criterion. Thus, when the inventory is below or at s , it is optimal to switch on production, and when the inventory is above or at S , switch off. The other goal is numerical and constructive, namely, to devise a numerical procedure to efficiently compute an optimal policy for general one-product production-inventory systems.

A similar, but simpler, production-inventory problem has been considered earlier by Gavish and Graves [8]. They study the case with unit Poisson demand arrivals, and holding and backlogging costs that are linear in the number of items on stock or in backlog, and derive a numerical procedure to compute optimal switching levels. This approach depends critically on two assumptions. The first is that the Poisson demands arrive as single units, so that the inventory process becomes *two-sided skip-free*, that is, skip-free to the right *and* skip-free to the left. This assumption allows Gavish and Graves [8] to use the optimality result of Sobel [22] who proves that a two-critical-number policy (i.e., an (s, S) policy) is optimal for such two-sided skip-free queueing and inventory processes. The second is that the costs are linear in the inventory and backorder level, so that it becomes possible to express the cost functions as a constant times the average inventory (backlog) level. Graves and Keilson [10] extend the model of [8] such that the demand sizes are exponentially distributed rather than deterministic. An immediate consequence is that the inventory is no longer skip-free to the left, and the result of Sobel no longer applies. Thus, Graves and Keilson [10] restrict their search for the optimal policy from the onset to the class of (s, S) -policies, but do not prove that this optimal policy is stationary optimal.

To the best of our knowledge, there are as yet no structure results known for the more general production-inventory systems with compound Poisson demand with generally distributed single demand sizes and quasi-convex inventory costs. Dropping the skip-freeness property rules out the use of Sobel's result. Moreover, the current numerical methods, which require (a substantial amount of) ingenious and analytic work, appear to be incapable to efficiently analyze cost functions other than linear. Thus, the analysis of the more general inventory system requires a new approach to, on the one hand, *prove* structure results for the optimal policy in the class of stationary policies, and, on the other hand, *compute* the optimal policy by *simple* means.

The principal result of this paper is a framework that addresses simultaneously these three challenges, i.e., structure proofs, computational efficiency, and conceptual simplicity. It achieves this by combining three critical elements in a new way. First, as in Graves and Keilson [10] or, more generally, for Markov-modulated stochastic processes, the production-inventory process is defined on two lines to incorporate the state of production as being on or off. These two lines are such that when the state of the production-inventory process (P_t, I_t) is $(1, x)$, production is on and the inventory level is x , and when the state is $(0, x)$ production is off and the inventory level is x . It is quite straightforward to prove that it is optimal at the on-line to remain on until some critical level S is reached. The second element deals with the hard part: to prove that at the off-line it is optimal to only switch on at or below some level s . To establish this, we formulate the switching decision at the off-line as an optimal stopping problem to minimize the g -revised holding cost rather than just the holding costs, and prove that there exists a g -revised optimal policy for general $g > 0$. The third element is to use a bisection on g to construct a sequence of optimal policies that converge to a limiting policy. Since there exists a policy for any $g > 0$, the limiting policy exists and is optimal.

The combination in this paper between optimal stopping theory with the bisection method is a new and important extension to the work of Wijngaard and Stidham Jr. [28, 29] and an essential ingredient in the optimality proof of (s, S) -policies. Contrary to our work, Wijngaard and Stidham Jr. use the bisection method principally as a *numerical* means to efficiently compute the long run average cost of skip-free Markov decision processes on a finite or countable state space. We extend it such that it can be used to *prove* structure results for optimal policies. In passing, we extend their method such that it also applies to Markov processes on continuous, rather than countable, state spaces, and we slacken the condition that the transition probabilities $p_{i,i+1} > 0$ for all states i .

The marriage between the numerical and theoretical part of the bisection method allows us to achieve numerous results at the same time. The theoretical part makes the method more generally applicable to deriving structure results in other settings. For instance, our results can also be used to prove the optimality of, so-called, D -policies, c.f. Section 2, in the space of stationary policies for the workload process of the $M/G/1$ queue as studied by Feinberg and Kella [4]. We can adapt it to handle inventory systems with lost sales under various non-trivial rejection policies, such as complete rejection (reject the demand when it cannot be covered from on-hand inventory), complete acceptance (reject only demand that arrives when the inventory level is negative), or partial acceptance (accept that part of the demand that can be covered by on-hand inventory, and reject the remaining part of the demand). The method can also be applied to periodic-review systems. For instance, in Section 5 we show how to apply it to a periodic-review inventory system described by Graves [11], c.f., Section 2.3, and prove that an (s, S) -policy is optimal.

With regard to the numerical part of the method, it enables us to construct optimal policies for a wide range of problems even when the structure is not simple, e.g., not monotone, and certainly not (s, S) , as is the case in one of the examples of Section 7. This is an important contribution by itself, in that none of the other methods discussed up to now or below in Section 2, offer this possibility. Next, the method is entirely straightforward, and certainly does not require any of the ingenious insights that have been developed to obtain numerical results. Thus, it is not necessary to find expressions for the average costs and times of regenerative cycle times under some *posited* policy structure. This is also a major distinctive and simplifying feature of our approach. Finally, the numerical aspect of the bisection method brings the off-spin that the optimal policy can be computed very efficiently, typically exponentially fast.

The structure of the paper is as follows. In Section 2 we discuss related work. Section 3 introduces the model. Then, in Section 4, we prove the existence of an optimal stationary policy under quite general conditions, and provide insight into when this (s, S) policies can be optimal. In Section 5 we show how to apply our method to a capacitated periodic-review inventory model of Graves [11]. Section 6 discusses

some numerical issues. These are used to in the examples of Section 7 to illustrate, on the one hand, the procedure by means of a concrete example, and, on the other hand, to provide counter examples against the conjecture that (s, S) -policies are optimal for all combinations of demand distribution and quasi-convex inventory costs. Section 8 concludes and provides an outlook on potential further applications of our method.

2 Discussion and Related Literature

In this section, we discuss related literature and motivate why the methods developed earlier are not suitable to obtain the results of this paper. We group in Sections 2.1–2.4 the literature in accordance to similarity in modeling approach, and then, in Sections 2.5–2.7, in accordance to method of proof.

2.1 Relation to Controlled Queueing Systems

As mentioned above, Sobel [22] proves that two-critical-number policies are long-run average optimal for two-sided skip-free stochastic processes defined on the space of integers. Examples are birth-death processes, the $G/G/1$ queue length process, and inventory processes that change by single units. Gavish and Graves [8] map the inventory level $I(t)$ (in items) at time t to the queue-length $Q(t) = S - I(t)$ of the $M/D/1$ queue, where S is the inventory level at which production switches off. Then, using the result of Sobel [22], it suffices for the identification of the optimal policy to find an algorithm that computes the optimal critical numbers. Others, see below, use similar approaches.

Given this versatile equivalence between controlled queue-length processes and inventory level processes, it is of interest to try to extend it to more general cases. In particular, the sample paths of the (s, S) -controlled inventory level for our production-inventory system can be made to correspond in a one-to-one way to the sample paths of the (virtual) workload process of an $M/G/1$ queue with a removable server that is controlled by a, so called, D -policy. Such D -policies switch on the server when the *workload*, rather than the *queue length*, exceeds a level D , and switch off the server until the workload becomes zero. To establish the mapping, note that the capacity-constrained production rate of the production-inventory system is analogous to the finite service rate of the server, and the (virtual) workload process $V(t) = S - I(t)$ at time t when the continuous inventory level is $I(t)$ at time t . Thus, S —the level at which production is switched off in the inventory system—is mapped to 0 in the queueing system, and s —the level at which production is switched on—to $D = S - s$.

In the queueing context, Federgruen and So [1] prove for the $M^X/G/1$ queue the optimality of N -policies, policies that switch on a removable server when the queue-length exceeds some integer N , and switch off when the system becomes empty. Federgruen and So [1] conjecture that D -policies are optimal to control the workload when service becomes known upon arrival. Feinberg and Kella [4] prove in full generality the long-run average optimality of D -policies in the space of all policies, not just the stationary policies, for the $M/G/1$ queue with a controllable server, fixed switching costs of the server, and non-decreasing holding costs.

Given these optimality results for these queueing systems, it would be nice to try to carry over these results to the production-inventory system under consideration. This, for instance, would prove that (s, S) -policies would be optimal for the case of Graves and Keilson [10]. However, this is not possible in general. Even though the sample paths of workload and inventory level processes are in one-to-one correspondence, the cost structures are certainly not the same. For the $M/G/1$ queue, the holding costs is non-decreasing in the workload (or the queue-length), while the inventory cost function $h(\cdot)$ for the inventory system is generally not monotone, but (quasi)-convex on $(-\infty, S]$, and such that $h(S) > h(0) = 0 < h(z)$ for some sufficiently large z .

Precisely this difference in the cost functions motivates Federgruen and Zheng [2] to generalize the queueing model of Federgruen and So [1] to production-inventory systems with unit production and vacations, and point out that the queueing system is a special case of the inventory system under an $(s, 0)$ -policy, i.e., an inventory system that is purely make-to-order, and no orders are produced to inventory. Similarly, our model generalizes that of Feinberg and Kella [4] to more general cost structures. Our result is not a

full generalization, though, as we consider optimality in the class of stationary policies, while Feinberg and Kella [4] prove that a D -policy is optimal in the class of all policies.

