University of Groningen

Optimal policies for production-clearing systems under continuous-review

Germes, Remco; van Foreest, Nicky D.; Kilic, Onur A.

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2016.05.004](https://doi.org/10.1016/j.ejor.2016.05.004)

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:
2016

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

Citation for published version (APA):

Germes, R., van Foreest, N. D., & Kilic, O. A. (2016). Optimal policies for production-clearing systems under continuous-review. *European Journal of Operational Research*, 255(3), 747-757.
<https://doi.org/10.1016/j.ejor.2016.05.004>

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.

Production, Manufacturing and Logistics

Optimal policies for production-clearing systems under continuous-review

Remco Germs^a, Nicky D. Van Foreest^{b,*}, Onur A. Kilic^{b,c}^a Logistics, Decision and Information Sciences, Wageningen University, Wageningen, The Netherlands^b Department of Operations, University of Groningen, Groningen, The Netherlands^c Institute of Population Studies, Hacettepe University, Ankara, Turkey

ARTICLE INFO

Article history:

Received 1 October 2015

Accepted 3 May 2016

Available online 11 May 2016

Keywords:

Production-clearing systems

Inventory systems

Stochastic demand

Continuous review

Optimal policies

ABSTRACT

In this paper, we consider a production-clearing system with compound Poisson demand under continuous review. The production facility produces one type of item without stopping and at a constant rate, and stores the product into a buffer to meet future demand. To prevent high inventory levels, a clearing operation occasionally removes all or part of the inventory from the buffer. We prove that an (m, q) -policy, i.e., a policy that clears the buffer to level m as soon as the inventory hits a level q , minimizes the long run average holding and clearing cost. We also derive a numerically very efficient approach to compute the optimal parameters of the (m, q) -policy for models with backlogging and models with lost sales. With these numerical methods we show that tuning the clearing levels m and q in concert can lead to substantial cost savings.

© 2016 Published by Elsevier B.V.

1. Introduction

Stochastic production-clearing systems were introduced and studied in (Berman, Parlar, Perry, & Posner, 2005) and (Perry, Stadje, & Zacks, 2005) to generalize the classical clearing models investigated by Stidham (1974; 1977; 1986). In these production-clearing systems, a single machine produces at a fixed rate and without stopping a certain product into a buffer. The cumulative demand, for instance generated by retailers, is modeled as a compound Poisson process. Demand that cannot be satisfied from on-hand stock is backordered or lost, depending on the details of the model. Typically, the production rate is larger than the demand rate. Since production never stops, it is necessary to prevent inventory levels to grow without bound. This is realized by a clearing policy that occasionally prescribes to clear a part or the complete inventory, for instance by selling the surplus to a large wholesaler. Costs are incurred for holding and backordering (or rejection) demand, and there is a fixed cost and a variable cost related to the clearing operation to represent the difference in price per unit stock as offered by the wholesaler and the retailers (c.f., Kim & Seila, 1993). The problem is to identify the clearing policy that, on the long run, optimally balances the inventory and clearing costs.

1.1. Background

Production-clearing inventory models are useful to analyze inventory systems in which the supply of raw materials typically exceeds the demand or when it is too expensive or difficult to switch off or reduce the production rate. As the examples below show, such situation typically occurs when a number of production steps are tightly coupled, and one of these steps is confronted with the excess supply produced as a by-product of one of the other steps. Since these other steps are economically dominant, the excess supply cannot be reduced or avoided. The only way to keep the inventories of the by-products within reasonable limits is therefore to occasionally dispose of it or sell it at a discount.

Observe that production-clearing situations are very different from 'regular' inventory systems: in such systems the problem is to control the input, in the form of replenishment orders, such that customer demand is met and holding and ordering costs are minimized. In production-clearing system the input rate of material or the production cannot be controlled, and there are only two types of decision available: when to clear and to what level to clear the inventory. Interestingly, it might seem that a production-clearing system behaves like a production-inventory system run backwards in time. (Fig. 1 below may suggest this) This is, however, not the case, since in the latter type of system production can be switched off, for instance.

Production-inventory systems controlled by an (s, S) -policy are related to production-clearing systems in which the clearing level

* Corresponding author. Tel.: +31 50 363 51 78.

E-mail address: n.d.van.foreest@rug.nl (N.D. Van Foreest).

Fig. 1. A sample path for the inventory level process $(I_t)_{t \geq 0}$ under the (m, q) policy. The duration of one regenerative cycle is τ_q .

$q = S$ and the remaining level $m = s$. In both systems an action is chosen when the inventory level hits S , or equivalently q . The fundamental difference is that in production-clearing systems, inventory is cleared instantaneously to level m , whereas in production-inventory systems the inventory level decreases gradually as demand arrives and the production capacity is switched off.

1.2. Examples

A first, simple example of a system in which the rate of raw material cannot be controlled is the dairy industry in which animals (e.g. cows, goats) are the principle producers. It is obviously not possible to reduce the production of milk overnight, just because the inventory levels of milk or cheese are too high or the demand is lagging.

The production of cheese and powdered milk forms another interesting example. A major by-product of making cheese is whey. The whey is, typically, transported to another, downstream, factory which turns the whey into powdered milk. Up to several years ago, the demand for cheese exceeded the demand for whey, so that the cheese-making step was the main profit source in the dairy industry. The downstream factory was confronted with a steady input of whey, which they could not control, but still had to handle. Interestingly, the last years have seen a meteoric increase in demand for powdered milk, in fact, to the extent that the cheese-making step has become just a necessary pre-processing step for whey, at least from an economic perspective. Now the upstream, cheese-making, factory is confronted with an input that exceeds the demand for cheese. Thus, in either one of the situations, one factory has to deal with an inflow of raw materials whose rate they cannot control, except by occasionally clearing the inventory of raw materials or the inventory of finished items.

Yet another example is the production and storage of electricity by small combined heat and power systems that interact with either a heat or electricity inventory system, [Larsen, Van Foreest, and Scherpen \(2014\)](#). When there is heat demand, electricity is co-produced at a steady rate, and, as long as there is heat demand, the production of electricity cannot be stopped. This electricity can be consumed, if there is demand by local devices, or stored, but only up to certain extent. When the electricity buffer is full, (part of) the stored electricity can be sold to the electricity grid, thereby implementing a clearing action. Similar clearing policies can be applied to the production and storage of electricity generated by solar or wind sources. Here the production can also not be stopped, and it may be economically viable to clear the electricity buffers when full.

In other industry branches, such as in paper-making, steel-casting or glass-production, the setup costs or setup times of the factories are (very) large. In such continuous production environments switching off is, although possible in principle, economically undesirable, and the costs of occasionally clearing the inventory are much less than any costs or down-times associated with switching off the plant. A case in point is the occasional selling steel, or other bulk material, below production prices.

Finally, [Barron \(2015\)](#) also discusses a number of interesting examples.

1.3. Related literature

[Berman et al. \(2005\)](#) consider a production-clearing system with exponentially distributed demand sizes and controlled by a policy that clears all inventory in the buffer when the inventory level hits or exceeds a critical level q . For this system, they minimize the long-run average cost as a function of the clearing level q . The average cost function is stated in terms of the steady-state distribution of the inventory level, which is obtained by using level crossing arguments. This work is a useful starting point in the analysis of continuous-review production-clearing systems, but it is of paramount importance to generalize it in at least three directions. First, it is important to consider policies that clear the inventory to a positive level. For instance, instead of selling all inventory to the wholesaler, it might be more profitable to sell only part of the inventory and keep the remaining inventory to forestall high backlog or lost-sales cost (c.f., [Stidham, 1977](#)). Second, demand sizes often have variation coefficients considerably different from 1. As a consequence, it is necessary to study the influence of the distribution of the demand sizes on the optimal policy and the associated minimal costs. Third, the inventory cost model can be considerably more complicated than the linear cost functions considered by [Berman et al. \(2005\)](#). For instance, the cost functions considered in ([Porteus, 2002, Exercise 9.5](#)) are piecewise continuous but not quasi-convex. [Barron \(2015\)](#) discusses a production-clearing system in which the production switches between a number of rates. However, demand cannot be backlogged and the author assumes that the inventory is totally cleared. We allow for backlogging, deal with various different loss policies, and consider more general clearing policies.

For more information on related production-inventory systems we refer to [Liu and Cao \(1999\)](#), [Perry and Posner \(2002\)](#), [Wang, Cao, and Liu \(2002\)](#), [Van Foreest and Wijngaard \(2014\)](#) and [Germs and Van Foreest \(2013b\)](#).

