DISCOUNT REPLENISHMENT OPPORTUNITIES WITH LIMITED CAPACITY: OPTIMAL POLICY AND PRICING DYNAMICS

Bram J. De Moor, Robert N. Boute, Stefan Creemers, Joren Gijsbrechts

REPRINT | 3357

CORE

Voie du Roman Pays 34, L1.03.01

B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email:lidam-library@uclouvain.be

https://uclouvain.be/en/research-institutes/lidam/core/core-reprints

ARTICLE IN PRESS

European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/eor

Discount replenishment opportunities with limited capacity: Optimal policy and pricing dynamics

Bram J. De Moor [a,*], Robert N. Boute [b,c,d], Stefan Creemers [e], Joren Gijsbrechts [f]

[a] Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands
[b] Research Center for Operations Management, KU Leuven, Belgium
[c] Technology and Operations Management Area, Vlerick Business School, Belgium
[d] Flanders Make@KU Leuven, Belgium
[e] Center for Operations Research and Econometrics (CORE), Université Catholique de Louvain, Belgium
[f] Department of Operations, Innovation and Data Sciences, Esade, Ramon Llull University, Spain

ARTICLE INFO

Keywords:
Inventory
Discount opportunities
Capacity constraints

ABSTRACT

We study an inventory model with occasional discounts on the fixed order cost, where a random capacity constraint limits the quantity ordered at the discount. We prove that the optimal policy resembles a can-order policy, characterized by a unique reorder, can-order, and order-up-to level. However, the random capacity constraints on the discount offers add complexity: the optimal policy features additional thresholds for inventory levels between the reorder and can-order point that determine when to use the discount. As these inventory-specific thresholds may render the optimal policy less straightforward to implement, we develop three heuristic replenishment policies inspired by the optimal policy structure. We compare their performance in a numerical study and derive an analytical bound on their worst-case optimality gaps in settings with an infinitely high backlog cost rate. With the buyer's optimal response to capacitated discounts in mind, we investigate the supplier's decision to set the discount price. A high discount relative to the regular order cost increases the probability that the buyer will accept the offer, but erodes the supplier's margin on discount sales and may reduce regular sales. We show that a mutually beneficial outcome requires buyer-to-supplier compensation.

1. Introduction

Analogous to carpooling platforms that connect drivers with empty seats to passengers, digital freight platforms aim to improve supply chain efficiency by advertising surplus transportation capacity at discounted rates on the fixed order cost (see, e.g., Teleroute). Shippers who wish to use this spare capacity face a critical trade-off: accepting the discount offers lowers truck operating costs, but doing so when inventories are high may lead to overstocking, while delaying orders in anticipation of discounts risks stockouts.

The shipper's optimal response to such discount offers on the fixed order costs is further complicated by capacity constraints, as the quantity available at the reduced rate is limited by the empty truck space. When the offered capacity is small, a shipper may choose to pass on the opportunity and wait for a future offer with greater volume or place an order at regular cost instead. The combination of uncertain discount timing and random capacity constraints creates a novel operational question: When faced with a time-sensitive opportunity to purchase a limited amount of goods at a reduced fixed order cost, should you place an order – and if so, in what quantity?

While digital freight platforms provide a motivating example for our work, the underlying decision problem is general and arises in various contexts. For instance, a manufacturer might temporarily reduce the setup cost for limited-size orders to incentivize buyers to use their spare production capacity during low-demand periods, thereby smoothing production. In such cases, the buyer has to decide whether or not to use this capacitated discount opportunity. Similarly, suppliers with excess inventory may offer temporary discounts on transportation costs to incentivize buyers to order. Again, the buyer has to analyze whether the discount is worth it, taking into account capacity constraints such as maximum discount order quantities set by the supplier. Despite their contextual differences, the examples share a common structure: a buyer[1] occasionally receives an opportunity from a supplier to

* Corresponding author.
E-mail addresses: b.j.k.de.moor@tue.nl (Bram J. De Moor), robert.boute@kuleuven.be (R.N. Boute), stefan.creemers@uclouvain.be (S. Creemers), joren.gijsbrechts@esade.edu (J. Gijsbrechts).