As a matter of fact, the difference in the cost functions has severe consequences for the structure of the optimal policies for the production-inventory system. In Section 7 we provide a counter example that shows that (s, S) -policies are not optimal for certain combinations of quasi-convex inventory costs and exponentially distributed demand. Thus, the optimality of the D -policy for the D -controlled $M/G/1$ queueing system does not imply the optimality of (s, S) -policies for the production-inventory system, and as an immediate consequence, D -policies are not optimal under holding costs that are quasi-convex, rather than non-decreasing, in the workload. Moreover, in full generality, (s, S) -policies are not optimal for the production-inventory system under consideration here; hence, one of the objectives of this paper is to find criteria on the demand distribution and inventory cost function such that (s, S) -policies are optimal.

This brings us to the question why the result of Sobel [22] apparently do not generalize to, for instance, the D -controlled workload process of the $M/G/1$ -queue. The reason is that Sobel is concerned with two-sided skip-free processes, i.e., simple random walks on the integers that only make transitions to neighboring states. When the server is off, the service-queue length process $(S(t), Q(t))$ ‘lives’ on the off-line $1 \times \{\dots, -1, 0, 1, \dots\}$. Suppose that $Q(t) < M$, where M is the larger of the two-number critical policy, then, eventually, the queueing process will hit M , the server will switch on, and the server/queue length process jumps to the on-line $1 \times \{-1, 0, 1, \dots\}$. As a consequence, the process cannot enter the set $0 \times \{M + 1, M + 2, \dots\}$. Reasoning similarly, $(S(t), Q(t))$ will never enter the set such that the server is off and $Q(t) < m$, where m is the smaller of the two critical numbers. Thus, no matter how pathological the structure of any policy is at the sets $0 \times \{M + 1, M + 2, \dots\}$ and $1 \times \{\dots, m - 2, m - 1\}$ (where m and M are allowed to be policy-dependent), the characteristics of the policy on these sets cannot affect the behavior of the process.

In contrast, the inventory level under compound Poisson demand with generally distributed demand sizes (or the workload process of the $M/G/1$ -queue) is no longer two-sided skip-free. In our case, while indeed the production-inventory will never exceed S (save an initial transient phase that depends on the starting condition of the inventory level process), the inventory level can make transitions of arbitrary size to the left. Now, suppose that it is optimal on the off-line to switch on in some finite interval $(s - a, s]$ for some finite $a > 0$, remain off in $(s - b, s - a]$, $b > a$, and on again in $(-\infty, s - b]$. When the inventory process would have been skip-free to the left, it is impossible to hit the ‘wrong’ sets $(s - b, s - a]$. However, when the demand distribution has with infinite support, for instance, this is not the case. In fact, the example in Section 7 mentioned before, shows that such alternating policies are optimal for some production-inventory systems, hence, (s, S) -policies are not optimal in general.

2.2 Single-Item Systems with General Processing Times

Gavish and Graves [9] and Lee and Srinivasan [16] derive algorithms to compute the optimal policy in the class of (s, S) policies for the production-inventory system with stochastic production times rather than constant production times. Thus, they are concerned with the production-inventory analogy of the $M/G/1$ -queue, rather than the $M/D/1$ -queue as Gavish and Graves [8]. Srinivasan and Lee [23] consider the more difficult case of a batch Poisson arrival process in which the demands have general processing times and linear holding costs. Their review model is different in that the inter-arrival time between successive review epochs during off-periods is a random variable. Srinivasan and Lee [23] confine the analysis to (s, S) -policies and derive an algorithm to find the optimal s and S .

Production-inventory systems with generally distributed production times are similar to our system in that in both cases it takes a duration D to replenish the inventory when a demand of size D arrives, but the difference is that in our case (and that of Feinberg and Kella [4]) the inventory level is physically reduced by an amount of D while in the other cases it is reduced by just a single item.

2.3 Capacitated Periodic-Review Systems

Attempts have also been made to incorporate bounds on the order quantity for periodic-review stochastic inventory systems. This resembles the finite production capacity of the continuous-review production-inventory process. Federgruen and Zipkin [3] prove that the optimal policy is an order-up-to policy, i.e., an

(s, S) policy with $s = S$ when the order quantity is bounded by some number A and there are no ordering cost. In contrast, when the ordering cost K is positive, one would expect the optimal policy to have the following structure: do not order if the inventory $I(t) > s$, and order $\min\{A, S - I(t)\}$ if $I(t) < s$. However, this is not true, see Wijngaard [26] for a straightforward counterexample. In fact, the general optimal policy does not have a simple structure as shown by Shaoxiang and Lambrecht [20], Gallego and Scheller-Wolf [6], and Shaoxiang [19], but is characterized by critical levels X and Y , $X \leq Y$, on the inventory: when $I(t) < X$ order A , when $I(t) > Y$ do not order, but when $X < I(t) < Y$ the decision structure is more difficult. A related approach is to consider only full-capacity orders, see Wijngaard [27]. More recently Gallego and Toktay [7] also consider this problem and show that the optimal policy *under discounting* is a threshold policy. Under the long-run average expected cost optimality criterion, however, this inventory problem appears to be somewhat less interesting. Since the ordering cost per time unit simply becomes $K \lambda E[D]/A$, the behavior of the optimal policy does not change if K increases. This aspect of this model is also somewhat undesirable in comparison to the ‘normal’ periodic-review model with unlimited order size. There, if K increases, the difference $S - s$ typically also increases, thereby making the trade-off between ordering cost and inventory cost explicit. If, however, the ordering quantity is fixed to A , as in Gallego and Toktay [7], this trade-off is no longer of relevance.

An interesting related model is studied by Graves [11]. We discuss this in considerable detail in Section 5.

2.4 Continuous-review Systems with Stochastic Demand

In the context of continuous-review inventory models with unlimited production capacity (or order-size independent delivery times), Hordijk and Van der Duyn Schouten [14] show that (s, S) -policies are optimal, under the discounted and average cost criterion, for an Economic Order Quantity (EOQ) model with a demand process that is the sum of a deterministic (state-dependent) demand rate and a compound Poisson process. Presman and Sethi [18] obtain the same results for the slightly less general system with constant, rather than state-dependent, deterministic demand. The main difference with our model is that these authors consider unlimited production capacity while we assume finite capacity, and they consider a more general stochastic demand process as it contains, besides the compound Poisson demand, a constant demand term. Presman and Sethi [18] show that, in the average cost case, the average cost per cycle of the best (s, S) -policy cannot be improved by any other policy. However, as their approach is not constructive, it might be difficult to apply to more general inventory models in which it may be inferred that the optimal policy is more complex than (s, S) .

Moreover, the method does not require ingenious expressions for the average cost and cycle times of regenerative cycles associated with a specific (s, S) -policy.

2.5 Monotone Policies

In numerous settings, monotonicity properties of candidate optimal policies have been used to prove structure results. Such arguments, however, cannot be applied to our case. The numerical examples in Section 6 show that, simply put, optimal policies are not monotone in general. In one of the examples, the off-line indeed splits into *multiple* disjoint intervals under the optimal policy. On some of these intervals it is optimal to stay off, while on the complement of these intervals it is optimal to switch on. To still establish the optimality of (s, S) -policies in some of these cases, we can prove that it is optimal below s to switch on, and to stay off anywhere in an interval $(s, S]$, and that it is optimal to keep the inventory below or at S . Since the ‘wrong’ intervals of the off-line lie above S , these will never be hit again once the inventory enters the set $(-\infty, S]$.

2.6 Sample Path Arguments

Perhaps, more generally, sample-path arguments may be exploited to obtain our results, but it is not so clear how. In periodic-review systems, Huh and Janakiraman [15] successfully apply sample-path arguments to prove the optimality of (s, S) -policies for the classical period-review inventory system with stochastic demand and fixed setup costs, but under the discounted cost criterion, not under the average-cost criterion.

Fu [5] uses sample-path derivatives to find optimal s and S in the class of (s, S) -policies, but does not prove the optimality of (s, S) -policies in the space of stationary policies. For the $M/G/1$ queue, Feinberg and Kella [4] use sample-path arguments (and coupling) to establish that, by judiciously switching the server on and off, (parts of) the sample paths can be shifted downward. Since they assume non-decreasing holding costs, a downward shift in the workload also decreased the costs. As a consequence, ‘serve-to-exhaustion’, i.e., serve until the workload is zero, is long-run average optimal. However, this type of argument does not simply carry over to inventory problems with the more general cost functions, as discussed above. In such cases, then, shifting (part of) a sample path downward need not result in overall lower costs.

2.7 Markov Decision Processes

Finally, we remark that the continuous-review production-inventory system can also not be easily formulated as a discrete-time Markov decision problem (MDP). (As discrete-time MDP require less technical detail, it is of interest to try to obtain our results via this approach.) The point is that under any (s, S) -policy the decision epoch to switch off is not a jump epoch of the demand process. (Sample paths are by assumption right-continuous. As S has to be hit from the left, and jumps are also to the left, S can never be hit at jump epochs.) Thus, the controlled production-inventory process cannot be formulated as an MDP with a chain embedded at jump epochs. It seems possible to ‘repair’ this problem by including fast but small jumps to the right into the Markov chain, and study the sequence of MDPs that results when taking the jump rate larger and the jump sizes smaller. We infer that such a sequence of Markov decision problems converges to the production-inventory system, but proving this may not be easy, neither that the structure of the optimal policies is maintained in the limit, see [21] for details on such limiting procedures. Hence, to prove our results as a limit of discrete-time MDPs appears not to be straightforward.