1.4. Overview and results

In view of the points above, the main concerns of this paper are (1) to derive a numerically efficient approach for finding the average-cost optimal policy for a production-clearing system, (2) to allow for generally distributed demand sizes, and (3) to allow for piecewise continuous inventory cost functions. To achieve these goals, we first of all show that the optimal clearing policy is of (m, q) -type. Such policies restore the inventory level after clearing to a fixed level $m \leq q$. To efficiently compute the optimal (m, q) policy for general production-clearing systems we generalize a method recently developed by [Van Foreest and Wijngaard \(2014\)](#) and [Germs and Van Foreest \(2013b\)](#). This method uses the concept of g -revised holding cost, as introduced by [Wijngaard and Stidham \(1986\)](#), to reformulate an optimization problem into an equivalent optimal stopping problem which is (numerically) much easier to analyze. We show in [Section 4](#) how the optimal (m, q) clearing policy can also be computed in principle by the level crossing approach discussed in ([Berman et al., 2005](#)), and we provide closed form expressions for the stationary distribution of the

inventory level process when the demand sizes are exponentially distributed. However, the procedure based on level crossing involves a two dimensional search for m and q to find an optimal policy. In contrast, our method only requires a one-dimensional search in a bounded interval which can be done efficiently by using a bisection scheme.

We apply our approach to a production-clearing model in which demand is backlogged if it cannot be met from on-hand stock. We show that when the utilization is high, the coefficient of variation of the individual demand sizes has a large effect on the optimal policy parameters m^* and q^* . Moreover, it turns out that it is important to tune m and q simultaneously. Even relatively small deviations from the optimal parameters choices can lead to substantial cost increases.

When the load of the system is high and the variability in demand is outside the producer's control, the producer basically has two options to improve the performance of the system. First, the producer can buy additional capacity, thereby reducing the load of the system and the sensitivity of the average cost towards variability in demand. Second, the producer can reduce the load of the system by rejecting, rather than backlogging, part of the demand. We also analyze clearing systems with order rejection and show numerically that judiciously rejecting demand can reduce average cost up to some three times as compared to a situation in which demand is backlogged, and this cost is less sensitive to the coefficient of variation of the demand. Thus, order rejection seems to be a viable policy to keep the average cost within reasonable limits. We should, however, make the proviso that systems with backlogging and systems with lost sales are difficult to compare due to fundamental differences in the performance measures.

1.5. Structure

The structure of the paper is as follows. In Section 2 we formulate models for production-clearing systems with backlogging and lost sales and prove that (m, q) policies are long-run average optimal. In Section 3 we demonstrate our computational approach to identify optimal clearing policies for backlog model and lost sales models. In Section 4 we consider the special case of a production-clearing system with backlogging and use level crossing arguments to derive a system of integral equations to be satisfied by the stationary distribution of the inventory level process. For the case with exponentially distributed demand sizes, we derive a closed-form expression for the average cost per unit time under an (m, q) policy. Section 5 provides an extensive numerical analysis of models with backlogging and models with lost sales. In Section 6 we summarize our results and provide directions for future research.

2. Models and optimality

In Section 2.1 we describe the production-clearing system with backlogging in detail and introduce the optimization problem that is to be solved. In Section 2.3 we modify this model such that demand that cannot be satisfied from on-hand stock is lost, rather than backlogged. In the last section we prove that an optimal stationary policy has an (m, q) form.

2.1. Models with backlogging

We consider the setting in which one machine produces a certain product at continuous and constant rate and the product is stored in a buffer. As the inflow is assumed to be deterministic and proportional to time we can, without restriction of generality, take the input up to time t as equal to t . Customers (retailers) arrive at the buffer according to a Poisson process; the demand up

to time t is given by a compound Poisson process

$$X_t = \sum_{n=1}^{N_t} Y_n,$$

where $X_0 = 0$, the $(Y_n)_{n \geq 1}$, are i.i.d. positive random variables with common distribution function F and expectation $\mathbb{E}Y$, and $(N_t)_{t \geq 0}$ is an ordinary Poisson process of intensity $\lambda > 0$ starting at $N_0 = 0$. For later use we define the survival function $G(\cdot) = 1 - F(\cdot)$. We require for the backlog model that $1 > \lambda \mathbb{E}Y$, for otherwise production cannot keep up with demand, and the inventory level will drift to $-\infty$.

Let $(I_t)_{t \geq 0}$ be the inventory process. If the inventory starts at level I_0 the level at time t , provided no clearing action occurs before t , is given by

$$I_t = I_0 + t - X_t, \quad (1)$$

as the machine produced t units of product up to time t and the total demand is X_t . To avoid high inventory levels, we assume that the system is cleared as soon as the inventory reaches some pre-specified threshold q . Let τ_q be the (stopping) time that the inventory is cleared for the first time, i.e.,

$$\tau_q = \inf\{t \geq 0 \mid I_t \geq q\}.$$

Under the (m, q) clearing policy, the clearing operation resets the inventory to the level m so that $I_{\tau_q} = m$ and

$$I_t = m + t - \tau_q - (X_t - X_{\tau_q}) \quad (2)$$

for $t \geq \tau_q$ until the next clearing action, and so on. A typical realization of $(I_t)_{t \geq 0}$ is depicted in Fig. 1.

Observe that under an (m, q) clearing policy, the inventory process $(I_t)_{t \geq 0}$ is a regenerative process with expected cycle time $\mathbb{E}_m \tau_q$. Here, $\mathbb{E}_x$ be the expectation of functionals of the inventory process $(I_t)_{t \geq 0}$ starting at $I_0 = x < q$.

The cost structure for the model consists of two parts. First, each time the buffer is cleared, the system incurs a cost given by the function $K(m, q)$ which depends on the reset level m and clearing threshold q as

$$K(m, q) = K + c(q - m), \quad (3)$$

where K is some positive constant and c the cost associated with discounting an amount of $q - m$ of inventory, (c.f., Kim & Seila, 1993). Second, cost accrue at rate $h(\cdot)$; when the inventory level is $x \geq 0$, $h(x)$ is the cost per unit stock on hand per unit time, when $x < 0$, $h(x)$ corresponds to the cost per unit backlog per unit time. We require that h is piece-wise continuous, non-negative, $h(0) = 0$, and $h(x) = o(|x|^n)$ for some $n > 0$ as $|x| \rightarrow \infty$.

For a given (m, q) clearing policy the expected cost of one regenerative cycle, starting at $I_0 = m$ and clearing the inventory at the instant τ_q , is given by

$$V(m, q) = \mathbb{E}_m \left(\int_0^{\tau_q} h(I_s) ds \right) + K(m, q). \quad (4)$$

Our problem is to determine

$$\min_{m, q} \frac{V(m, q)}{\mathbb{E}_m \tau_q}, \quad (5)$$

and to identify the optimal (m^*, q^*) policy that achieves this minimum.

Remark 1. After hitting q it is formally possible to clear the inventory level to some negative level $m < 0$. However, this appears rather odd from a practical perspective. The aim of clearing is to do away with temporary excess inventory, not to satisfy demand at the expense of backlogging cost. Moreover, why would a wholesaler, who is supposed to absorb the excess inventory, clear the inventory to a negative level? In view of this, even though it may

be optimal in some parameter settings to clear to negative levels, we constrain the reset level m to the interval $[0, q]$ in Section 5 in which we study the characteristics of the optimal policy.

We emphasize that our method to find the optimal policy, c.f., Section 3, does not depend on m being positive.

Remark 2. The optimization procedure as set out in Section 3 allows $K(m, q)$ to be of quite general form, but the linear form in (3) appears to be the most natural from a practical point of view.

Remark 3. There is subtle problem with the formal construction of the (sample paths of) (I_t) in Eqs. (1) and (2). For the inventory to actually hit level q it is necessary that the inventory process is left-continuous at q . However, we define $I_{\tau_q} = m$, as if the process is right-continuous. We note that this type of inconsistency can be repaired by constructing the inventory process as a process that is multi-valued at the moments of a jump, c.f., Peskir and Shiryaev (2006).

2.2. Structure of optimal policies

Up to now we have been concerned with policies of (m, q) -type, but the problem remains whether an (m, q) -policy can be optimal in the class of stationary policies. The following theorem answers this in the affirmative.

Theorem 1. *There exists an (m, q) -policy that is optimal in the class of stationary policies.*

Proof. We prove that any stationary policy with finite average cost can be reduced to an (m, q) -policy. The immediate consequence is that the optimal stationary policy is of (m, q) type.

Consider first a stationary policy without a clearing level q , in other words, this policy never clears the inventory. This leads to infinite cost since the production rate is larger than the average demand rate, i.e., $1 > \lambda \mathbb{E}Y$, so that the inventory drifts to ever larger levels.

Thus, any policy that has finite average cost must provide for every possible starting level x a clearing level $q(x)$ that is larger than x . Moreover, since the inventory can start at any level x , we require that $\sup q(x) = \infty$.