[1] To preserve generality, we refer to the "supplier" as the party offering the discount, and the "buyer" the one receiving the discount.

https://doi.org/10.1016/j.ejor.2026.04.016
Received 11 June 2025; Accepted 10 April 2026

replenish a limited quantity of a good at a reduced fixed order cost and needs to decide whether or not to use this opportunity.

To address how buyers should deal with capacity-constrained, reduced-rate offers, we develop an inventory model featuring both regular and discount orders. Regular orders are available at any time and incur a fixed cost, whereas discount orders, offered at random times and subject to stochastic capacity constraints, allow buyers to benefit from a lower fixed cost. Motivated by the examples above, we focus our analysis exclusively on discounts applied to the fixed order cost, while the variable (per-unit) ordering cost remains constant across both scenarios.

Although discount opportunities and capacity-constrained orders have been explored individually, we integrate and simultaneously address both complexities. We show that the random capacity constraints on the discount offers complicate the optimal policy: it resembles a can-order policy (characterized by a reorder, can-order, and order-up-to level), but features additional thresholds for inventory levels between the reorder and can-order point that determine when to use the discount. If the inventory level drops to the reorder point, it is optimal to place an order to replenish up to the order-up-to level, even without a discount. If the inventory level is between the reorder and can-order point, it is only optimal to accept a discount offer if the quantity that can be ordered at the discount exceeds an inventory-level specific threshold. In that case, it is optimal to replenish up to the order-up-to level if the discount's capacity constraint allows to do so. If not, we conjecture that it is optimal to replenish inventory as close to the order-up-to level as possible, given the capacity constraint.

As the optimal policy's thresholds limit its practical ability, we propose three heuristic replenishment rules. The heuristics differ in computational complexity and ease of implementation. We numerically demonstrate they perform well and provide a worst-case guarantee in settings with an infinitely high backlog cost rate.

Studying how buyers respond to discount offers with random capacity constraints enables an analysis of the magnitude of the discount, a key decision variable for the supplier. A higher discount on the fixed order cost increases the likelihood that a buyer will take advantage of the offer. However, excessive price reductions erode the supplier's profit margin from discount sales, and may reduce revenue from regular sales. Our analysis reveals that the supplier's optimal discount leads to a low acceptance rate of discount offers, below the level that would be ideal for the buyer and supplier combined. Through an additional buyer-to-supplier compensation, a mutual beneficial scenario can emerge, where both the supplier and the buyer achieve better outcomes compared to the supplier's standalone optimum.

Our main contributions can be summarized as follows: (1) we partially characterize the optimal replenishment policy for an inventory problem with occasional discount opportunities on the fixed order cost with random capacity constraints; (2) we propose three novel replenishment heuristics for this inventory problem; and (3) we study the dynamics that arise when setting the discount price.

2. Related literature

We focus on stochastic-demand inventory models with fixed order costs and discount opportunities, but note that substantial work has also been conducted on their deterministic counterparts (see, e.g., Chaouch (2007), Goh and Sharafali (2002), Moinzadeh (1997), and Tajbakhsh et al. (2011), among others). However, in both the stochastic- and deterministic-demand literature streams, the impact of capacity constraints on discount orders remains unexplored.

Stochastic inventory models with fixed order costs have been researched extensively, dating back to the seminal work of Scarf (1960) introducing K-convexity to establish the optimality of (s, S)-policies: if inventory drops below the reorder point s, it is optimal to incur the fixed order cost and replenish up to the order-up-to level S. Beckmann (1961) extends this result for continuous-review systems. Recent reviews of stochastic inventory systems with fixed order costs are provided by Perera and Sethi (2023a, 2023b). Henceforth, we focus on models with discount opportunities and capacity constraints.

In a system with Poisson discount opportunities, at which the fixed order cost temporarily drops, Zheng (1994) proves the optimality of the can-order (s, c, S)-policy. First studied by Friend (1960) and later coined by Balintfy (1964), the can-order (s, c, S)-policy places an order when the inventory is at or below the reorder point s to replenish the inventory up to the order-up-to level S. When a discount opportunity arises, an order is placed when the inventory is at or below the can-order point c to replenish up to the order-up-to level S. Compared to Zheng (1994), our model assumes that the discount opportunities are capacity-constrained by a random capacity size. We find that the optimal policy is a generalization of the can-order policy, where the decision to use a discount opportunity depends on its size and the current inventory level. The can-order policy has also been used to synchronize replenishments across multiple SKUs in the well-studied joint replenishment problem. In that setting, discount opportunities decompose the multi-SKU problem into single-SKU problems, where a discount opportunity represents an order triggered by another SKU. Our problem can thus also be interpreted as a capacitated variant of the joint replenishment problem, with the random capacity serving as a shared, limited-ordering resource.