3 Model, Notation and Preliminaries.

There is one stock keeping unit. Orders arrive according to a Poisson process with parameter λ . The order size D is distributed according to $F(\cdot)$ with mean $\mathbb{E}[D] > 0$ such that $\lambda \mathbb{E}[D] < 1$. It is assumed that D is light-tailed, i.e., there exists an $\alpha > 0$ such that $\int_0^\infty e^{\alpha x} G(x) dx < \infty$, where $G(x) = 1 - F(x)$ is the survival function, and that $\mathbb{P}(D = 0) = 0$. Demand is filled from stock, or backlogged if no stock is available. Inventory and stock-out costs are represented by the function $h(\cdot)$, which is assumed to be quasi-convex, non-negative with $h(0) = 0$, and $h(x) = O(|x|^n)$ for some $n > 0$ as $|x| \rightarrow \infty$. The supply comes from production. When production is on, the supply is produced continuously at a fixed rate, which is set to 1 without loss of generality. When production is off, the production rate is 0. Thus, the only way to control production is by switching it on and off. Each time the production is switched on, a fixed cost K is incurred. (In case there is a switch-off cost, this cost can be trivially absorbed in K .) Switching occurs instantaneously, so that if production is switched on, the inventory immediately starts to increase. Because of backlogging, it is possible to relax the assumption to have a fixed throughput time, i.e., the time lag between the start of production and the inventory starting to grow, but we do not do so here.

The state space of the controlled production-inventory process $\{P(t), I(t)\}$ can be visualized as two lines, the on-line with $(P(t), I(t)) = (1, x)$ and the off-line with $(P(t), I(t)) = (0, x)$. Any stationary policy π can be fully characterized by a subset O_1^π of the on-line and a subset O_0^π of the off-line. If $P(t) = 1$ and $I(t)$ hits O_1^π production is switched off, while if $P(t) = 0$ and $I(t)$ hits O_0^π production is switched on. In the sequel we use this correspondence between policy and these subsets interchangeably. Observe that as production is continuous, any subset of the on-line is entered from the left. Therefore it is necessary that O_1^π for any stationary policy π is left closed, in other words, $\liminf x_i \in O_1^\pi$ for any sequence $\{x_i\}$ in O_1^π . Next, to be able to properly define expected hitting times we require the sets O_0^π and O_1^π to be Borel sets. We impose some further conditions on O_0^π and O_1^π to ensure that only recurrent cycles with finite expected costs can result. First, the skip-freeness-to-the-right of the inventory process on the on-line and the stability condition $\lambda \mathbb{E}[D] < 1$ imply that the set O_1^π has finite expected hitting time if it can be reached from any initial condition $I(0) = x$. Second, on the off-line the inventory process $\{I(t)\}$ behaves according to a compound Poisson process. We require that O_0^π is such that it is hit with finite expected time for any $I(0) = x$. (Providing tight necessary conditions for this to be true is actually quite technical and

dependent on the specific properties of $F(\cdot)$. A simple, but certainly more than necessary, condition is that there exists some z^π for each policy π such that $(-\infty, z^\pi] \subset O_0^\pi$.) Third, O_0^π and O_1^π should be disjoint, for otherwise it may happen that $I(t) \in O_0^\pi \cap O_1^\pi$ and production switches on and off incessantly, resulting in an infinite cost.

In the proof we consider first a space of ‘tamed’ stationary policies consisting of policies such that $O_1^\pi = [R, \infty)$ for a given, fixed, R , while O_0^π can still be arbitrary. Typical recurrent production cycles induced by such policies may be assumed to start at $(P(0), I(0)) = (0, R)$, see Figure 1. The inventory decreases in jumps at each arriving order until O_0^π is hit. Then production is switched on. The inventory increases now continuously, except when an order arrives, until it hits R . Then production is switched off again, and the process returns to the point $(0, R)$. We write \mathcal{H}_R to denote the class of policies π such that $O_1^\pi = [R, \infty)$, and $\mathcal{H} = \bigcup_R \mathcal{H}_R$. Thus, the policies in \mathcal{H} are such that they control pure jump processes at the off-line and jump processes with drift at the on-line.




Figure 1: The state space of the inventory process. The thick bar on the on-line represents the subset in which the policy π_R decides to switch production off. The off-line is dashed to indicate that the inventory does not decrease continuously, but in jumps.

4 Analysis.

In Section 4.1 we present an outline of the proof to clarify how all steps of the proof relate and provide a list of steps that require further proof. These technical points are addressed in the subsequent Sections 4.2–4.5.

4.1 Notation and Proof Outline.

We start with introducing some concepts related to the notion of *g-revised cost*. Take some arbitrary R and consider some arbitrary policy $\pi \in \mathcal{H}_R$. Let

$$\tau^\pi(x) = \inf\{t > 0; (P(t), I(t)) = (1, x)\},$$

where E_x is the expectation of functionals of the π -controlled process $\{P(t), I(t)\}$ given that $(P(0), I(0)) = (0, x)$, $x \leq R$. By the assumptions of Section 3, the sets O_1^π and O_0^π are such that $E_x \tau^\pi(x) < \infty$. Let $C^\pi(x)$ denote the expected cost (inventory, stock-out and set-up cost) for the process to move from state $(0, x)$ to $(1, x)$, i.e.,

$$C^\pi(x) = E_x \left[\int_0^{\tau^\pi(x)} h(I(t)) dt \right] + K. \quad (1)$$

We define the (expected) *g-revised cost* to move from $(0, x)$ to $(1, x)$ as

$$C_g^\pi(x) = C^\pi(x) - g E_x \tau^\pi(x), \quad (2)$$

where $g > 0$ is some arbitrary *revision rate*. Clearly, $C_g^\pi(x) = E_x \left[\int_0^{\tau^\pi(x)} (h(I^\pi(t)) - g) dt \right] + K$, and we may interpret the *g-revised cost* as the expected cost resulting from an inventory cost $h(x) - g$ rather than $h(x)$.

We next introduce the g -minimizing policy in \mathcal{H} . Suppose that

$$C_g(x) = \inf_{\pi \in \cup_{R \geq x} \mathcal{H}_R} \{C_g^\pi(x)\} \quad (3)$$

exists. Then $C_g(x)$ is the *least* g -revised cost to move from $(0, x)$ to $(1, x)$. Assume further that for any given g , $C_g(\cdot)$ attains it (finite) minimum, and let S_g be the left-most minimizer. Then $C_g(S_g)$ is the *least overall* g -revised cost to move from some level on the off-line to the same level on the on-line. Suppose finally that there exists a policy π_g that establishes the infimum in (3) at S_g , so that it follows that

$$C_g^{\pi_g}(S_g) = C_g(S_g) = \min_x C_g(x) = \min_x \min_{\pi \in \cup_{R \geq x} \mathcal{H}_R} C_g^\pi(x). \quad (4)$$

We call π_g the g -minimizing policy.

It is interesting to note that $C_g^\pi(S)$ is not just the expected cost of to move from $(0, S)$ and to stop at $(1, S)$ under policy π . As the cost to switch off is absorbed in K , $C_g^\pi(S)$ can also be seen as the long-run expected cost of a *recurrent cycle* that starts and stops in $(0, S)$. This insight allows us to relate the revision cost g to the expected recurrence cost of a cycle under the policy π . It is well known, see for instance Tijms [24], that the average recurrence cost g^π equals $C^\pi(R)/E_R \tau^\pi(R)$, where the cost function C is defined in (1), for a cycle induced by policy $\pi \in \mathcal{H}_R$. Rewriting this in the form of (2), it means that g^π and π are related such that

$$C_{g^\pi}^\pi(R) = C^\pi(R) - g^\pi E_R \tau^\pi(R) = 0. \quad (5)$$

To find the recurrence cost by this approach, first *choose* some policy π , and then *compute* the cost rate g^π such that (5) is satisfied. Now observe that in the reasoning leading to (4) we actually reverse this procedure. There we first *choose* some initial g , and then *construct* the policy $\pi_g \in \mathcal{H}$ that achieves the minimal cost $C_g^{\pi_g}(S_g)$. If we can find a $g^* > 0$ and a policy π_{g^*} such that

$$C_{g^*}^{\pi_{g^*}}(S_{g^*}) = 0, \quad (6)$$

then, by (4), $C_{g^*}^{\pi_{g^*}}(x) \geq 0$ for any other $\pi \in \mathcal{H}$ and any other x . Next, by taking $g^\pi = g^*$, $\pi = \pi_{g^*}$ and $R = S_{g^*}$ in (5), it follows from (6) that g^* is the long-run cost rate associated with π_{g^*} . All in all these observations imply, in words, that π_{g^*} is the optimal policy in \mathcal{H} with minimal long-run average expected cost rate g^* , and S_{g^*} is the optimal switching level (from on to off).

To actually find C_g in (3) and the g -minimizing policy in \mathcal{H} we derive a dynamic programming equation (DPE) that C_g has to satisfy. It turns out that this DPE has a nice form and is easy to solve numerically for any g . To identify the ‘right’ g , i.e., the revision cost that solves (6), we use bisection. Choose some arbitrary g . It may then happen that this choice for g is not equal to the recurrence cost \bar{g} of the cycle induced by π_g . If $g > \bar{g}$, which intuitively would mean that the cost compensation rate g is larger than the actual average running cost \bar{g} , then

$$\begin{aligned} C_g^{\pi_g}(S_g) &= C_g^{\pi_g}(S_g) - g E_{S_g} \tau^{\pi_g}(S_g) \\ &< C_g^{\pi_g}(S_g) - \bar{g} E_{S_g} \tau^{\pi_g}(S_g) = 0. \end{aligned}$$

As a second trial we choose a smaller value, g' say, as a revision rate, and construct the associated policy $\pi_{g'} \in \mathcal{H}$. When $C_{g'}^{\pi_{g'}}(S_{g'}) > 0$, the next choice, g'' say, must be larger than g' , while if $C_{g'}^{\pi_{g'}}(S_{g'}) < 0$, g' is still too large, and so on.