Suppose that the inventory starts at some level x . As the inventory process drifts upwards with rate $1 - \lambda \mathbb{E}Y < 1$, the skip-free-to-right property of the production process ensures that the inventory will hit level $q(x)$ with expected finite time. (In other words, the continuity of production guarantees that the inventory cannot ‘jump over’ level $q(x)$.) When the inventory hits the clearing level $q(x)$, the inventory is cleared to a reset level $m(x)$. Note that, since the policy is stationary by assumption, $m(x)$ can only depend on $q(x)$ and is therefore fixed. As a consequence, any stationary clearing policy can be characterized by the two functions $q(x)$ and $m(x)$.

Next, observe that when the process starts at $x < q(x)$ the process can reach any level $y < x$. To see this, let $\tilde{t} = q(x) - x$ be the smallest amount of time in which the inventory can move from x to $q(x)$ (recall that the production rate is taken to be 1). Since $\mathbb{P}(N(\tilde{t}) = n) > 0$ for any n , we have for the total demand $X_{\tilde{t}}$ that $P(X_{\tilde{t}} > y - x) > 0$. Thus, level y will be reached or crossed with positive probability.

This implies that $q = \inf q(x)$ must be finite, for if this would not be true, the inventory would drift to ever smaller levels: once below some level x it cannot cross level $q(x)$ anymore, and any level y below x will be hit eventually. As a second consequence, the set $(-\infty, q]$ is the single absorbing set of the inventory process (I_t) .

For the purpose of finding stationary policies that minimize long-run average cost, it therefore suffices to consider only policies that bound the inventory to sets of the type $(-\infty, q]$. Whenever the inventory hits q , the stationarity of the policy implies that

the inventory always moves to level m . All in all, it follows that all finite-cost stationary policies are of (m, q) type. But then also the optimal stationary policy must be of (m, q) type. $\square$

2.3. Models with lost sales

When backlogging is not allowed, demand has to be rejected. We consider two commonly used rejection/acceptance policies, c.f., Perry and Asmussen (1995), Germs and Van Foreest (2013a).

Under the partial rejection policy, when the demand y is larger than the on-hand stock x , the order is partially satisfied, i.e., up to level x , and the inventory is emptied. The remainder of the demand, i.e., $y - x$, is lost at the expense of a cost $l(y - x)$. In view of this, the construction of the inventory process in Eq. (1) has to be updated accordingly. Let $\tau_1, \tau_2, \dots$ be the arrival times of the demands $Y_1, Y_2, \dots$, then for the first cycle, i.e., for $t < \tau_q$,

$$I_t = I_0 + t - \sum_{n=1}^{N_t} \min\{Y_n, I_{\tau_{n-}}\},$$

where $I_{\tau_{n-}} = \lim_{h \uparrow t} I_h$. We define the total cost due to loss during the first cycle as

$$R(t) = \sum_{n=1}^{N_t} l([Y_n - I_{\tau_{n-}}]_+),$$

where $[y]_+ = \max\{y, 0\}$.

Under the complete rejection policy, only when an order of size y can be satisfied completely from on-hand stock x , it is accepted, and otherwise rejected entirely at the expense of a cost $l(y)$. Thus, for the first cycle,

$$I_t = I_0 + t - \sum_{n=1}^{N_t} Y_n 1_{\{Y_n \leq I_{\tau_{n-}}\}},$$

$$R(t) = \sum_{n=1}^{N_t} l(Y_n) 1_{\{Y_n > I_{\tau_{n-}}\}}.$$

We assume for both cases that the lost-sales cost function $l(x)$ is non-negative, $l(0) = 0$, and increasing in x .

For both policies we can write the cost until τ_q in the form

$$V(m, q) = \mathbb{E}_m \left(\int_0^{\tau_q} h(I_t) dt \right) + \mathbb{E}_m R(\tau_q) + K(m, q),$$

where $R(t)$ is the cost function related to the chosen rejection policy. Note that the optimization problem has the same form as (5), except that the above expression for $V(m, q)$ has to be used rather than (4). Further, as the reasoning in the proof of Theorem 1 also holds for the models with lost sales, the optimal stationary policy is also an (m, q) policy for the lost sales models.

Remark 4. In view of Remark 1 we do not consider the ‘Complete Acceptance’ policy that accepts each demand in its entirety, possibly with a backlog, but, once the inventory is below a certain level, rejects the demands.

Note also that under the rejection policies, all orders are rejected when the on-hand inventory level $I_t \leq 0$, so that, as a consequence, it is no longer necessary to require that the production rate r exceeds the average demand rate $\lambda \mathbb{E}Y$.

3. Finding the optimal policy

In this section, we provide a powerful procedure to compute the minimizing policy (m^*, q^*) for generally distributed demand for models with backlogging. The optimization procedure consists out of three basic steps. The first is a method to find the optimal, i.e., minimal, average cost. The second is a method to construct the

optimal policy, i.e., the policy that achieves the minimal cost. Note that, in view of [Theorem 1](#), we can constrain the search for an optimal policy right away to an (m, q) policy. We use this fact from now on. We remark in passing the obvious fact that the existence of an optimal policy is guaranteed once we have constructed it. The third step is to most technical one and consists of the study of the—very loosely speaking—derivative of the function $m \rightarrow V(m, q)$. In [Section 3.4](#) we provide an example, and [Section 3.5](#) discusses some useful computational aspects.

As it turns out to be straightforward to modify the model with backlogging to a model with lost sales, we refer to [Appendix B](#) for the derivations and main results for the lost sales case.

3.1. Finding the minimal average cost

Let us assume that the producer receives a reward $g > 0$ per unit time to compensate for the costs that are incurred by the system. Thus, when the inventory level is x , the inventory cost rate becomes $h(x) - g$ rather than $h(x)$. Analogous to [\(4\)](#), let

$$V(m, q; g) = \mathbb{E}_m \left(\int_0^{\tau_q} (h(I_s) - g) ds \right) + K + c(q - m)$$

be the expected g -revised cost of one regenerative cycle that starts at $I_0 = m \geq 0$ and clears the inventory at the stopping time τ_q . Since $\mathbb{E}_m(\int_0^{\tau_q} g ds) = g \mathbb{E}_m \tau_q$, we can write the above as

$$V(m, q; g) = \mathbb{E}_m \left(\int_0^{\tau_q} h(I_s) ds \right) - g \mathbb{E}_m \tau_q + K + c(q - m). \quad (6)$$

Suppose that we can find an (m, q) policy and a g such that $V(m, q; g) = 0$. Then [\(6\)](#) implies

$$\mathbb{E}_m \left(\int_0^{\tau_q} h(I_s) ds \right) + K + c(q - m) = g \mathbb{E}_m(\tau_q),$$

and thereby,

$$g = \frac{\mathbb{E}_m \left(\int_0^{\tau_q} h(I_s) ds \right) + K + c(q - m)}{\mathbb{E}_m(\tau_q)}. \quad (7)$$

From the renewal-reward theorem, c.f., [Tijms \(2003\)](#), it then follows that g is precisely the long-run average cycle cost of the (m, q) policy.

We now use this observation to solve the optimization problem [\(5\)](#). Start with choosing some arbitrary, but positive, g , and assume it is possible to construct an (m, q) policy that minimizes $V(m, q; g)$. Thus, provided such a policy exist, let (m_g, q_g) be this minimizing policy and define $V(g)$ as the minimal value of [Eq. \(6\)](#) for our chosen g , that is,

$$V(g) = V(m_g, q_g; g) = \inf_{m, q} V(m, q; g). \quad (8)$$

Suppose now that our choice for g is such that the minimal value $V(g) > 0$, in other words, the value under the minimizing policy (m_g, q_g) is $V(g) = V(m_g, q_g; g) > 0$. Then it follows from [Eq. \(6\)](#) that

$$\frac{\mathbb{E}_{m_g} \left(\int_0^{\tau_{q_g}} h(I_s) ds \right) + K + c(q_g - m_g)}{\mathbb{E}_{m_g}(\tau_{q_g})} > g.$$

That is, the average cost of the minimizing policy (m_g, q_g) , which is the left hand side of the above, is larger than g . As any other policy has at least the cost of the minimizing policy, the average cost of any policy cannot be equal to g . Thus, our chosen g is simply too small.

What if g were such that under the minimizing policy $V(g) < 0$? In that case, [Eq. \(6\)](#) becomes

$$\frac{\mathbb{E}_{m_g} \left(\int_0^{\tau_{q_g}} h(I_s) ds \right) + K + c(q_g - m_g)}{\mathbb{E}_{m_g}(\tau_{q_g})} < g.$$

In words this means that there is a policy, in fact, policy (m_g, q_g) , that has smaller average cost than g , because the average cost at the left hand side of the equation is smaller than g . Thus, this choice of g cannot be the minimal average cost: the average cost associated with policy (m_g, q_g) is smaller.