Capacity constraints on the order quantities complicate the optimal replenishment policy. When there are no fixed order costs, Federgruen and Zipkin (1986a, 1986b) show that the optimal policy changes from a conventional base-stock policy to a modified base-stock policy when the orders are constrained by a fixed capacity: if the inventory is below a certain base-stock level, it is optimal to replenish to the base-stock, or as close to the base-stock as possible, given the capacity constraint. This led to the conjecture that under a fixed order cost and capacity constraints, a modified (s, S)-policy is optimal, where you order as close as possible to the order-up-to level S if inventory is at or below s (Federgruen & Zipkin, 1986a, 1986b). However, Wijngaard (1972), Shaoxiang (2004) and Shaoxiang and Lambrecht (1996) provide counterexamples, showing that the optimal policy is hard, if not impossible, to characterize. In particular, Shaoxiang and Lambrecht (1996) establish the existence of XY-bands. If inventory drops below X, it is optimal to order up to capacity. If inventory is above Y, it is optimal to order nothing. Between X and Y, the optimal policy is complex. Gallego and Scheller-Wolf (2000) extend (Scarf, 1960)'s K-convexity to CK-convexity, where C is the fixed capacity, to further (partially) characterize the optimal policy between the X- and Y-band. Recently, Rossi et al. (2024) show that a modified (s, S)-policy with multiple reorder and order-up-to levels is optimal, provided that the "continuous order property" holds, which states that if it is optimal to order at an inventory level x, it is also optimal to order at inventory levels $y < x$. In our work, only the discount opportunity is capacity-constrained by a random capacity size. The randomness in both the occurrence and size of the capacity constraint complicates our replenishment problem further. Since we do not have a single capacity constraint C, it is challenging, if not impossible, to exploit existing insights derived using CK-convexity. Therefore, we resort to alternative proof techniques to (partially) derive the optimal response to capacitated discount opportunities.

3. An inventory problem with capacitated discount opportunities

3.1. Model description

To address how buyers should deal with capacity-constrained, reduced-rate offers, we formulate a discounted, infinite-horizon continuous-time Markov decision process that is driven by two distinct types of independent events: (1) demand, which follows a Poisson process with rate λ; and (2) discount offers, which occur according to a Poisson process with rate μ and indicate the possibility to order at a lower fixed order cost. The time between two events follows an exponential distribution with rate $(\lambda + \mu)$. The next event is triggered by a demand arrival with probability $\lambda/(\lambda + \mu)$, and by a discount opportunity with probability $\mu/(\lambda + \mu)$.

Upon each nth event, the buyer must decide how many units $q_n \in \mathbb{Z}$ to order based on the state (x_n, r_n), consisting of the current inventory level x_n and the available capacity of r_n units that can be ordered at a discount cost k. For the sake of notational brevity, we assume zero lead time, so that an order instantaneously raises the inventory level to $x_n + q_n$. We note, however, that our results continue to hold under a constant, positive lead time (see Section 3.4). When the next decision epoch is triggered by a demand arrival, with probability $\lambda/(\lambda + \mu)$, the system transitions to state $(x_{n+1} = x_n + q_n - 1, r_{n+1} = 0)$. Alternatively, when a decision epoch is triggered by a discount opportunity arrival, with probability $\mu/(\lambda + \mu)$, a transition is made to state $(x_{n+1} = x_n + q_n, r_{n+1} = R)$. Here, R is a random variable with probability mass function $\phi_R(r) = \mathbb{P}(R = r)$, cumulative distribution function $\Phi_R(r) = \mathbb{P}(R \le r)$, and finite positive support $\bar{r}$.

We exclusively analyze scenarios involving unit demand arrivals and positive integer discount opportunities, denoted as $r_n \in \mathbb{Z}^+$. Consequently, we restrict our focus to integer inventory and order levels. Order quantities can assume negative values, indicating the possibility of discarding inventory, which incurs no cost. The following assumption encapsulates this principle:

Assumption 1. Inventory can be discarded at no cost.