Now that we have described the overall procedure of the proof we enumerate the steps that need to be filled in to obtain the structure results. In Section 4.2 we first derive a useful expression for the derivative $\gamma(\cdot)$ of the expected cost to move on the on-line from some level r to a level $x > r$ for quite general inventory cost functions $c(\cdot)$. We next show, in Section 4.3, that the g -minimal cost $C_g(x)$ of (3) is the unique solution of a DPE related to an optimal stopping problem with cost-to-go function $\gamma(\cdot)$, and that a policy exists that achieves the minimum. This optimal policy prescribes to switch on production if the inventory is less than s_g , which is the left-most root of γ . We then show that $C_g(x)$ attains its minimum for all $g > 0$, and that the left-most minimizer S_g of C_g lies in the interval $(s_g, t_g]$, where t_g is the right-most root of γ . In Section 4.4 we show that $C_g(S_g)$ can be made arbitrarily small, hence there exists a g such that $C_g(S_g) < 0$. Since $C_g(S_g)$ is continuous in g and $C_0(S_0) > 0$ we conclude that there must exist a g^*

such that $C_{g^*}(S_{g^*}) = 0$. As there exists a g -minimizing policy for each g , this implies there must exist a g^* -minimizing policy $\pi_{g^*} \in \mathcal{H}$. Thus, the bi-section on g indeed yields most of the results. The last step of the procedure, c.f. Section 4.5, is to prove that no stationary policy can improve π_{g^*} , hence π_{g^*} is an optimal stationary policy. A simple corollary of the above is that when $h(\cdot)$ is convex, γ is convex, hence has two roots. It is then immediate that the optimal policy has an (s, S) structure.

In the rest of this section, we deal with the technical details to prove the structure results. In Section 6 and Appendix A we discuss some numerical points of concern, and in Section 7 we illustrate the bisection procedure for a concrete example, and provide some counter-examples.

4.2 Costs on the On-line.

Consider a policy $\pi_R \in \mathcal{H}_R$. If the process starts in some state $(P(0), I(0)) = (1, r)$ with $r < R$, π_R prescribes to continue producing until state $(1, R)$ is reached, and then to switch off. In this section we are concerned with the cost function $V(r, x)$ under π_R of ‘moving’ on the on-line from $(1, r)$ to $(1, x)$ for $r < x \leq R$.

Consider a given cost rate function $c(\cdot) \in \mathcal{B}_{\alpha, R}$, where $\mathcal{B}_{\alpha, R}$, $\alpha > 0$, is the Banach space of real-valued continuous functions f such that the weighted supremum norm $\|f\| = \sup\{|f(x)e^{\alpha x}|; x \leq R\}$ is finite. Let $V^n(r, x)$ be the cost incurred to start in $(1, r)$ and continue producing until either point $(1, x)$ is reached or the n -th order arrives, whichever is first; let $V^0(r, x) = \int_r^x c(y)dy$. When $n > 0$, we have for sufficiently small Δ ,

$$V^n(r, x + \Delta) = V^n(r, x) + c(x)\Delta + \lambda\Delta \mathbb{E}[V^{n-1}(x - D, x)] + o(\Delta).$$

By induction it follows that $V^n(r, \cdot)$ is differentiable, and its derivative $\gamma^n(\cdot)$ satisfies

$$\gamma^n(x) = c(x) + \lambda \mathbb{E} \left[\int_{x-D}^x \gamma^{n-1}(z) dz \right]. \quad (7)$$

Observe that $\gamma^n(x)$ is independent of r .

Lemma 4.1. *As the demand is light-tailed, there exists $\alpha > 0$ such that $\beta := \lambda \int_0^\infty e^{\alpha x} G(x) dx < 1$. If $c(\cdot) \in \mathcal{B}_{\alpha, R}$ and $c(x) \rightarrow \infty$ as $|x| \rightarrow \infty$, the sequence $\{\gamma^n\}$ converges to a limit function $\gamma \in \mathcal{B}_{\alpha, R}$, which is the unique solution of the integral equation*

$$\begin{aligned} \gamma(x) &= c(x) + \lambda \mathbb{E}[V(x - D, x)] \\ &= c(x) + \lambda \mathbb{E} \left[\int_{x-D}^x \gamma(z) dz \right] \\ &= c(x) + \lambda \int_0^\infty \gamma(x - z) G(z) dz, \end{aligned} \quad (8)$$

where $V(r, x) = \int_r^x \gamma(z) dz$. The function γ is decreasing (increasing) in a neighborhood of $-\infty$ ($+\infty$).

Proof. Define the linear operator

$$(Pf)(x) = \lambda \int_0^\infty \int_{x-y}^x f(z) dz dF(y),$$

and rewrite (7) as $\gamma^n(x) = c(x) + (P\gamma^{n-1})(x)$, with $\gamma^0 = c$. The operator P is a contraction on the Banach space $\mathcal{B}_{\alpha, R}$, that is, $\|P\| = \sup\{\|Pf\| : \|f\| = 1\} \leq \beta$. To see this, note that

$$(Pf)(x) = \lambda \int_0^\infty \int_{x-y}^x f(z) dz dF(y) = \lambda \int_0^\infty f(x - z) G(z) dz,$$

which implies

$$(Pf)(x)e^{\alpha x} = \lambda \int_0^\infty f(x - z)e^{\alpha(x-z)} e^{\alpha z} G(z) dz \leq \lambda \|f\| \int_0^\infty e^{\alpha z} G(z) dz = \beta \|f\|,$$

hence $\|Pf\| \leq \beta \|f\|$. From this it follows that $\gamma^n \rightarrow \gamma \in \mathcal{B}_{\alpha, R}$. This limit function γ is the fixed point of the operator $c + P$, that is, γ is the unique solution of the integral equation (8). As a further consequence, $V^n(r, x)$ converges to $V(r, x) = \int_r^x \gamma(z) dz$.

Since $c(x) \rightarrow \infty$ if $x \rightarrow \pm\infty$ by assumption, γ must also be decreasing (increasing) in a neighborhood of $-\infty$ ($+\infty$). \square

4.3 Minimal Expected Costs from $(0, x)$ to $(1, x)$.

We next characterize the minimal cost $C(x)$ to move from state $(0, x)$ to $(1, x)$.

Suppose that $(P(0), I(0)) = (0, x)$, i.e., the process starts at inventory level x and production is off. Choose $R > x$. Let us compare the expected costs of two alternative routes to move from $(0, x)$ to $(1, x)$. The first alternative is to switch on immediately, which costs clearly K . The second alternative is to wait for the next arrival, switch on, and then replenish the inventory to level x again. The cost for this second route is

$$c(x)\lambda^{-1} + K + EV(x - D, x) = \gamma(x)/\lambda + K,$$

where the equality follows from (8) and the definition of $V(r, x)$. Now observe that if x is such that $\gamma(x) < 0$, it must be better to await a new arrival than to switch on immediately. We therefore write

$$C^1(x) = \min \{K, \gamma(x)/\lambda + K\}, \quad (9)$$

for the cost to move from $(0, x)$ to $(1, x)$ associated with the policy that either switches on immediately, or awaits at most one arrival before switching on, whichever is best. This reasoning applies of course recursively. Rather than awaiting just one arrival, it may be interesting, cost-wise, to await two arrivals. This would result in a cost $\gamma(x)/\lambda + E[C^1(x - D)]$ for the second term in the minimization in (9) rather than $\gamma(x)/\lambda + K$. Continuing like this, we arrive at a sequence of cost functions $\{C^i(\cdot)\}$ which satisfy the recursions

$$C^i(x) = \min \{K, \gamma(x)/\lambda + E[C^{i-1}(x - D)]\}, \quad i \geq 1, \quad (10)$$

where we define $C^0(x) = K$ for all x .

In the sequel of this section, assume that the cost rate $c(\cdot)$ is such that $\gamma(x) < 0$ for at least one x , for otherwise production will be switched on everywhere on $(-\infty, R]$, which is relatively uninteresting. Since γ is continuous and increases to ∞ as $|x| \rightarrow \infty$, it then has at least two roots. Let s and t be its left- and right-most roots, respectively.

Lemma 4.2. *The sequence $\{C^i(\cdot)\}$ has the following properties:*

1. $C^i(x) = K$ for all $x \leq s$.
2. $\{C^i(\cdot)\}$ is point-wise decreasing as $i \rightarrow \infty$, i.e., $C^i(x) \leq C^{i-1}(x)$.
3. $\{C^i(\cdot)\}$ is bounded from below, i.e., there exists $M > -\infty$ such that $C^i(x) > M$.
4. $C^i(\cdot)$ is continuous for all i on $(-\infty, R]$.

Proof. (i) This follows immediately from combining $\gamma(x) > 0$ for all $x < s$ with the recursion (10), and (9) which gives that $C^1(x) = K$ whenever $x \leq s$.

(ii) If $x \leq s$, $C^i(x) = K$, for all i . For $x > s$ we use induction. Suppose $C^j(x) \leq C^{j-1}(x)$ for all $j < i$. Use (10) to see that $C^i(x) - C^{i-1}(x) \leq EC^{i-1}(x - D) - EC^{i-2}(x - D)$. Then, by induction

$$\begin{aligned} C^i(x) - C^{i-1}(x) &\leq EC^{i-1}(x - D) - EC^{i-2}(x - D) \\ &\leq EC^1(x - D_{i-1}) - K \\ &\leq 0, \end{aligned} \quad (11)$$

by (9), where we write D_{i-1} to denote the cumulative demand of $i - 1$ orders.