The last step is to combine the above ideas. Provided that we can construct a policy that minimizes $V(m, q; g)$ for any given g , we can simply use bisection on g to ‘home-in’ on the optimal reward g^* and optimal policy (m_{g^*}, q_{g^*}) . Specifically, we use bisection to construct a sequence of rewards $g_1, g_2, \dots$ and associated policies $(m_{g_1}, q_{g_1}), (m_{g_2}, q_{g_2}), \dots$ such that $g_n \rightarrow g^*$ and $(m_{g_n}, q_{g_n}) \rightarrow (m_{g^*}, q_{g^*})$ as $n \rightarrow \infty$, where g^* and (m_{g^*}, q_{g^*}) solve [\(5\)](#). As a result, the limiting value g^* is the minimal long-run average cost and its associated clearing policy (m_{g^*}, q_{g^*}) is the optimal policy.

3.2. Constructing the optimal policy

Clearly, the procedure described above hinges on being able to construct a policy that minimizes $V(m, q; g)$ for any given $g > 0$. We now show how to do this.

Below we show that $V(m, q; g)$ can be written as

$$\begin{aligned} V(m, q; g) &= K + \int_m^q \gamma_g(x) dx + c(q - m) \\ &= K + \int_m^q (\gamma_g(x) + c) dx \end{aligned} \quad (9)$$

where γ_g is a convex function if the inventory cost function h is convex, and $c(q - m)$ is the cost of clearing the inventory. It turns out that γ_g is decreasing in g in the sense that if $g_1 > g_2$ then $\gamma_{g_1}(x) < \gamma_{g_2}(x)$ for all x . Thus, by making g sufficiently large, $\gamma_g(\cdot) + c$ will have two roots, s_g and t_g . The above expression for $V(m, q; g)$ then immediately shows how to minimize $V(m, q; g)$: just choose $m = s_g$ and $q = t_g$ since, by convexity, $\gamma_g(x) + c < 0$ for $x \in (s_g, t_g)$, and $\gamma_g(x) + c \geq 0$ elsewhere. Thus, $V(g)$ as defined in [Eq. \(8\)](#) becomes

$$V(g) = \inf_{m, q} V(m, q; g) = V(s_g, t_g; g).$$

Thus, the optimization problem [\(5\)](#) can be reduced to constructing the function γ_g and studying its properties.

3.3. Constructing γ_g

In this section we show how to compute γ_g , i.e., the integrand in [Eq. \(9\)](#). To this end, consider the following function

$$V(m; g) = \inf_{q \geq m} V(m, q; g). \quad (10)$$

Observe that, if we have this function, it is easy to solve the minimization problem [\(8\)](#), because

$$V(g) = \inf_m V(m; g) = \inf_m \inf_{q \geq m} V(m, q; g).$$

Next, let

$$\gamma_g(x) = \lim_{\delta \downarrow 0} \frac{V(x + \delta, q; g) - V(x, q; g)}{\delta}, \quad (11)$$

be the right-derivative of $V(x, q; g)$ at x , so that $V(m; g)$ can be written as

$$V(m; g) = K + \inf_{q \geq m} \int_m^q (\gamma_g(x) + c) dx. \quad (12)$$

Observe that [\(12\)](#) is in fact an optimal stopping problem. This plays a crucial role in the proof of the next theorem.

The most important properties of γ_g can be summarized as follows.

Theorem 2. Let γ_g be defined in [Eq. \(11\)](#).

Fig. 2. The left panel shows the graph of γ_6 and $V(m; 6)$ for $g = 6$ as a function of the reset inventory level $l_0 = m$. The right panel shows graphs of $V(m; g)$ as a function of the reset level m for various values of g .

(i) γ_g is the unique solution of the integral equation

$$\gamma_g(x) = h(x) - g + \lambda \int_0^\infty \gamma_g(x-y)G(y) dy. \quad (13)$$

(ii) $\lim_{x \rightarrow \pm\infty} \gamma_g(x) = \infty$;

(iii) $\gamma_{g_1}(x) < \gamma_{g_2}(x)$ if $g_1 > g_2$.

(iv) $\min_x \gamma_g(x) \rightarrow -\infty$ if $g \rightarrow \infty$.

(v) If h is convex then $\gamma_g(\cdot)$ is also convex for all g .

Proof. We refer to Appendix A for a proof. $\square$

With these properties the optimization procedure becomes entirely clear. First observe that γ_g , being the solution of an integral equation, does not depend on the policy parameters m or q . Then, once we have obtained γ_g from this integral equation, the computation of $V(m; g)$ in Eq. (12) only depends on the upper limit q of the integral.

Next, suppose that g is so small that $\gamma_g(\cdot) + c \geq 0$, where c is the cost per unit cleared. Then Eq. (12) implies that it is optimal to set $q = m$, i.e., clear the inventory right away, and pay K . Hence, if g is too small, $V(m; g) = K$ for all m . Therefore, assume henceforth that g is such that $\gamma_g(\cdot) + c < 0$ on at least one interval. Observe that properties (iii) and (iv) of Theorem 2 imply that this is indeed possible.

By (v), if h is convex, then $\gamma_g + c$ is convex. As a consequence, if g is sufficiently large, $\gamma_g(\cdot) + c$ has exactly two roots s_g and t_g with $s_g < t_g$. Clearly, $\gamma_g(x) + c > 0$ when $x > t_g$. Therefore, when the inventory process starts at $m \geq t_g$ it follows from (12) that it is best to clear immediately; hence $V(m; g) = K$ for all $m \geq t_g$. When $m < t_g$, we can obtain $V(m; g)$ by integrating (12). By the convexity of $\gamma_g(\cdot) + c$, $V(m; g)$ is minimized by taking $m_g = s_g$ as the lower limit of the integral. Hence, when $h(\cdot)$ is convex, the minimal cost must be

$$V(g) = \inf_m V(m; g) = K + \int_{s_g}^{t_g} (\gamma_g(x) + c) dx, \quad (14)$$

and the clearing policy with $m_g = s_g$ and $q_g = t_g$ is the g -revised minimizing policy.

Remark 5. It might be that h is not a convex function. To handle these cases, it is necessary to assume that $\gamma_g(\cdot) + c$ has a finite number of roots, so that the roots of $\gamma_g(\cdot) + c$ can be represented by the set of root pairs $\{(s_g^1, t_g^1), \dots, (s_g^n, t_g^n)\}$, such that $s_g^1 < t_g^1 <$

$\dots < s_g^n < t_g^n$ and $\gamma_g(s_g^i-) + c \geq 0 \geq \gamma_g(s_g^i+) + c$ and $\gamma_g(t_g^i-) + c \leq 0 \leq \gamma_g(t_g^i+) + c$. As a result, $V(g)$ is given by

$$V(g) = K + \min_{s_g^i < t_g^i} \left\{ \int_{s_g^i}^{t_g^i} (\gamma_g(z) + c) dz \right\}. \quad (15)$$

Now the minimization comes down to choosing the best root pairs such that above integral becomes as small as possible.

3.4. An example

To illustrate the procedure we plot the functions $\gamma_6(m)$ and $V(m; 6)$ for $g = 6$ in Fig. 2. The parameters are such that the production rate $r = 1$, the fixed clearing cost $K = 10$, there are no variable clearing cost (i.e., $c = 0$), and the inventory cost

$$h(x) = -4 \min\{0, x\} + \max\{0, x\}.$$

The arrival rate of the demand $\lambda = 0.8$, and the demands are distributed on $[0, 2]$. Thus, $\lambda \mathbb{E}Y = 0.8 \cdot 1 < 1 = r$.

The figure shows that γ_6 preserves the convexity of the inventory cost function h . It also confirms that $V(m; 6)$ obtains its minimum at the left root s_6 of γ_6 and that the optimal clearing level q is attained at the right root t_6 of γ_6 . Since $V(s_6; 6) < 0$, it follows from (9) that $g = 6$ is not the minimal long-run average cost for this system.

Using bisection we compute the g -revision cost that solves $V(g^*) = 0$. The right panel of Fig. 2 shows the graphs of $V(m; g)$ for $g = 5.5$, $g = 5.75$ and $g = 6$. Since $g = 6$ was too large, we try $g = 5.5$ as a next guess. As we see, $V(5.5) > 0$, thus $g = 5.5$ is not large enough to cover the average cycle cost. By bi-section $g = 5.75$ should be the next guess. This leads to the (very close to) optimal solution, i.e. $V(5.75) = 0$ and thereby $g^* = 5.75$, $m_g^* = 2.66$ and $q_g^* = 6.94$.