We allow negative orders to prove Lemma 1 and characterize the structure of the two-dimensional value functions that define our problem. It is worth noting, however, that the optimal policy does not include negative order quantities in the recurring states. This suggests that the optimal policy we derive in Section 3.2 is also optimal if we do not allow negative order quantities.

Orders not exceeding the discount opportunity size, $0 < q_n \le r_n$, incur a discounted fixed cost k, while larger orders, $q_n > r_n$, incur the regular fixed order cost K, with $k \le K$. The ordering cost in decision epoch n thus equals:

$$k\mathbb{1}_{0<q_n\le r_n} + K\mathbb{1}_{r_n<q_n<\infty},$$

where $\mathbb{1}_{y\le q_n\le y'}$ equals 1 if $q_n \in [y, y']$, and 0 otherwise.

Inventory costs between two decision epochs are incurred at a holding cost rate $h > 0$ or a backlog rate $b > 0$, depending on whether $x_n + q_n$ is positive or negative, respectively. Holding and backlog costs are discounted continuously using a discount rate $\alpha > 0$, so that a cost incurred at time t is weighted by $e^{-\alpha t}$, similar to Puterman (1994). Since the time τ between two decision epochs follows an exponential distribution with rate $(\lambda + \mu)$, the expected discounted holding and backlog cost accumulated over this interval can be written as:

$$\begin{aligned} L(x_n + q_n) &= \mathbb{E}\left[\int_0^\tau \left(h[x_n + q_n]^+ - b[x_n + q_n]^-\right) e^{-\alpha t} dt\right], \\ &= \left(h[x_n + q_n]^+ - b[x_n + q_n]^-\right) \mathbb{E}\left[\frac{1 - e^{-\alpha\tau}}{\alpha}\right], \\ &= \frac{h[x_n + q_n]^+ - b[x_n + q_n]^-}{\lambda + \mu + \alpha}, \end{aligned}$$

with $[y]^+ = \max(0; y)$ and $[y]^- = \min(0; y)$.

Thus, the expected (present-valued) cost incurred between two decision epochs when visiting state (x_n, r_n) and ordering q_n is:

$$C\left((x_n, r_n), q_n\right) = L(x_n + q_n) + k\mathbb{1}_{0<q_n\le r_n} + K\mathbb{1}_{r_n<q_n\le\infty}. \quad (1)$$

The objective of our inventory problem is to find an admissible, stationary order policy π^* that minimizes the expected infinite horizon present-valued costs for every state. The policy π is a mapping from state space to order quantities: $\pi : (\mathbb{Z}, \mathbb{Z}^+) \mapsto \mathbb{Z}$. That is, based on the inventory and the size of the discount opportunity, (x_n, r_n), the policy π determines the quantity ordered in decision epoch n by $q_n = \pi(x_n, r_n)$. Since the time τ between two decision epochs follows an exponential distribution with rate $(\lambda+\mu)$, the expected discount factor applied to the cost incurred at the next decision epoch is $\mathbb{E}\left[e^{-\alpha\tau}\right] = (\lambda+\mu)/(\lambda+\mu+\alpha)$. For brevity, we define $\gamma = (\lambda+\mu)/(\lambda+\mu+\alpha)$. Then, the expected present value of future costs when starting in state (x_n, r_n) and henceforth following policy π can be expressed as:

$$v^\pi(x_n, r_n) \triangleq \lim_{N\to\infty} \mathbb{E}^\pi_{(x_n,r_n)}\left[\sum_{i=0}^{N} \gamma^i C\left((x_{n+i}, r_{n+i}), \pi(x_{n+i}, r_{n+i})\right)\right]. \quad (2)$$

Note that we can refrain from using time subscripts on the states, since we assume stationary rewards and transition probabilities and $0 \le \gamma < 1$. The optimal value function $v^*(x, r)$ in state (x, r) is then defined as the infimum of $v^\pi(x, r)$ and denotes the expected present value of future costs when starting in state (x, r) and following the optimal policy π^*:

$$v^*(x, r) = v^{\pi^*}(x, r) \triangleq \inf_{\pi\in\Pi} v^\pi(x, r),$$

with Π the set of all admissible policies. The optimal policy is known for two special cases: in the absence of discount opportunities, (s, S)-ordering policies are optimal (Beckmann, 1961; Scarf, 1960). When the discount opportunities are uncapacitated, the can-order (s, c, S)-policy is optimal (Zheng, 1994). In the next section, we partially derive the optimal policy when the discount opportunities are capacitated.