(iii) Define $\|v\| = \sup\{|v(x)|; x \in [s, R]\}$. We prove that $\|C^i - C^{i-1}\| \rightarrow 0$ geometrically fast. Using (11), $\|C^i - C^{i-1}\| \leq \sup_{x \in [s, R]} |\mathbb{E}C^1(x - D_{i-1}) - K|$. By (9), $\mathbb{E}C^1(x - D_{i-1}) - K = \int \min\{0, \gamma(x - y)\} dF^{i-1}(y)$, where $F^i(\cdot)$ is the i -fold convolution of $F(\cdot)$. Since $\gamma(x) > 0$ for $x < s$, this integral becomes $\int_0^{x-s} \min\{0, \gamma(x - y)\} dF^{i-1}(y)$. Therefore,

$$\|C^i - C^{i-1}\| \leq \sup_{x \in [s, R]} \int_0^{x-s} |\gamma(x - y)| dF^{i-1}(y) \leq \|\gamma\| \int_0^{R-s} dF^{i-1}(y).$$

Now take some arbitrary $\beta > 0$, and write $\tilde{F}(\beta)$ for the Laplace transform of $F(\cdot)$. Then,

$$\begin{aligned} \|C^i - C^{i-1}\| &\leq \|\gamma\| e^{\beta(R-s)} \int_0^{R-s} e^{-\beta y} dF^{i-1}(y) \\ &\leq \|\gamma\| e^{\beta(R-s)} \int_0^\infty e^{-\beta y} dF^{i-1}(y) \leq \|\gamma\| e^{\beta(R-s)} (\tilde{F}(\beta))^{i-1}. \end{aligned}$$

Since $\tilde{F}(\beta) < 1$, $\|C^i - C^{i-1}\| \rightarrow 0$ geometrically fast. Therefore $C(x) > -\infty$.

(iv) C^i is continuous as follows from the recursion (10) and the fact that γ is continuous as $\gamma \in \mathcal{B}_{\alpha, R}$. \square

Now Lemma 4.2.ii,iii imply that $C(\cdot) = \lim_i C^i(\cdot)$ exists. If we can show that in (10) the limit in i and the expectation can be reversed, the limit function $C(\cdot)$ is a solution of an *optimal stopping problem* with $\gamma(\cdot)/\lambda$ as the *cost-to-go*, K as the *stopping cost*, and $(-\infty, s]$ as stopping set. Moreover, if $C(\cdot)$ is the unique solution it must represent the minimal cost to move from $(0, x)$ to $(1, x)$. The next theorem provides an answer in the affirmative.

Theorem 4.3. *The limiting function $C(x) = \lim_{i \rightarrow \infty} C^i(x)$ is the unique solution of the dynamic programming equation*

$$C(x) = \begin{cases} K, & \text{if } x \leq s, \\ \min\{K, \gamma(x)/\lambda + \mathbb{E}C(x - D)\}, & \text{if } x > s. \end{cases} \quad (12)$$

Proof. First fix some $R > \min\{x, s\}$. Lemma 4.2.(ii, iii) imply the existence and uniqueness of a limiting function $C(\cdot)$. Lemma 4.2.i shows that $C(x) = K$ for $x \leq s$. To see that $C(\cdot)$ satisfies (12) we use monotone convergence to justify the interchange of limit and expectation. The only formal condition to verify here is that $C^i(x) < \infty$ for all i , which is satisfied as from (10) it follows that $C^i(x) \leq K$ for all i and all $x \leq R$. Thus, the existence of a solution of (12) is established.

To prove that $C(\cdot)$ is the unique solution we use that the operator T , defined as

$$(Tv)(x) = \min\{K, \gamma(x)/\lambda + \mathbb{E}v(x - D)\},$$

is a contraction on the Banach space of bounded functions v on $[s, R]$ with norm $\|v\| = \sup\{|v(x)|; x \in [s, R]\}$. By the above we already have that $(TC)(x) = C(x)$ for $x \in [s, R]$. To see that T is a contraction, observe that since $|\min\{K, a\} - \min\{K, b\}| \leq |a - b|$ for any a, b ,

$$|(Tv)(x) - (Tw)(x)| \leq |\mathbb{E}[v(x - D) - w(x - D)]| \leq \|v - w\| e^{\beta(R-s)} \tilde{F}(\beta),$$

by similar arguments as used in the proof of Lemma 4.2.iii. The contraction follows since $\tilde{F}(\beta) < 1$.

The form of (12) shows that $C(x)$ is in fact independent of R for all $x \leq R$. Hence, we can take R as large as necessary, and in particular larger than the global minimizer of $C(\cdot)$. (Recall that in (4) we need to find the global minimum of $C(\cdot)$.) \square

In the sequel we need two further properties of $C(\cdot)$.

Lemma 4.4. *$C(\cdot)$ is continuous and attains its minimum. Its left-most minimizer $S \in (s, t]$.*

Proof. With regard to the continuity, use the triangle inequality: $|C(x) - C(y)| \leq |C(x) - C^i(x)| + |C^i(x) - C^i(y)| + |C^i(y) - C(y)|$. Since $C^i \rightarrow C$ (by Lemma 4.2) it suffices to take i sufficiently large to make the left and right term arbitrarily small. Use the continuity of $C^i(\cdot)$ to make the middle term small. Next, C is bounded from below as follows from Lemma 4.2.iii. The continuity and boundedness imply that C attains its minimum.

We next show that $S \in (s, t]$. By the above $C(x) = K$ for $x \leq s$. Therefore $S > s$. Furthermore, $S \leq t$, for suppose the contrary: $S > t$. Since S is a minimizer it must be that $C(S) < K$. Therefore (12) implies that $C(S) = \gamma(S)/\lambda + EC(S - D)$. Since $\gamma(S) > 0$, as $S > t$, it then follows that $C(S) > EC(S - D)$. But $EC(S - D) \geq \inf_{y \leq S} \{C(y)\}$. This would imply that $C(S) > \inf_{y \leq S} \{C(y)\}$, but this is impossible since $C(\cdot)$ attains its minimum at S . Contradiction. \square

4.4 Structure of Optimal Policy in \mathcal{H} .

In this section we use the properties of the solution of the dynamic programming equation (12) to obtain insight into structural properties of the optimal policy in \mathcal{H} .

We now take as cost function $c(x) = h(x) - g$ where $h \in \mathcal{B}_{\alpha, R}$ (α as in Lemma 4.1). By Lemma 4.1 there exists a unique solution $\gamma_g(\cdot)$ of (8) with this cost function. Write $C_g(\cdot)$ for the solution of (12) with cost-to-go function $\gamma_g(\cdot)$. We next prove that there is a revision rate g^* and a policy π_{g^*} such that g^* is equal to the average cost rate of a recurrent cycle induced by π_{g^*} ; recall the reasoning in Section 4.1.

Theorem 4.5. *There exist $g^* > 0$ such that $S_{g^*} \in (s_{g^*}, t_{g^*}]$, the minimal revised cost function $C_{g^*}(x) \geq 0$ for all x , and $C_{g^*}(S_{g^*}) = 0$. Moreover, there exists an optimal (g^* -minimizing) policy $\pi_{g^*} \in \mathcal{H}$, and this policy is such that $O_0^{\pi_{g^*}} \supset (-\infty, s_{g^*}]$ and $O_1^{\pi_{g^*}} = [S_{g^*}, \infty)$.*

Proof. We first show that, for fixed x , $\gamma_g(x)$ can be made arbitrarily negative as a function of g . Inserting the relation $\gamma_g(x) = \gamma_0(x) + \alpha$ into both sides of (8), where γ_0 is the solution of (8) with $g = 0$, and solving for α results in $\alpha = -g/(1 - \lambda E[D])$, since $\lambda \int_0^\infty \alpha G(y) dy = \lambda E[D]\alpha$. Therefore,

$$\gamma_g(x) = \gamma_0(x) - \frac{g}{1 - \lambda E[D]}. \quad (13)$$

Thus, we can assure there exists a g such that $\gamma_g(x) < 0$ for at least one x . The continuity of $\gamma_g(\cdot)$ ensures then that γ_g has at least two roots, with s_g and t_g as its left-most and right-most roots.

We next prove that there exists a g^* such that $C_{g^*}(x) \geq 0$ for all x , and $C_{g^*}(S_{g^*}) = 0$. By Lemma 4.4 $C_g(\cdot)$ attains its left-most minimum at some S_g . The continuity of $\gamma_g(x)$ in g implies the continuity of $C_g(x)$ in g . If $C_g(S_g) < 0$ for sufficiently large g the continuity in g implies the existence of a g^* such that $C_{g^*}(S_{g^*}) = 0$. (Recall that if $g = 0$, $C_0(x) \geq 0$ for all x .) A sufficiently large g can be found by noting that $\gamma_g(x)$ can be as small as we like by (13). Next, by subtracting K from both sides of (12), we see that $C_g(S_g) - K = \gamma_g(S_g) + E[C_g(S_g - D) - K]$. Now (12) implies also that $E[C_g(S_g - D) - K] \leq 0$. Therefore, $C_g(S_g) - K \leq \gamma_g(S_g)$. Hence, $C_g(S_g)$ is negative if we chose g such that $\gamma_g(S_g) \leq -K$.

By applying Theorem 4.3 the existence of the optimal policy π_{g^*} is guaranteed. From Lemma 4.4 we conclude that $S_{g^*} \in (s_{g^*}, t_{g^*}]$. Since $C_{g^*}(S_{g^*}) = 0$ it is best to switch off at S_{g^*} , so that $O_1^{\pi_{g^*}} = [S_{g^*}, +\infty)$, and $O_0^{\pi_{g^*}} \supset (-\infty, s_{g^*}]$ as $C_{g^*}(x) = K$ for $x < s_{g^*}$. \square

From the above we can conclude that the optimal policy π_{g^*} in \mathcal{H} prescribes to switch off when the inventory level is larger than or equal to S_{g^*} and to switch on when the inventory level is less than s_{g^*} .