3.5. Computational aspects

For the numerical integration of (13), we approximate γ_0 over the continuous domain $\mathbb{R}$ by a discrete set of function values at a discrete set of points in the domain. That is, we discretize the state space by reducing the real numbers to the grid $\{\dots, k\delta, (k+1)\delta, \dots\}$ for $k \in \mathbb{Z}$, where $\delta > 0$ denotes the grid size. Writing

$\gamma_k = \gamma_0(k\delta)$, $h_k = h(k\delta)$, and $G_k = G(k\delta)$, Eq. (13) reduces to

$$\gamma_k = h_k + \lambda\delta \sum_{i=0}^{\infty} \gamma_{k-i} G_i. \quad (16)$$

Taking the term $\gamma_k G_0$ out of the summation and to the left, we obtain the recursion

$$\gamma_k = \frac{h_k}{1 - \lambda\delta G_0} + \frac{\lambda\delta}{1 - \lambda\delta G_0} \sum_{i=1}^{\infty} \gamma_{k-i} G_i.$$

Thus, once we have all γ_j , $j < k$ we can compute γ_k and so on.

This, however, is a problem: Suppose we want to start the recursion at some index k , then the values γ_j for $j < k$ are unknown. To resolve this, consider some $N \ll 0$ and take $\gamma_{k,N} = h_k/(1 - \lambda\delta G_0)$ for $k \leq N$. Then, for $k > N$ the above recursion becomes

$$\gamma_{k,N} = \frac{h_k}{1 - \lambda\delta G_0} + \frac{\lambda\delta}{1 - \lambda\delta G_0} \sum_{i=1}^{k-N-1} \gamma_{k-i,N} G_i + \frac{\lambda\delta}{1 - \lambda\delta G_0} \sum_{i=k-N}^{\infty} h_{k-i} G_i. \quad (17)$$

It can be shown, c.f. Van Foreest and Wijngaard (2014, Appendix A), that $\gamma_{k,N} \rightarrow \gamma_k$ exponentially fast as $N \rightarrow -\infty$.

The numerical procedure involves to bisect on g , so that we need to compute $\gamma_g(\cdot)$ for various values of $g > 0$. This seems to require to solve the integral Eq. (13) for each g separately, but it turns out that it can be avoided by the following idea. Let us assume that $\gamma_g(\cdot)$ can be written in the form

$$\gamma_g(x) = \gamma_0(x) - g\beta(x), \quad (18)$$

where γ_0 is the solution of (13) with $g = 0$. Substituting this expression in (13) leads to the following condition on $\beta(x)$:

$$\beta(x) = 1 - \lambda \int_0^{\infty} \beta(x-y) G(y) dy.$$

Next, by filling in the guess $\beta(x) \equiv \beta$, i.e., $\beta(x)$ is constant, we see that $\beta = 1/(1 - \lambda\mathbb{E}Y)$. Thus,

$$\gamma_g(x) = \gamma_0(x) - \frac{g}{1 - \lambda\mathbb{E}Y}, \quad (19)$$

is a solution of (13), and by the uniqueness of Theorem 2.i, it is the only solution. As a result, we only need to compute γ_0 , and then γ_g follows trivially. Similar arguments can be applied to the models with lost sales. We refer to Appendix B for the results.

Finally, to determine the minimal long run average cost g^* , we use bisection. As initial lower bound g_- on g^* we take $g_- = 0$, and as initial upper bound g_+ we take the average cost of an arbitrary clearance policy, for instance the average cost related to the (m, q) -policy with $m = 0$ and $q = 1$. The optimal policy must outperform this arbitrary policy, so $g^* \leq g_+$.

4. Level crossing analysis

With the method developed above it is possible to prove structure results for the optimal policy and numerically analyze production-clearing systems with general demand sizes. For the special case of exponentially distributed demand sizes and a given (m, q) clearing policy, we can use level crossing to obtain closed-form expressions for the long-run average cost in terms of the steady-state density function $\pi(\cdot)$ of the inventory process. In this section we derive this cost function. Besides the intrinsic value of having closed-form results, we can use the cost formula as a test case to validate the software implementation of the numerical method of the previous sections.

The result below is a straightforward generalization of the result of Berman et al. (2005) in that we allow m to take any value rather than fixing it to $m = 0$. We therefore refer to this work

for technical details regarding the existence of the density function and other related points; here we are content with a heuristic derivation based on level crossing arguments.

Let $\pi_\delta(x) = \mathbb{P}(I \in [x - \delta, x))$ where I is the inventory process in steady-state. Consider some level $x \in (m, q)$. Then level x is up-crossed with rate $\delta\pi_\delta(x) + o(\delta)$, recall that the production facility works at rate δ . The down-crossing rate is $\delta\pi_\delta(q) + \delta \int_x^q \pi_\delta(y) G(y-x) dy + o(\delta)$, due to the clearing action at level q , so that the inventory moves to level m (hence crosses level x), plus all transitions due to demand of sufficiently large size. Equating the up- and down-crossing rates, dividing by δ and taking $\delta \rightarrow 0$, we arrive at the fact that the density $\pi(\cdot)$ of $(I_t)_{t \geq 0}$ should satisfy a Pollaczek-Khintchine equation of type:

$$\pi(x) = \lambda \int_x^q G(w-x) \pi(w) dw + \pi(q), \quad m \leq x \leq q \quad (20a)$$

$$\pi(x) = \lambda \int_x^q G(w-x) \pi(w) dw, \quad x < m. \quad (20b)$$

While it is not possible (to the best of our knowledge) to obtain a closed-form solution of (20) for general G , we can find an expression for the case of exponentially distributed demand.

Lemma 3. *If demand is exponentially distributed with mean $\mu^{-1} = \mathbb{E}Y$, the steady-state density function π of the inventory process under an (m, q) clearing policy is given by*

$$\pi(x) = \begin{cases} Ae^{-(\lambda-\mu)x} + B, & m \leq x \leq q, \\ Ce^{-(\lambda-\mu)x}, & x < m, \end{cases}$$

with

$$A = -\frac{1}{q-m} \frac{\lambda}{\mu} e^{(\lambda-\mu)q}, \quad B = \frac{1}{q-m}, \quad C = \frac{1}{q-m} (1 - e^{-(\lambda-\mu)(q-m)}).$$

Proof. We refer to Appendix C for a proof. $\square$

As $\pi(x)$ is the steady-state density at the inventory level $x \leq q$, it follows from level crossing theory that $\pi(q)$ is the clearing rate. Hence, $1/\pi(q)$ is the expected cycle length. Moreover, as $h(x)$ is the inventory cost rate when the inventory level is x , so that inventory costs accrue at rate $h(x)\pi(x)$, and as $\pi(q)$ is the clearing rate, so that the cost rate associated with clearing is $(K + c(q-m))\pi(q)$, the average cost per unit time of the (m, q) policy equals

$$C(m, q) = \int_{-\infty}^q h(x) \pi(x) dx + [K + c(q-m)] \pi(q). \quad (21)$$

Remark 6. By similar methods it is possible to include the effect of the lost-sales policies as discussed in Section 2.3, c.f., Berman et al. (2005). We remark in passing that for the case with complete rejection the second equation on page 214 in (Berman et al., 2005) contains a minor error. Instead of

$$\mathbb{E}(\text{shortage cost}) = \frac{\pi\lambda}{f_2(0)} \left[\frac{1}{\mu} + I_2(q) \right] \int_0^q [1 - G(x)] f_2(x) dx,$$

the correct expression is

$$\mathbb{E}(\text{shortage cost}) = \frac{\pi\lambda}{f_2(0)} \int_0^q \left[\frac{1}{\mu} + x \right] [1 - G(x)] f_2(x) dx, \quad (22)$$

where, in the notation of Berman et al. (2005), $f_2(\cdot)$ is the steady-state density of the inventory process and $I_2(q)$ the expected inventory level.

Remark 7. It is important to realize that finding the optimal m^* and q^* by means of (21) requires a search in two dimensions, i.e., along m and q . If $C(q, m)$ is not convex, finding the minimum is typically a difficult problem. With our method developed in Section 3, the search procedure is retrained to a one-D search, i.e., in g .

Table 1Optimal clearing policies and cost for the backlog model with non-negative reset level $m \geq 0$.

c_b	K	CV	$\rho = 0.5$						$\rho = 0.9$					
			$\lambda = 5, \mathbb{E}Y = 0.1$			$\lambda = 1, \mathbb{E}Y = 0.5$			$\lambda = 9, \mathbb{E}Y = 0.1$			$\lambda = 1, \mathbb{E}Y = 0.9$		
			m^*	q^*	g^*	m^*	q^*	g^*	m^*	q^*	g^*	m^*	q^*	g^*
2	4	0.50	0.00	2.01	1.95	0.00	2.23	1.93	0.03	1.50	1.08	4.17	7.15	6.23
2	4	1.00	0.00	2.03	1.93	0.00	2.48	2.10	0.26	1.97	1.44	7.22	10.73	9.98
2	4	2.00	0.00	2.18	1.94	0.00	3.29	3.15	1.44	3.75	3.11	18.71	23.49	23.80
2	40	0.50	0.00	6.33	6.27	0.00	6.41	6.09	0.00	3.16	2.61	2.67	9.14	6.96
2	40	1.00	0.00	6.34	6.23	0.00	6.56	6.06	0.00	3.56	2.74	5.40	12.98	10.60
2	40	2.00	0.00	6.38	6.14	0.00	7.43	6.44	0.38	5.44	4.05	16.09	26.41	24.26
4	4	0.50	0.00	2.02	1.95	0.00	2.37	2.08	0.34	1.81	1.40	6.98	9.96	9.04
4	4	1.00	0.00	2.05	1.95	0.00	2.80	2.45	0.77	2.48	1.95	11.81	15.32	14.56
4	4	2.00	0.00	2.29	2.06	0.44	4.48	4.45	2.75	5.06	4.44	30.11	34.90	35.18
4	40	0.50	0.00	6.33	6.27	0.00	6.46	6.15	0.00	3.37	2.82	5.48	11.95	9.77
4	40	1.00	0.00	6.34	6.24	0.00	6.70	6.21	0.12	4.02	3.20	9.98	17.57	15.18
4	40	2.00	0.00	6.43	6.19	0.00	8.18	7.25	1.70	6.75	5.38	27.50	37.82	35.64

5. Numerical analysis

Here we demonstrate the power of our numerical procedure by applying it first to the analysis of models with backordering, and then to models with lost sales.¹ In the sequel we impose the condition that the reset level $m \geq 0$, c.f., Remark 1.