3.2. Optimal replenishment policy

We prove that the optimal ordering policy for the inventory problem described in Section 3.1 works as follows. Equivalent to the (s, S)-policy, if demand depletes inventory to a threshold s, the buyer initiates an order at regular cost K, raising the inventory level to S. This corresponds to the blue area in Fig. 1. A more complex structure arises when an order can be placed at reduced order cost k that is capacity-constrained by the discount opportunity size r. The buyer uses the discount offer if the inventory level x is at or below c *and* if the size of the discount opportunity r exceeds an inventory level-specific threshold e_x units, visualized by the green zone in Fig. 1.

The can-order point c and thresholds e_x determine when it is optimal to use a discount opportunity. Deriving the optimal order quantity when using a discount opportunity, however, is more challenging. We can show that if it is optimal to use a discount opportunity and the buyer can reach the order-up-to-level S given the capacity constraint, i.e., if $S - x \le r$, then ordering up to S is optimal. Yet, if the size of the discount opportunity does not permit reaching the order-up-to level S, i.e., if $S - x > r$, we cannot analytically prove the optimal order quantity. In this case, we can only conclude that the optimal order quantity lies between the threshold e_x and the available capacity r. We thus present a partial characterization of optimal policy:

Theorem 1 (*Optimal Ordering Policy*). *The optimal order quantities for the inventory system with fixed costs and capacitated discount opportunities are given by:*

$$\pi^*(x, r) = \begin{cases} S - x & \text{if } x = s, \\ S - x & \text{if } s < x \le c \text{ and } r \ge S - x \text{ and } r \ge e_x, \\ \pi^*(x, r) \in \{e_x, e_x + 1 \cdots, r - 1, r\} & \text{if } s < x \le c \text{ and } r < S - x \text{ and } r \ge e_x, \\ 0 & \text{if } s < x \le c \text{ and } r < e_x, \\ 0 & \text{if } c < x \le S, \end{cases} \quad (3)$$

with the reorder point s, the can-order level c, the order-up-to level S, and the vector $\mathbf{e} = (e_{s+1}, e_{s+2}, \ldots, e_{c-1}, e_c)$ its parameters.

In our numerical experiments, we did not find any counterexample in which the optimal policy orders less than the available capacity when the discount opportunity is large enough to trigger an order (i.e., $r \ge e_x$