Corollary 4.6. *If $C_{g^*}(x) < K$ on $x \in (s_{g^*}, S_{g^*}]$ the optimal policy $\pi_{g^*} \in \mathcal{H}$ has an (s, S) -structure with $s = s_{g^*}$ and $S = S_{g^*}$.*

Proof. An immediate consequence of theorem 4.5. \square

Corollary 4.7. *If $\gamma_{g^*}(x) < 0$ on $x \in (s_{g^*}, S_{g^*}]$ the optimal policy $\pi_{g^*} \in \mathcal{H}$ is an (s, S) -policy.*

Proof. Use (12) to see that $\gamma_{g^*}(x) < 0$ on $x \in (s_{g^*}, S_{g^*})$ implies $C_{g^*}(x) < 0$ on $x \in (s_{g^*}, S_{g^*})$. Next use Corollary 4.6. \square

Theorem 4.8. *The optimal policy $\pi_{g^*} \in \mathcal{H}$ has an (s, S) structure if $h(\cdot)$ is convex.*

Proof. If $h(\cdot)$ is convex, the function $c(\cdot) = h(\cdot) - g$ is convex. From (7) and induction we see that $\gamma_g^n(\cdot)$ is also convex for all n . This in turn implies that $\lim_n \gamma_g^n(\cdot) = \gamma_g(\cdot)$ is convex for all g . The convexity implies that γ_g has precisely two roots s_g and t_g for sufficiently large g . Since $\gamma_g(S_g) < 0$ it follows that $\gamma_g(x) < 0$ for all $x \in (s, S_g]$. Now invoke Corollary 4.7. \square

4.5 Structure of the Optimal Stationary Policy.

The first goal of this section is to show that no policy in the class of all stationary policies can improve $\pi_{g^*} \in \mathcal{H}$; thus there exists at least one optimal stationary policy, which is π_{g^*} . It then follows trivially from Theorem 4.8 that π_{g^*} has an (s, S) -structure when the inventory function $h(\cdot)$ is convex.

Theorem 4.9. *The optimal policy $\pi_{g^*} \in \mathcal{H}$ is also optimal in the class of stationary policies. If $h(\cdot)$ is convex, this policy has an (s, S) -structure.*

Proof. Define $S^\pi(x) = \min \{O_1^\pi \cap [x, \infty)\}$. Thus, $x \leq S^\pi(x)$ for all x , and if $(P(t), I(t)) = (1, x)$ for some t , production will eventually switch off at $S^\pi(x)$, and the inventory can never exceed $S^\pi(x)$ again. Note that $S^\pi(x) < \infty$ for all x , for otherwise the inventory would drift to infinity if the inventory were to start at or above x ; since it is also not optimal for the inventory to drift to $-\infty$, $S^\pi = \lim_{x \rightarrow -\infty} S^\pi(x) = \min\{x; x \in O_1^\pi\} > -\infty$ for any sensible stationary policy.

If F has infinite support the set $(-\infty, S^\pi)$ will be hit eventually, and from then on, level S^π will not be crossed any more, so that π cannot improve the minimizing policy in \mathcal{H}_{S^π} .

If F has finite support, however, the process may get stuck in some interval $(S^\pi(x), S^\pi(y)]$ if $I(t) = y > x$ for some t . Let $S^\pi = \operatorname{argmin}_x \{C_g(S^\pi(x))\}$, i.e., S^π is the inventory level at which $C_g(\cdot)$ achieves its global minimum. The stationary policy π can again not improve the best policy in \mathcal{H}_{S^π} .

Theorem 4.8 implies that π_{g^*} has an (s, S) -structure if h is convex. \square

5 An Application to Capacitated Periodic-review Inventory models

As the production-inventory model introduced by Graves [11] can be seen as the periodic-review analogon of the model presented in Section 3, we discuss in this section how to prove structure results for this model.

The model is as follows. At the beginning of each period the state of production is chosen to be on or off. If production is off at the end of a period, it can be switched on at the beginning of the next period at the expense of a setup cost K . However, when production is on in a period, keeping it on is free. Hence, setup costs are only incurred at the first period of a run of consecutive periods during which production is on. Production capacity is, without loss of generality, in single units. The demand size is distributed as a generic random variable D defined on the integers. Demand is met by end-of-period inventory, and, if necessary, backlogged. Thus, if the inventory level is I at the beginning of a period, production is on, and the demand is d , the inventory is $I + 1 - d$ at the end of the period. Inventory costs are accrued at the end of a period according to the function $h(\cdot)$. We assume that $h(\cdot)$ and D are such that $Eh(x - D) < \infty$ for all x .

Graves [11] conjectures that (s, S) -policies are optimal for this model. Here we show how to prove this conjecture under the assumption that inventory costs are convex and the expected inventory cost $Eh(x - D) < \infty$ for all x .

Similar to Section 4.2, let $V(x, y)$ be the cost to move from inventory level x to level y along the on-line. Then

$$V(x, y + 1) = V(x, y) + Eh(y + 1 - D) + EV(y + 1 - D, y + 1).$$

Letting

$$\gamma(y + 1) = V(x, y + 1) - V(x, y),$$

the above can be rewritten as

$$\gamma(y) = Eh(y - D) + EV(y - D, y).$$

Next, let $C(x)$ be the minimal cost to move from x at the off-line to level x at the on-line. Then C must satisfy the dynamic programming equations

$$C(x) = \min\{K, Eh(x - D) + EV(x - D, x) + EC(x - D)\},$$

since either production is switched on at cost K , or production remains off for at least one more period. The latter decision costs $Eh(x - D)$ due to the inventory, plus $EV(x - D, x)$ to return to level x after the demand occurred, plus the minimal cost $EC(x - D)$ that will be incurred after the demand while production is off. Using the above definition of γ , we can rewrite this such that

$$C(x) = \min\{K, \gamma(x) + EC(x - D)\}. \quad (14)$$

Observe the similarity to Eq.(12).

The last step is to replace $h(\cdot)$ in the above by the g -revised cost $h(\cdot) - g$. With these notions well-established, the rest of the procedure as explained in the previous sections can be applied straightaway.

Issues related to the existence and uniqueness of (optimal) value functions are simpler now since the process here is a discrete-time, discrete state process. As a first step, since $h(\cdot)$ is convex, γ is also convex. Take any x smaller than the left root s of γ . Since $\gamma(x) > 0$ for all $x < s$ and the inventory process can only drift to the left, it is evident from (14) that it is optimal to switch on immediately. Thus, $C(x) \equiv K$ for all $x < s$. Then, on the set $x \in \{s, s + 1, \dots, t\}$, where t is the right root of γ , the optimal control problem can be seen as the problem to minimize the expected cost until the target set $\{x < s\}$ hit. Since the cost-to-go γ on the set $\{s, \dots, t\}$ is not positive, provided g is chosen suitably, stopping before the set $\{x < s\}$ is not optimal. Finally, since the hitting time to the set $\{x < s\}$ has finite expectation for any policy and γ is finite between s and t , the cost to the set $\{x < s\}$ has finite expectation for any policy. The rest of the details are easy to provide. Finally, extending the above to rational-valued demand is also straightforward.

6 Numerical Considerations.

From the procedure explained in Section 4 and the results of Theorems 4.5, 4.9, and 4.8 it is apparent that the optimal cost and the optimal policy can be determined by bi-section. Suppose we have numbers g_l and g_u such that $g_l < g^* < g_u$. Let $g = (g_l + g_u)/2$. Equations (8) and (12) can then be used to calculate $\gamma_g(\cdot)$ with $c(\cdot) = h(\cdot) - g$ and $C_g(\cdot)$. If the minimum $C_g(S_g) \leq 0$ ($C_g(S_g) \geq 0$) the revised cost of the corresponding policy is smaller (larger) than or equal to g , and g should become the new upper-bound (lower-bound).

Obtaining initial values for g_l and g_u is easy. Pertaining to the former, since $h(x) \geq 0$ for all x , $g_l = 0$ is a proper choice. With respect to the latter, the average cost under any feasible stationary policy is an upper-bound on the optimal g^* . Hence, we can take the policy in \mathcal{H}_0 that switches the production on after the first order arrival and switches off as soon as point $(1, 0)$ is reached again. The expected cost from $(1, 0)$ until $(1, 0)$ is in this case equal to $K + \gamma(0)/\lambda$. The expected recurrence time is equal to $1/\{(1 - \lambda E[D])\lambda\}$. Hence, the average cost under this policy is equal to $(K\lambda + \gamma(0))(1 - \lambda E[D])$. Setting g_u to this value would do for our purpose.

With respect to the computation of $\gamma_g(x)$ at some point x , observe that (8) requires to know $\gamma_g(y)$ for all $y < x$, which complicates the computational procedure. When $h(\cdot)$ is a (convex) polynomial, this problem is still relatively easy to resolve. By assumption h attains its minimum at $x = 0$. Then substitute in (8) a polynomial of suitable degree for $\gamma_g(y)$ for $y \leq 0$ and solve for the coefficients on the half-line $(-\infty, 0]$. For instance, if $h(x) = bx$, $b < 0$, on $x \leq 0$, this results in

$$\gamma_g(x) = -\frac{bx + g}{1 - \lambda E[D]} + \frac{\lambda}{2} \frac{bE[D^2]}{(1 - \lambda E[D])^2}, \quad \text{on } x \leq 0. \quad (15)$$

Now we can use this part of γ_g to obtain γ_g , possibly numerically, on $(0, \infty)$. However, when $h(\cdot)$ is not a polynomial, solving (8) is more complicated. To handle such situations we refer to the appendix for an approximation method based on a discretization of the model. There we also prove that this approximation converges exponentially fast (as a function of a suitable parameter) to the correct solution of (8).