5.1. Models with backlogging

To illustrate our approach, we compute the optimal clearing policy (m^*, q^*) and cost g^* for a clearing system with (convex) inventory cost function

$$h(x) = c_b x^- + c_h x^+,$$

where $x^- = \max\{-x, 0\}$ and $x^+ = \max\{x, 0\}$. The common demand size Y is gamma-distributed with coefficient of variation $CV = \sigma/\mathbb{E}Y$, where σ denotes the standard deviation. Specifically, the distribution function of Y is given by

$$F(y) = \int_0^y \frac{x^{k-1}}{\theta^k \Gamma(k)} e^{-x/\theta} dx,$$

where $\theta = \mathbb{E}Y/k$, $k = 1/CV^2$, and $\Gamma(k)$ is the gamma function evaluated at k .

Our initial set of experiments consists of 48 test cases, c.f., Table 1, which are based on combinations of the backorder cost c_b , fixed clearing cost K , arrival rate λ , mean demand size $\mathbb{E}Y$, and CV. The unit holding cost c_h is set to 1 for all test cases. It is important to remark that when the demand size is exponentially distributed the CV is 1, so that we can use Eq. (21) to validate our results. As expected, the results of our method correspond with those of the closed-form expression.

We remark in passing that when the demand process is fairly constant, e.g., when λ is large and $\mathbb{E}Y$ small, the sample paths appear to be similar to the sample paths of the Economic Order Quantity (EOQ) model run 'backwards' in time, c.f., Fig. 1. Therefore, the EOQ model might provide reasonable estimates for the optimal policy parameters q and m . We will verify this assumption below.

The results in Table 1 show some interesting aspects of the model. In case of low demand, i.e., when $\rho = \lambda \mathbb{E}Y = 0.5$, the CV hardly affects the optimal m^* , q^* , and g^* . This means that both the EOQ model with backlogging and the production-clearing model with exponential demand can be used as a good approximation

for production-clearing models with general demand distributions when the load is low.

However, when $\rho = 0.9$ and the demand size is large, i.e., $\mathbb{E}Y = 0.9$, the influence of the CV is dramatic: g^* varies for instance from 9.77 to 35.64. This implies that the analysis of the case with exponential demand, i.e., when $CV = 1$, cannot be used to obtain valuable insights: the estimates for m^* , q^* and g^* with $CV = 1$ are more than a factor 2 off from the values for the case with $CV = 2$.

Second, when $\rho = 0.5$, the optimal reset level $m^* = 0$ for almost all instances. This suggests that the optimal cost can be lowered if we would allow m to take negative values. From a practical point of view, however, it appears to be rather odd to backlog the demand of the wholesaler, i.e., the customer that clears the inventory when the inventory level grows to undesirable levels. Note that the $m \geq 0$ constraint has the most effect on instances with a low load which explains the counter intuitive result that for some instances the optimal cost is decreased when the load is increased. For example, for the case $\rho = 0.9$, $\lambda = 9$, $\mathbb{E}Y = 0.1$, $CV = 2$, $c_b = 4$ and $K = 40$, the average optimal cost g^* is actually lower than for the similar cases with $\rho = 0.5$.

Our next objective is a sensitivity analysis of the influence of the policy parameters m and q on the average cost, c.f., Fig. 3. In the left panel we consider a model with $\lambda = 1$, $c_b = 4$, $K = 40$, and gamma distributed demand with $\mathbb{E}Y = 0.9$, and $CV = 1/2$, and compute the optimal policy parameters q^* , m^* , and the minimal average cost g^* . Then we compare the average cost for three different policies (m, q^*) , $(m, q^*/2)$ and $(m, 2q^*)$ as a function of m . Specifically, we compute $g(m, q^*)/g^*$ for policy (m, q^*) (and likewise for $(m, q^*/2)$ and $(m, 2q^*)$) and vary m from 0 to $0.95q$. The approach for the right panel is identical, except that $CV = 2$ rather than $1/2$. The results show clearly that the cost is quite sensitive to the value of q , i.e. a factor 2 deviation in the optimal q results in a cost increase of at least 20 percent. This is unlike the insensitivity of the EOQ model, in which ordering deviation of a factor 2 leads to a cost increase of just 5 percent compared to the minimal cost. From the results we conclude that it is also necessary to tune m and q in concert. Finally, fixing the reset level m to zero is quite costly: the cost increases by some 20 percent as compared to the cost with optimal m^* .

5.2. Results for lost-sales models

We now analyze the influence of lost sales policies on the parameters of the optimal policy and the associated cost. For our numerical design we use the same clearing system and combinations of model parameters as in Section 5.1 with the exception of back-order cost which we replace by lost sales costs as explained in

¹ The code is written in Python and can be obtained upon request at one of the authors.

Fig. 3. The graph labeled as $q^*/2$ shows the average cost of policy $(m, q^*/2)$, as a multiple of g^* , as a function of m . Similar graphs show the average cost of policy (m, q^*) and $(m, 2q^*)$. Note that for only when $m = m^*$ the average cost of policy (m, q^*) is equal to g^* , hence equals 1 when $m = m^*$. At the left panel, $CV = 0.5$, and at the right, $CV = 2$.

Table 2

Optimal clearing policies and cost for the lost-sales model with the partial acceptance and complete rejection issuing policy with unit holding cost $c_h = 1$.

c_l	K	CV	$\lambda = 9, \mathbb{E}Y = 0.1$						$\lambda = 1, \mathbb{E}Y = 0.9$					
			Partial acceptance			Complete rejection			Partial acceptance			Complete rejection		
			m^*	q^*	g^*	m^*	q^*	g^*	m^*	q^*	g^*	m^*	q^*	g^*
2	4	0.50	0.05	1.53	1.11	0.12	1.60	1.18	0.00	3.02	2.12	0.00	3.51	2.67
2	4	1.00	0.05	1.76	1.24	0.12	1.89	1.37	0.00	3.20	2.44	0.00	3.63	3.19
2	4	2.00	0.02	2.29	1.60	0.00	2.57	1.93	0.00	3.33	3.09	0.00	3.30	3.92
2	40	0.50	0.00	3.48	2.93	0.00	3.55	3.00	0.00	6.26	4.09	0.00	6.78	4.65
2	40	1.00	0.00	3.85	3.03	0.00	3.98	3.16	0.00	7.00	4.61	0.00	7.75	5.52
2	40	2.00	0.00	4.88	3.43	0.00	5.20	3.77	0.00	8.28	5.85	0.00	9.00	7.41
20	4	0.50	0.53	2.01	1.59	0.60	2.08	1.66	2.39	5.38	4.46	3.03	6.02	5.10
20	4	1.00	0.76	2.47	1.95	0.90	2.61	2.09	2.90	6.41	5.67	4.05	7.61	6.89
20	4	2.00	1.38	3.69	3.01	1.79	4.11	3.43	3.66	8.39	8.74	5.38	10.66	11.69
20	40	0.50	0.26	3.75	3.20	0.33	3.82	3.27	1.46	7.93	5.76	2.10	8.57	6.40
20	40	1.00	0.38	4.27	3.45	0.52	4.42	3.60	1.74	9.33	6.96	2.79	10.54	8.19
20	40	2.00	0.75	5.81	4.35	1.11	6.23	4.77	2.03	12.22	10.12	3.25	14.65	13.08

Section 2.3. In particular, we assume that the lost sales costs are linear in the amount lost, i.e.,

$$l(x) = c_l x, \quad (23)$$

where $c_l > 0$ is the cost per unit lost.

To validate our implementation we compared our results to the example of Berman et al. (2005) for the policies with partial acceptance and complete rejection. For both policies we obtain the same results, except for the fact that we use (22) rather than their expression to compute the expected cost.