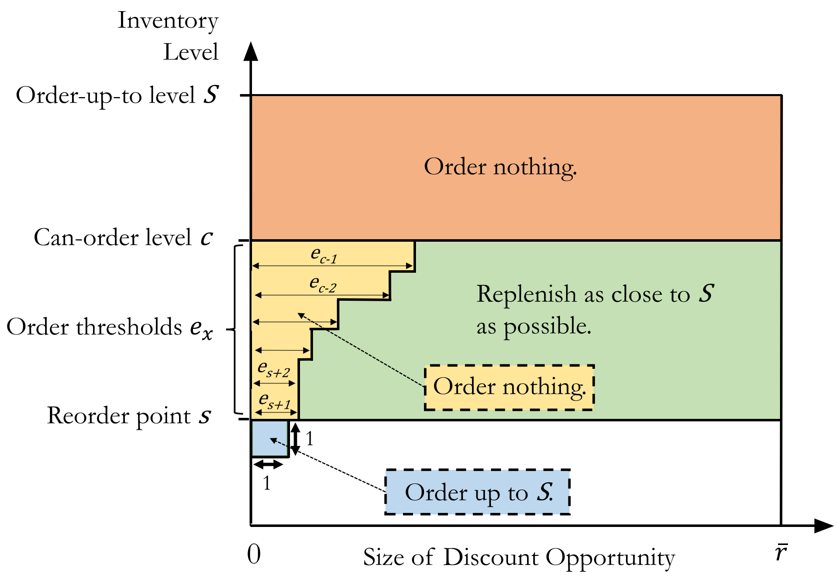

Fig. 1. The optimal replenishment policy, as per Theorem 1 and Conjecture 1. Note that the reorder point s is never positive in the case of zero lead time, so that the vertical axis does not necessarily start at 0. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and $s < x \leq c$) but insufficient to raise the inventory level to S (i.e., $r < S - x$). In other words, if it is optimal to use a discount opportunity but the capacity constraint does not allow to reach S, we conjecture that it is optimal to order at full capacity r to come as close as possible to S. We acknowledge that counterexamples to this claim exist in models with compound demand, fixed capacity constraints, and no discount opportunities (see, e.g., Shaoxiang (2004) and Shaoxiang and Lambrecht (1996)). However, we believe that our setting is sufficiently different, due to the random occurrences of stochastic capacity constraints, to justify the following conjecture:

Conjecture 1. *For states (x, r), with $s < x \leq c$, $r < S - x$, and $r \geq e_x$, the optimal order quantity $\pi^*(x, r)$ is equal to r.*

We relegate the proof of Theorem 1 to Appendix A and provide a sketch below. Our analysis relies on the well-known Bellman optimality equations for each state (x, r):

$$v^*(x,r) = \min_{q\in\mathbb{Z}} \Bigg[C((x,r),q) + \gamma \Bigg[\frac{\lambda}{\lambda+\mu} v^*(x+q-1,0) + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+q,r') \Bigg]\Bigg]. \tag{4}$$

The fact that our value function is two-dimensional complicates characterizing the value function's structure fully. We can, however, derive an important subset of results to develop a partial characterization of the optimal policy structure.

Our first result, formalized in Lemma 1, is intuitive: (a) given any discount opportunity size r, the minimum expected cost in states (x, r) with an inventory level x is always lower than or equal to the minimum expected cost in states (x', r) with an inventory level $x' < x$, because inventory can be discarded at no cost; and (b) given any inventory level x, the minimum expected cost in states with a discount opportunity size r is always lower than or equal to the minimum expected cost in states (x, r') with a discount opportunity size $r' < r$.

Lemma 1. *The value functions are non-increasing in inventory and size of discount opportunity:*

(a) $v^*(x,r) \geq v^*(x+\epsilon, r)$, *with* $\epsilon \in \mathbb{N}$, $\forall\,(x,r) \in (\mathbb{Z},\mathbb{Z}^+)$.
(b) $v^*(x,r) \geq v^*(x, r+\epsilon)$, *with* $\epsilon \in \mathbb{N}$, $\forall\,(x,r) \in (\mathbb{Z},\mathbb{Z}^+)$.

Next, we introduce $G(y)$, which is the expected sum of discounted costs given inventory level y if we follow the optimal policy from the next decision epoch onward:

$$G(y) = L(y) + \gamma \Bigg[\frac{\lambda}{\lambda+\mu} v^*(y-1,0) + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(y,r') \Bigg]. \tag{5}$$

The minimizer of $G(y)$ represents the inventory level we would like to have. Getting to this desired inventory level, however, may require one of the following actions: (i) discarding inventory; (ii) using a discount opportunity at cost k; or (iii) placing an order at regular cost K. Re-writing the value functions $v^*(x, r)$ in function of $G(y)$:

$$v^*(x,r) = \min\left[\min_{y\leq x} G(y),\ \min_{x<y\leq x+r} [G(y)+k],\ \min_{x+r<y} [G(y)+K] \right], \tag{6}$$

and using Lemma 1, allows to derive the following structural properties of $G(x)$:

Lemma 2. *We find that $G(x)$:*

(a) is strictly decreasing in the interval $x \in\ ]-\infty, 0\,]$.
(b) is strictly increasing in the interval $x \in [S+1,\ \infty[$, with $S \geq 0$ a global minimum of $G(x)$.
(c) has a unique global minimum in S with $S \geq 0$, except in the (pathological) case where $v^*(S-1,0) = \frac{h}{\lambda\alpha}(\lambda+\alpha+\lambda S)$ in which case $G(S) = G(S+1)$.

Lemmas 2(a) and 2(b) show that G is strictly decreasing for negative inventory levels and strictly increasing above a point $S+1 > 0$, respectively. These results allow us to show in Lemma 2(c) that there exists a unique inventory level $S \geq 0$ at which $G(x)$ is minimal. In one specific pathological case, i.e., if $v^*(S-1,0) = \frac{h}{\lambda\alpha}(\lambda+\alpha+\lambda S)$, $G(x)$ will have two adjacent global minima in S and $S+1$. Then, the optimal decision maker is indifferent between inventory levels S and $S+1$. We derive the condition to obtain this case in Section S.1 in the online Supplementary Material.