7 Examples

In this section we first illustrate the procedure to find an optimal policy as described in Section 4. We then discuss two examples from the literature. As a fourth example we show that the optimal policy to switch on is not monotone in general. The final example is a counter-example against the claim that (s, S) -policies are optimal for general quasi-convex inventory cost functions¹.

7.1 Bisecting the Revision Cost

As a first example, consider inventory cost

$$h(x) = b \min\{-x, 0\} + h \max\{x, 0\}, \quad (16)$$

with $b = 3$, $h = 1$, arrival rate $\lambda = 3/2$, setup cost $K = 5$, and uniformly distributed demand on $[0, 1]$ so that $E[D] = 1/2$ and $E[D^2] = 1/3$. In this case, (15) becomes $\gamma_g(x) = 12(1-x) - 4g$ on $x \leq 0$. The (numerical) evaluation of (8) and (12) is now easy for given g . (We use $\Delta = 0.005$ for the integrations.) Figure 2 shows the graphs of $\gamma_g(\cdot)$ and $C_g(\cdot)$ for $g = 2.1$, $g = 2.3$ and $g = 2.5$. The left panel shows that $g = 2.1$ is sufficiently small for $\gamma_g(\cdot)$ to have two roots, and indeed, $C_g(x) < K$ for some x . Apparently, $g = 2.1$ is not large enough to ensure that there exists S_g such that $C_g(S_g) = 0$. As a next guess, we try $g = 2.5$. The right panel shows now that C_g has two roots, implying that $g = 2.5$ must be too large. By bi-section $g = 2.3$ should be the next guess. Indeed, the middle panel shows that $C_{2.3}(S_{2.3})$ is nearly 0. Continuing the iteration leads ultimately to $g^* \approx 2.31$, $s = s_{g^*} = 0.480$ and $S = S_{g^*} = 2.610$.




Figure 2: Graphs of $\gamma_g(\cdot)$ and $C_g(\cdot)$ for $g \in \{2.1, 2.3, 2.5\}$.

7.2 Computing the optimal D for an $M/G/1$ Queue

As a second example, we show how to compute the optimal threshold for the D -policy considered by Feinberg and Kella [4]. We take as inventory cost function $h(x) = \infty$ for $x > 0$, and non-increasing for $x \leq 0$ such that $h(0) = 0$. In this way, $S = 0$, and $-I(t)$ equals the workload of the $M/G/1$ queue, and the level s below which production switches on is equal to (minus) the workload level D above which the server switches on. Now, it is easy to see that, since h is non-increasing, γ has just one left-most root for all $g > 0$. This optimal root s^* for the inventory system is equal to the optimal level $-D$ for the queueing system.

7.3 Results of Graves and Keilson [10]

Graves and Keilson [10] consider exponentially distributed demands with cost function (16) for various values of b , $h = 1$, setup costs and arrival rates. Our results correspond to theirs, but slightly improve the

¹The source code (in python/numpy) and all examples are available at the first author's homepage [25]

minimal average costs since we do not limit s and S to the integers as is the case in Graves and Keilson [10]. There appears to be a small typo in one of their examples: namely the case with $\lambda = 0.9$, $b = 2$ in (16), and $K = 25$. They claim that $s = -1$, while we get $s \approx 5.1$. For the same case with $K = 50$, they find $s = 6$. The result $s = -1$ when $K = 25$ is also at variance with their remark that an increase in K should result in an increase in $S - s$, and should hardly affect s . Thus, in both cases the s 's should be roughly the same.

7.4 An Optimal Policy that is not Monotone

The fourth example, see the left panel of Figure 3, proves that it is not true that for all g the optimal decision structure on the off-line is monotone. Here the demand D is uniform on $[1.02, 1.12]$, $\lambda = 0.2$, $h(\cdot)$ is given by (16) with $b = 4$ and $h = 1$, $K = 5$ and $g = 1.7$. Clearly, $C_g(x) < K$ for $x \in (s, S]$. Next, note that $C_g(3) = K$, but $C_g < K$ at the small dent at the right of $x = 3$. Thus, at $x = 3$ it is optimal to switch on, but at the dent it is optimal not to switch on. Thus, the decision structure is not monotone, but alternates between on and off several times. Since $C_g(x) < K$ for $x \in (s, S]$ the optimal policy does not alternate between on and off in $(s, S]$, so that the optimal policy still has an (s, S) -structure. Note also that C_g is not monotone decreasing between s and S .




Figure 3: The left panel shows C_g when the demand is deterministic $D \equiv 1$. The appearance of the small dent at the right of S implies that the optimal policy is not monotone. The right panel shows C_g for quasi-convex holding costs. The fact that the value function hits K between s and S implies that the optimal policy is not (s, S) .

Thus, even when the cost function is convex, the value function is for general g not monotone.

7.5 An Optimal Policy that is not (s, S)

The last example shows that the combination of the demand distribution and the cost function determine the structure of the optimal policy. In this particular case, the optimal policy is not (s, S) but more difficult. The parameters are the same as the previous example, the demand is exponentially distributed with mean 2 and the inventory cost $h(\cdot)$ is given by the quasi-convex function

$$h(x) = (4\{x \leq 0\} + \{x > 0\})\lceil x \rceil,$$

where $\{\cdot\}$ is the indicator function and $\lceil x \rceil$ is the smallest integer equal to or larger than x . The results are shown in the right panel of Figure 3.

We see that $s^* = 0.61$, and the minimum of C_{g^*} occurs at $S^* = 2.$, but $C_{g^*}(x) = K$ for some $x \in (s^*, S^*)$, for instance at $x \approx 1.$. This implies that the optimal policy cannot be an (s, S) -policy since at it is optimal to switch on between s^* and S^* .

Interestingly, when the demand is exponentially distributed with mean 1, the optimal policy is (s, S) again. Hence, convexity of the cost function is sufficient, but by no means necessary.

8 Summary and Suggestions for Further Research.

We summarize the basic elements of the procedure we developed in this paper to prove that (s, S) -policies are long-run average optimal in the class of stationary policies for the stochastic EPQ model. As this procedure allows the analysis of more general inventory problems with restricted production capacity, which we discuss below, we enumerate the involved steps.

1. Subtract a revision rate g from the inventory cost rate $c(\cdot)$ to obtain the g -revised inventory cost. Formulate the minimal cost problem to move from level x at the off-line to level x on the on-line under some given cost rate $c(\cdot)$ as an optimal stopping problem on the off-line. Prove that for all g a g -minimizing policy and function $C_g(x)$ for this optimal stopping problem exists.
2. Use properties of the optimal stopping problem to find the minimal overall cost to move from the off-line to the on-line, i.e., find S_g such $C_g(S_g) = \min_x C_g(x)$. If there is more than one such point, take S_g as the left-most point. At S_g it is best to stop producing when production is on. This, in effect, starts a new production cycle.
3. Exploit other structural properties of the optimal stopping problem to find the best point s_g at which to switch on production. For our problem we take s_g as the left root of the cost-to-go $\gamma_g(\cdot)$, provided the conditions of Corollary 4.7 are satisfied.
4. Then use bisection over g to show that a policy can be found that performs arbitrarily closely to the optimal policy. In the limit an optimal policy results. The structural properties derived in the previous two points can be finally used to prove the (s, S) -structure of the optimal policy.

Interestingly, this procedure not only enables us to derive formal results, it also provides an efficient method to compute the optimal (s, S) -policy to arbitrary precision. In fact, while completing the above steps typically involves numerous technical arguments to show the existence and uniqueness of solutions of dynamic programming equations, proper limits, and so on, for numerical purposes there appears no harm in simply skipping these technical desiderata. Simple graphical arguments involving the interplay between the revision rate g and the resulting numerical solution $C_g(\cdot)$ of the optimal stopping problem, such as in Section 7, generally provide considerable insight into the structure of the optimal policy. For instance, as Corollaries 4.6 and 4.7 make clear, the convexity of $h(\cdot)$ is quite more than we actually need to prove that an (s, S) policy is stationary optimal. Any other argument, possibly numerically, that ensures that $C_{g^*}(x) < K$ on $x \in (s_{g^*}, S_{g^*}]$, where s_{g^*} is the left-most root of the cost-to-go function $\gamma_{g^*}(\cdot)$ and S_{g^*} is the left-most most minimizer of $C_{g^*}(\cdot)$, suffices.

The above procedure seems to open roads to explore more complicated situations, such as more general arrival processes (e.g., Markovian Arrival Processes (MAPs), state dependent production rates or setup costs, or other control rules such as order acceptance. Introducing an order acceptance mechanism brings the system studied here closer to the stock rationing system investigated by Ha [12, 13]. Ha uses a queueing approach, but it seems also possible to use our on-off production approach. The case without set-up cost is non-trivial already, but maybe our approach makes it possible to combine stock rationing with set-up cost. Other directions for future research are to include set-up times and vacations. Incorporating multiple production levels is also an interesting theoretical extension, c.f., Lu and Serfozo [17], but it is unclear whether in such systems our approach with cycles and stopping problems still works.

9 Acknowledgments.

We are grateful to the (anonymous) referees and the associate editor for their many helpful insights and suggestions to improve the paper. We thank Remco Germs with providing us with one of the numerical examples.