Table 2 shows the results for the analysis of the case with $\rho = 0.9$ as in Section 5.1 and with $c_l = 2$ or $c_l = 20$ and the cost function as defined in (23). It is apparent that, as expected, the average cost under partial acceptance is lower than cost under complete rejection. When comparing the cases with $\lambda = 9, \mathbb{E}Y = 0.1$ to the cases with $\lambda = 1, \mathbb{E}Y = 0.9$ we see that the costs increase (significantly) as a function of the coefficient of variation of the demand. Finally, when c_l is large, i.e., 20 rather than 2, the optimal reset level m is positive.

Finally, comparing the results of Table 2, i.e., the average cost with lost sales, to those of Table 1, i.e., the average cost with backlogging, we see a remarkable difference. With lost sales, high load, i.e., $\rho = 0.9$, and large order sizes, i.e., $\mathbb{E}Y = 0.9$, the average costs are often much lower, even in case the rejection costs c_l are large.

As yet, we do not have a simple or intuitive explanation for this difference. An interesting direction for further research is to identify performance measures such that inventory systems with and without rejection can be compared in a consistent way.

6. Conclusions and topics for further research

In this paper, we studied continuous-review clearing policies for production-clearing systems with compound Poisson demand and piece-wise continuous cost functions. We show that the average-cost optimal clearing policy is of (m, q) -type. Such policies restore the inventory in the system to a level $m \leq q$ when the inventory level hits or exceeds a critical level q . We derive a new method to efficiently compute the parameters optimal (m, q) policy for models with backlogging and lost sales. For the special case with backlogging and exponentially distributed demand we formulate closed-form expressions for the average-cost of a given (m, q) policy using level crossing arguments. Our numerical experiments show that optimizing the clearing levels m and q can lead to substantial cost savings, in the order of 20 percent or more. Moreover, when the utilization of the system is high, the coefficient of variation of the individual demand sizes has a large effect on the optimal policy parameters and cost. Therefore, deterministic approximations for stochastic production-clearing models such as the EOQ model perform badly in cases with high utilization. A remarkable result is that for cases with lost sales, the inventory levels as well as the average costs are much lower than for cases with backlogging, even when the lost sales cost are (very) large. We do note however that it is difficult to compare the results of models with backlogging and models with lost sales because of the difference in cost structure.

In view of this last remark, an interesting direction for further research is to identify performance measures such that production-

clearing systems with and without lost sales can be compared in a consistent way. Another direction is to analyze a system in which the production capacity can be controlled at the expense of a fixed or variable cost. Then the problem is to balance the cost of producing at a certain rate against the inventory and clearing cost. Finally, an interesting modification of the model is to include constraints on the minimum and maximum difference between m and q , as such constraints could incorporate size constraints imposed by the trucks that clear the inventory, or, for instance, the size of one or more silo's that may need to be emptied completely during the clearing action.

Appendix A. Proof of Theorem 2

Proof.

- (i) In the proof, the infinitesimal operator of the inventory process plays an important role for deriving a useful expression for $V(m; g)$ in terms of the negative derivative of $V(m, q; g)$. Let $\mathbb{L}_I$ denote the infinitesimal operator of I_t . Then from the standard theory of continuous-time Markov processes it can be shown that for functions $f \in C^1$

$$\begin{aligned}\mathbb{L}_I f(x) &= \lim_{t \downarrow 0} \frac{\mathbb{E}_{(I_0=x)}(f(I_t) - f(x))}{t} \\ &= \lim_{t \downarrow 0} \frac{(1 - \lambda t)f(x+t) + \lambda t \mathbb{E}(f(x-Y) - f(x))}{t} \\ &= f'(x) + \lambda[\mathbb{E}(f(x-Y)) - f(x)]\end{aligned}$$

where $f'(x) = df(x)/dx$. Applying the infinitesimal operator $\mathbb{L}_I$ to $V(\cdot, q; g)$, as defined in (6), gives

$$\mathbb{L}_I V(x, q; g) = V'(x, q; g) + \lambda[\mathbb{E}(V(x-Y, q; g)) - V(x, q; g)]. \quad (\text{A.1})$$

We obtain a second expression for $\mathbb{L}_I V(x, q; g)$ by applying Dynkin's formula,

$$\mathbb{E}_x(f(I_\tau)) = f(x) + \mathbb{E}_x\left(\int_0^\tau \mathbb{L}_I f(I_s) ds\right),$$

to $V(I_t, q; g)$ at the stopping time τ_q . That is,

$$\begin{aligned}V(x, q; g) &= \mathbb{E}_x(V(I_{\tau_q}, q; g)) - \mathbb{E}_x\left(\int_0^{\tau_q} \mathbb{L}_I V(I_s, q; g) ds\right) \\ &= K + c(q-m) - \mathbb{E}_x\left(\int_0^{\tau_q} \mathbb{L}_I V(I_s, q; g) ds\right).\end{aligned}$$

Combining this with (6), we see that for all q

$$-\mathbb{E}_x\left(\int_0^{\tau_q} \mathbb{L}_I V(I_s, q; g) ds\right) = \mathbb{E}_x\left(\int_0^{\tau_q} (h(I_t) - g) dt\right),$$

from which follows that

$$\mathbb{L}_I V(m, q; g) = -h(x) + g. \quad (\text{A.2})$$

Hence, after combining (A.1) and (A.2) it follows that

$$-V'(x, q; g) = h(x) - g + \lambda[\mathbb{E}(V(x-Y, q; g)) - V(x, q; g)]. \quad (\text{A.3})$$

Defining $\gamma_g(x) = -V'(x, q; g)$ as the negative derivative of $V(x, q; g)$, it follows from (A.3) that γ_g is the solution of the equation

$$\gamma_g(x) = h(x) - g + \lambda \mathbb{E}(V(x-Y, q; g) - V(x, q; g)).$$

This can be rewritten as the integral equation

$$\gamma_g(x) = h(m) - g + \lambda \mathbb{E}\left(\int_{x-Y}^x \gamma_g(z) dz\right)$$

$$\begin{aligned}&= h(x) - g + \lambda \int_0^\infty \int_{x-y}^x \gamma_g(z) dz dF(y) \\ &= h(x) - g + \lambda \int_0^\infty \gamma_g(x-y) G(y) dy.\end{aligned}$$

Observe that γ_g is a function of only the starting inventory level x and hence independent of the clearing threshold q . As a consequence, we can reformulate $V(x, q; g)$ and $V(m; g)$ in terms of γ_g as follows

$$V(x, q; g) = \int_x^q \gamma_g(z) dz + K + c(q-x) = \int_x^q (\gamma_g(z) + c) dz + K$$

and

$$V(m; g) = K + \inf_{q \geq m} \int_m^q (\gamma_g(x) + c) dx.$$

Finally, the uniqueness follows the fact the right hand side of (13) can be seen to be a functional P such that

$$(Pf)(\cdot) = h(\cdot) - g + \int_0^\infty f(\cdot - y) G(y) dy. \quad (\text{A.4})$$

It can be shown that P is a contraction on a Banach space with a suitable norm, c.f., Germs and Van Foreest (2013b) and Van Foreest and Wijngaard (2014). This proves the existence and uniqueness of a solution of (13), i.e., there is a γ_g such that $P\gamma_g = \gamma_g$. Moreover, as P is a contraction, $\gamma = \lim_{n \rightarrow \infty} P^n h$, where $P^n h = PP^{n-1} h$ for $n > 1$, and $P^1 h = Ph$.

- (ii) First consider (13) with $g = 0$. From Eq. (A.4) lem:3 and $h \geq 0$ it follows that $Ph(\cdot) = h(\cdot) + \int_0^\infty h(\cdot - y) G(y) dy \geq 0$. Thus, $P^n h \geq 0$, and consequently $\gamma_0 \geq 0$. From (13) it follows then that $\gamma_0 \geq h$. Finally, as $h \rightarrow \pm\infty$, and $h \geq 0$, the claim follows for γ_0 . From Eq. (19) the result also holds for $g > 0$.
- (iii) Follows from Eq. (19).
- (iv) Follows straightaway away from Eq. (19) and the fact that $\gamma_0 \geq 0$.
- (v) It is simple to see that the contraction P , defined in Eq. (A.4), maintains convexity, that is, if h is convex, then Ph is also convex. But then $P^n h$ must also be convex. Hence, γ_g is convex. $\square$

Appendix B. Lost sales cases

We first consider the inventory process under the partial acceptance regime. Note that the partial acceptance policy acts as a reflection barrier on the inventory process. As a result, the infinitesimal operator $\mathbb{L}_I$ of the inventory process changes to

$$\begin{aligned}\mathbb{L}_I f(x) &= f'(x) + \lambda \mathbb{E}((f(x-Y) - f(x))1_{Y < x}) \\ &\quad + \lambda \mathbb{E}((f(x-Y) - f(x))1_{Y \geq x}) \\ &= f'(x) + \lambda \mathbb{E}((f(x-Y) - f(x))1_{Y < x}) \\ &\quad + \lambda \mathbb{P}(Y > x)(f(0) - f(x)).\end{aligned} \quad (\text{B.1})$$