We illustrate an example of G for a specific numerical setting in Fig. 2 to explain how the shape of G determines the optimal ordering policy. Without fixed order costs, ordering up to S in every decision epoch is optimal. With fixed costs, however, ordering to inventory level y from inventory level x is only optimal if the resulting gain in G, i.e., $G(x) - G(y)$, outweighs the order cost, K or k. From Lemma

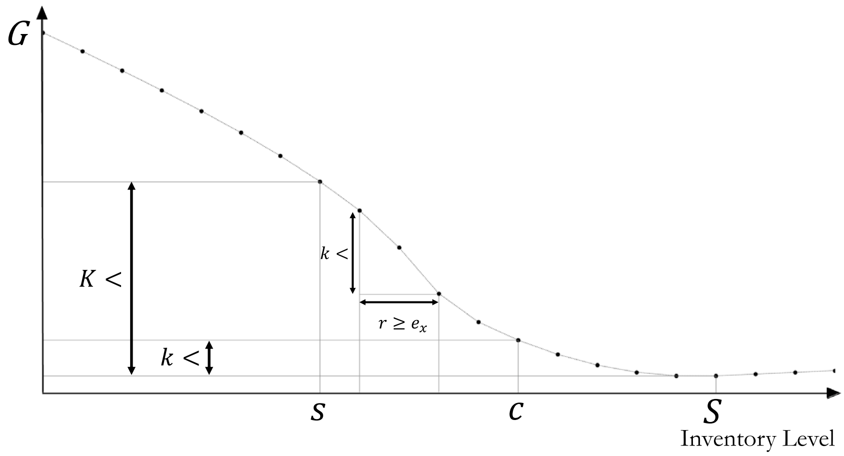

Fig. 2. Plot of $G(x)$, see Eq. (5), which can be interpreted as the expected sum of discounted costs of having an inventory level of x, and acting optimally from the next decision epoch on.

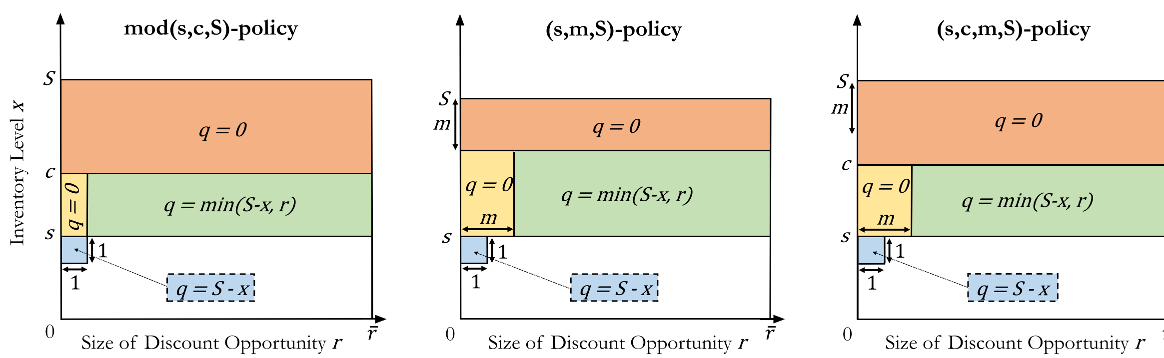

Fig. 3. Three new replenishment heuristics that allow buyers to respond to discount replenishment opportunities with limited capacity.

2(a), we derive that there must exist at least one point s such that $G(x) - G(S) \geq K$ for $x \leq s$. In Appendix A, we prove this point s is unique. Thus, when there is no discount opportunity, it is optimal to place an order to replenish up to S when inventory drops to s because the resulting decrease in G outweighs the fixed order cost K. Similarly, we prove that a unique can-order point c exists, such that $G(x) - G(S) \geq k$ for $x \leq c$. If there is a discount opportunity r at inventory level $x \leq c$ such that the size of the discount opportunity permits to reach S, i.e., $r \geq S - x$, it is optimal to replenish up to S because the gain in G outweighs the order cost k. If r is too small to reach S, placing an order may not be optimal. In Appendix A, we prove that for each inventory level x with $s < x \leq c$ there exists a threshold e_x: only if $r \geq e_x$, it is optimal to place an order. We conjecture that, in this case, it is optimal to place an order to get as close as possible to S.