References

- [1] A. Federgruen and K.C. So. Optimality of threshold policies in single-server queueing systems with server vacations. *Adv. Appl. Prob.*, 23:388–405, 1991.
- [2] A. Federgruen and Y.S. Zheng. Optimal control policies for stochastic inventory systems with endogenous supply. *Probability in the Engineering and Informational Sciences*, 7:257–272, 1993.
- [3] A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands I. the average-cost criterion. *Math. Oper. Res.*, 11(2):193–207, 1986.
- [4] E.A. Feinberg and O. Kella. Optimality of D -policies for an $M/G/1$ queue with a removable server. *Queueing Systems*, 42(4):355–376, 2002.
- [5] M.C. Fu. Sample path derivatives for (s, S) inventory systems. *Operations Research*, 42(2):351–364, 1994.
- [6] G. Gallego and A. Scheller-Wolf. Capacitated inventory problems with fixed order costs: Some optimal policy structure. *European Journal of Operation Research*, 126:603–613, 2000.
- [7] G. Gallego and L.B. Toktay. All-or-nothing ordering under a capacity constraint. *Operations Research*, 52(6):1001–1002, 2004.
- [8] B. Gavish and S. C. Graves. A one-product production/inventory problem with continuous review policy. *Operations Research*, 28:1228–1236, 1980.
- [9] B. Gavish and S. C. Graves. Production/inventory systems with a stochastic production rate under a continuous review policy. *Computers and Operations Research*, 8:169–183, 1981.
- [10] S. C. Graves and J. Keilson. The compensation method applied to a one-product production/inventory problem. *Math. Oper. Res.*, 6(2):246–262, 1981.
- [11] S.C. Graves. The multi-product production cycling problem. *AIIE Transactions*, 12(3):233–240, 1980.
- [12] A. Y. Ha. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science*, 43:1093–1103, 1997.
- [13] A. Y. Ha. Stock rationing in an $M/E_k/1$ make-to-stock queue. *Management Science*, 46:77–87, 2000.
- [14] A. Hordijk and F. A. Van der Duyn Schouten. On the optimality of (s, S) policies in continuous review inventory models. *SIAM J. Appl. Math.*, 46(5):912–929, 1986.
- [15] W.T. Huh and G. Janakiraman. (s, S) optimality in joint inventory-pricing control: an alternative approach. *Operations Research*, 56(3):783–790, 2008.
- [16] H.S. Lee and M. M. Srinivasan. The continuous review (s, S) policy for production/inventory systems with Poisson demands and arbitrary processing times. Technical Report TR 87-33, Dep. of Industrial and Operations Engineering, University of Michigan, Ann Arbor, 1989.
- [17] F. Lu and R. Serfozo. $M/M/1$ queueing decision process with monotone hysteretic optimal policies. *Operations Research*, 32(4):1117–1132, 1984.
- [18] E. Presman and S.P. Sethi. Inventory models with continuous and Poisson demands and discounted and average costs. *Production and Operations Management*, 15(2):279–293, 2006.
- [19] C. Shaoxiang. The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy. *Operations Research*, 52(3):409–421, 2004.

- [20] C. Shaoxiang and M. Lambrecht. $X - Y$ band and modified (s, S) policy. *Operations Research*, 44(6):1013–1019, 1996.
- [21] A.N. Shiryaev. *Optimal Stopping Rules*. Springer, 2007.
- [22] M. J. Sobel. Optimal average-cost policy for a queue with start-up and shut-down costs. *Operations Research*, 17:145–162, 1969.
- [23] M.M. Srinivasan and H.S. Lee. Random review production/inventory systems with compound Poisson demands and arbitrary processing times. *Management Sci.*, 37(7):813–833, 1991.
- [24] H. C. Tijms. *A First Course in Stochastic Models*. John Wiley & Sons, Chichester, 2003.
- [25] N.D. Van Foreest. <http://nicky.vanforeest.com>, 2012.
- [26] J. Wijngaard. An inventory problem with constrained order capacity. Technical report, Eindhoven University of Technology, 1972.
- [27] J. Wijngaard. *Stationary Markovian Decision Problems*. PhD thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, 1975.
- [28] J. Wijngaard and S. Stidham Jr. Forward recursion for Markov decision processes with skip-free-to-the-right transitions, part I: Theory and algorithm. *Math. Oper. Res.*, 11(2):295–308, 1986.
- [29] J. Wijngaard and S. Stidham Jr. Forward recursion for Markov decision processes with skip-free-to-the-right transitions, part II: Non-standard applications. *Statistica Neerlandica*, 54(2):160–174, 2000.
- [30] H.S. Wilf. *Generatingfunctionology*. Academic Press, 2nd edition, 1994.

A A Numerical Approximation for $\gamma_g(\cdot)$.

To cope with a situation in which $h(\cdot)$ is not a polynomial for $x \leq 0$ we use the following scheme to compute a numerical approximation for $\gamma_g(\cdot)$. First, discretize the inventory process (and the demand) to take values on a grid with points separated by some small $\Delta > 0$, and write $g_i = G((i+1)\Delta) - G(i\Delta)$, $d_i = \gamma_g(i\Delta)$, etc., where we temporarily drop the dependence on g for notational convenience. Second, as in Shaoxiang [19], assume that the demand has an upper-bound, so that there exists a finite $n := \min\{i : g_i = 0\}$. Third, for given $g > 0$, choose some z considerably smaller than the left root s_1^g of $h(\cdot) - g$. Finally, let

$$\tilde{d}_i = 0, \quad \text{on } i \leq z,$$

and define recursively

$$\tilde{d}_i = h_i - g + \lambda \sum_{j=0}^n \tilde{d}_{i-j} g_j, \quad \text{on } i > z. \quad (17)$$

Observe that it is easy to transform this into a simple recursion:

$$\tilde{d}_i = \frac{h_i - g}{1 - \lambda g_0} + \frac{\lambda}{1 - \lambda g_0} \sum_{j=1}^n \tilde{d}_{i-j} g_j, \quad i > z.$$

We next show that the difference $\delta_i = d_i - \tilde{d}_i$ for i fixed converges exponentially fast to 0 as $z \rightarrow -\infty$. Therefore, by taking z sufficiently small we can make the difference δ_i as small as we wish in the region of interest.

Subtracting (17) from the discretization of (8) leads to

$$\delta_i := d_i - \tilde{d}_i = \begin{cases} d_i, & i \leq z, \\ \lambda \sum_{j=0}^n \delta_{i-j} g_j, & i > z. \end{cases}$$

We therefore have to prove that $\delta_i \rightarrow 0$ exponentially fast as $z \rightarrow -\infty$.

This convergence is actually not immediate for the following reason. By the boundary conditions, $\delta_i = d_i$ for $i \leq z$. The assumption $\gamma_g \in \mathcal{B}_{\alpha,R}$, i.e., $\sup\{|\gamma_g(x)e^{\alpha x}|; x \geq 0\} < \infty$, gives the bound $d_i \leq Me^{-\alpha i \Delta}$ for some M . Thus, d_z may increase exponentially fast as $z \rightarrow -\infty$.

The main idea is to prove that $\delta_i = O(x_0^{-i})$ for $i \geq z$ for some constant $x_0 > 1$ to be determined below. Then it follows from the boundary condition $\delta_z = d_z$ that $\delta_i = O(\delta_z x_0^{z-i})$. Since $\gamma_g \in \mathcal{B}_{\alpha,R}$ it also follows that $d_z \leq Me^{-\alpha z \Delta}$ for some M . Combining these two estimates we see that $\delta_i \leq M' (x_0 e^{-\alpha \Delta})^z x_0^{-i}$ for some M' . The last step is to show that $e^{\alpha \Delta} < x_0$, so $e^{-\alpha \Delta} x_0 > 1$, and indeed $\delta_i \rightarrow 0$ as $z \rightarrow -\infty$.

To obtain x_0 , take the Z transform of both sides of the equation $\delta_i = \lambda \sum_{j=0}^n \delta_{i-j} g_j$, and write $\Delta(w) = \sum_{i=0}^{\infty} \delta_i w^i$. Some algebra leads to the equation $\Delta(w) = N(w)/D(w)$ where the numerator $N(w)$ is some polynomial in w , and the denominator $D(w)$ has the form

$$D(w) = 1 - \lambda \sum_{i=0}^n g_i w^i.$$

Let $x_0 \in \mathbb{R}$ solve $D(x_0) = 0$. Provided that x_0 is the root with the smallest modulus, it follows that $\delta_i = O(x_0^{-i})$, see e.g. Wilf [30, Theorem 2.4.3]. To see that indeed there is no w inside $C = \{z \in \mathbb{C}; |z| = x_0\}$ such that $D(w) = 0$, we apply Rouché's theorem. Recall that this states that if there exists some analytic f such that $|f(w) - D(w)| < |f(w)|$ for all w on some simple closed contour, then $f(\cdot)$ and $D(\cdot)$ have the same number of zeros within this contour. In our case, take $f \equiv 1$, take the contour as $C_\epsilon = \{z \in \mathbb{C}; |z| = x_0 - \epsilon\}$ for some $\epsilon > 0$, and take $w \in C_\epsilon$. Then, $|1 - D(w)| = |\lambda \sum_{i=0}^n g_i w^i| \leq \lambda \sum_{i=0}^n g_i |w|^i < \lambda \sum_{i=0}^n g_i x_0^i = 1 = |f(w)|$. Since f has no zeros in C_ϵ , and ϵ is arbitrary, we are done.

To see that $x_0 > 1$, note that since $g_j > 0$ the restriction of $D(w)$ to the real line must be decreasing when $w > 0$. Hence, since $D(1) = 1 - \lambda \sum_{i=0}^n g_i = 1 - \rho > 0$, its root $x_0 > 1$.

We finally show that $e^{\alpha \Delta} < x_0$. By the model assumptions α is such that $\lambda \sum_{i=0}^n (e^{\alpha \Delta})^i g_i = \beta < 1$, that is $D(e^{\alpha \Delta}) > 0$. Since D is decreasing on $[0, \infty)$ (and $e^{\alpha \Delta}$ is real), this implies $e^{\alpha \Delta} < x_0$, as $D(x_0) = 0$, and x_0 is the root with smallest modulus.

The proof is complete.