Similar derivations as in Section Appendix A yield that

$$\mathbb{L}_I V(x, q; g) = -h(x) - L(x) + g. \quad (\text{B.2})$$

where $L(x) = \lambda \mathbb{E}[I(Y-x)1_{Y > x}]$ is cost rate due to loss of demand. From (B.1) and (B.2) it then follows that $\gamma_g(x) = -V'(x, q; g)$ is the solution of the equation

$$\begin{aligned}\gamma_g(x) &= h(x) + L(x) - g + \lambda \mathbb{E}((V(x-Y, q; g) - V(x, q; g))1_{Y < x}) \\ &\quad + \lambda \mathbb{P}(Y > x)(V(0, q; g) - V(x, q; g)).\end{aligned}$$

We can rewrite this to an integral equation that is (numerically) easier to solve

$$\gamma_g(x) = h(x) + L(x) - g + \lambda \int_0^x \gamma_g(x-y) G(y) dy,$$

by observing that

$$\begin{aligned} & \lambda \mathbb{E}((V(x - Y, q; g) - V(x, q; g))1_{Y < x}) \\ &= \lambda \int_0^x \gamma_g(z) \int_{x-z}^x dF(y) dz \\ &= \lambda \int_0^x \gamma_g(z) (G(x - z) - G(x)) dz \\ &= \lambda \int_0^x \gamma_g(x - y) (G(y) - G(x)) dy \end{aligned}$$

and

$$\lambda \mathbb{P}(Y > x) (V(0, q; g) - V(x, q; g)) = \lambda G(x) \int_0^x \gamma_g(y) dy.$$

To speed up the computations, we again use (18). This results in that β should solve

$$\beta(x) = 1 - \lambda \int_0^x \beta(x - y) G(y) dy.$$

Most of the discussion of the previous paragraph can be carried over to the other two lost sales models; we mention only the results. For complete rejection we obtain

$$\begin{aligned} L(x) &= \lambda \mathbb{E}(I(Y)1_{Y > x}), \\ \mathbb{L}_I f(x) &= f'(x) + \lambda \mathbb{E}((f(x - Y) - f(x))1_{Y < x}), \\ \gamma_g(x) &= h(x) + L(x) - g + \lambda \mathbb{E}[(V(x - Y, q; g) - V(x, q; g))1_{Y < x}] \\ &= h(x) + L(x) - g + \lambda \int_0^x \gamma_g(x - y) (G(y) - G(x)) dy, \\ \beta(x) &= 1 - \lambda \int_0^x \beta(x - y) (G(y) - G(x)) dy. \end{aligned}$$

For complete acceptance, similar reasoning yields

$$\begin{aligned} L(x) &= \lambda 1_{x \leq 0} \mathbb{E}(I(Y)) \\ \mathbb{L}_I f(x) &= f'(x) + \lambda 1_{x > 0} \mathbb{E}((f(x - Y) - f(x))), \\ \gamma_g(x) &= h(x) + L(x) - g + \lambda 1_{x > 0} \mathbb{E}((V(x - Y, q; g) - V(x, q; g))) \\ &= h(x) + L(x) - g + \lambda 1_{x > 0} \int_0^\infty \gamma_g(x - y) G(y) dy, \\ \beta(x) &= (1 - \lambda 1_{x > 0} \mathbb{E}(Y))^{-1}. \end{aligned}$$

Appendix C. Proof of Lemma 3

By taking $G(x) = e^{-\mu x}$ Eq. (20) simplifies to

$$\pi(x) = \lambda \int_x^q e^{-\mu(w-x)} \pi(w) dw + \pi(q), \quad m \leq x \leq q \quad (\text{C.1a})$$

$$\pi(x) = \lambda \int_x^q e^{-\mu(w-x)} \pi(w) dw, \quad x < m. \quad (\text{C.1b})$$

To solve this, we differentiate the second equation with respect to x ,

$$\pi'(x) = -(\lambda - \mu)\pi(x), \quad x < m.$$

This and the argumentation of Berman et al. (2005) imply that π must be of the form

$$\pi(x) = \begin{cases} Ae^{-(\lambda-\mu)x} + B, & m \leq x \leq q, \\ Ce^{-(\lambda-\mu)x}, & x < m. \end{cases} \quad (\text{C.2})$$

From (C.1a) and (C.1b) it follows that $\pi(m+) = \pi(m-) + \pi(q)$, where $\pi(m\pm) = \lim_{h \downarrow 0} \pi(m \pm h)$. Therefore

$$Ae^{-(\lambda-\mu)m} + B = Ce^{-(\lambda-\mu)m} + Ae^{-(\lambda-\mu)q} + B.$$

This results in

$$C = B(1 - e^{-(\lambda-\mu)(q-m)}).$$

To determine A and B we note that by (C.2) we have $\pi(m+) = Ae^{-(\lambda-\mu)m} + B$ and $\pi(q) = Ae^{-(\lambda-\mu)q} + B$. Using this and Eq. (C.1a) to compute $\pi(m+)$ we obtain

$$B = -A \frac{\mu}{\lambda} e^{-(\lambda-\mu)q}.$$

Using the normalization condition

$$\int_{-\infty}^q \pi(x) dx = 1,$$

we find that

$$A = -\frac{1}{q-m} \frac{\lambda}{\mu} e^{(\lambda-\mu)q}, \quad B = \frac{1}{q-m}.$$

References

- Barron, Y. (2015). Clearing control policies for MAP inventory process with lost sales. *European Journal of Operational Research*, 251, 495–508. doi:10.1016/j.ejor.2015.12.028.
- Berman, O., Parlar, M., Perry, D., & Posner, M. J. M. (2005). Production/clearing models under continuous and sporadic reviews. *Methodology and Computing in Applied Probability*, 7, 203–224.
- Germs, R., & Van Foreest, N. D. (2013a). Analysis of finite-buffer state-dependent bulk queues. *OR Spectrum*, 35(3), 563–583.
- Germs, R., & Van Foreest, N. D. (2013b). *Optimal control policies for production-inventory systems with constant replenishment rates*. SOM Research Report, 14001-OPERA, University of Groningen.
- Kim, K., & Seila, A. F. (1993). A generalized cost model for stochastic clearing systems. *Computers & Operations Research*, 20(1), 67–82.
- Larsen, G. K. H., Van Foreest, N. D., & Scherpen, J. M. A. (2014). Distributed MPC applied to a network of households with μ -CHP and heat storage. *IEEE Transactions on Smart Grid*, 5(4), 2106–2114.
- Liu, B., & Cao, J. (1999). Analysis of a production-inventory system with machine breakdowns and shutdowns. *Computers & Operations Research*, 26, 73–91.
- Perry, D., & Asmussen, S. (1995). Rejection rules in the M/G/1 queue. *Queueing Systems*, 19, 105–130.
- Perry, D., & Posner, M. J. M. (2002). Production-inventory models with an unreliable facility operating in a two-stage random environment. *Probability in the Engineering and Informational Sciences*, 16(3), 325–338.
- Perry, D., Stadje, W., & Zacks, S. (2005). Sporadic and continuous clearing policies for a production/inventory system under an M/G demand process. *Mathematics of Operations Research*, 30(2), 354–368.
- Peskir, G., & Shiryaev, A. (2006). *Optimal stopping problems and free-boundary problems*. Basel: Birkhäuser Verlag.
- Porteus, E. L. (2002). *Foundations of stochastic inventory theory*. Stanford University Press.
- Stidham, S. (1974). Stochastic clearing systems. *Stochastic Processes and their Applications*, 2(1), 85–113.
- Stidham, S. (1977). Cost models for stochastic clearing systems. *Operations Research*, 25(1), 100–127.
- Stidham, S. (1986). Clearing systems and (s, S) inventory systems with nonlinear costs and positive leadtimes. *Operations Research*, 34(2), 276–280.
- Tijms, H. C. (2003). *A first course in stochastic models*. Chichester: John Wiley & Sons.
- Van Foreest, N. D., & Wijngaard, J. (2014). On optimal policies for production-inventory systems with compound Poisson demand and setup costs. *Mathematics of Operations Research*, 39(2), 533–560.
- Wang, J., Cao, J., & Liu, B. (2002). Unreliable production-inventory model with a two-phase Erlang demand arrival process. *Computers & Mathematics with Applications*, 43, 1–13.
- Wijngaard, J., & Stidham, S. (1986). Forward recursion for Markov decision processes with skip-free-to-the-right transitions, part I: Theory and algorithm. *Mathematics of Operations Research*, 11(2), 295–308.