We note that e_x is not always increasing in the inventory level x, as suggested by Fig. 1. Situations arise where it is sub-optimal to use a discount opportunity of size r for a low inventory level, while it is optimal to use a discount opportunity with the same size r for higher inventory levels. This behavior can be explained by further analyzing the structure of G. In Fig. 2, we observe that G has an S-shape in $[s, c]$: the slope of G is less steep close to s, then becomes steeper, after which it flattens out close to c. While we cannot analytically prove this structure, we can provide an intuition. For inventory levels close to s, instead of using a discount opportunity, it may be better to wait until an uncapacitated order at regular cost K is placed. Discount opportunities are thus less attractive for inventory levels close to s, which explains the shallow slope of G in that domain. For higher inventory levels, waiting until the inventory level hits s is more costly and using a discount opportunity is more attractive, resulting in a steeper slope for G. For inventory levels close to c, however, a discount opportunity is again less attractive as the inventory level is already high. This explains the shallow slope of G close to c.

3.3. Three simple decision rules

As the optimal policy may not be straightforward to implement due to the inventory-specific thresholds e_x, we introduce three more practical decision rules based on the characteristics of the optimal policy. We visualize the heuristics in Fig. 3. In what follows, we use superscript "*" to refer to the parameters of the optimal policy. The first heuristic is an extension of the can-order (s, c, S)-policy that accounts for the capacity constraints of the discount opportunity: the modified (s, c, S)-policy, or in short mod(s, c, S)-policy.

Definition 1 (*The* mod(s, c, S)*-Policy*)**.** The order placed by a mod(s, c, S)-policy with parameters s, c, and S in state (x, r) is:

$$q(x, r) = \begin{cases} S - x & \text{if } x \leq s, \\ \min(S - x, r) & \text{if } s < x \leq c, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

The mod(s, c, S)-policy is intuitive and easy to implement. A downside of the mod(s, c, S)-policy is that it also orders in states where the discount opportunity size is too small to be cost-optimal (i.e., it assumes $e_x = 1$ for all $s < x \leq c$). As a result, we observe that the can-order level c in the mod(s, c, S)-policy is below the can-order level c^* in the optimal policy to prevent excessive ordering using discount opportunities. To respond to this deficiency, we introduce a second heuristic that features a minimum order quantity m:

Definition 2 (*The* (s, m, S)*-Policy*)**.** The order placed by an (s, m, S)-policy with parameters s, m, and S in state (x, r) is:

$$q(x,r) = \begin{cases} S - x & \text{if } x \leq s, \\ \min(S - x, r) & \text{if } r \geq m \text{ and } S - x \geq m \text{ and } x > s, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

A minimum order parameter m now determines the decision to seize a discount opportunity: when the discount opportunity has size r, an order of size $\min(S{-}x, r)$ is placed *if* we can order at least m units. In contrast to the optimal policy, the threshold to use the discount opportunity is independent of the current inventory level x in the (s, m, S)-policy: a single parameter m is used to capture the vector $\mathbf{e}$ (i.e., $e_x = m$ for all $s < x \leq S - m$). The (s, m, S)-policy has no can-order parameter c, and the decision at which inventory level to use a discount opportunity only depends on $S - m$. Therefore, the order-up-to level S in the (s, m, S)-policy will be lower than the order-up-to level S^* in the optimal policy. We can avoid this discrepancy by using both a can-order parameter c and a minimum order parameter m. This leads to our third heuristic that combines (and thus dominates) the mod(s, c, S)- and (s, m, S)-policies.

Definition 3 (*The* (s, c, m, S)*-Policy*)**.** The order placed by an (s, c, m, S)-policy with parameters s, c, m, and S in state (x, r) is:

$$q(x,r) = \begin{cases} S - x & \text{if } x \leq s, \\ \min(S - x, r) & \text{if } x \leq c \text{ and } r \geq m \text{ and } S - x \geq m \text{ and } x > s, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

In this case, a discount opportunity of size r is used if the inventory level x is at or below c and if at least m units can be ordered. Observe that the optimal policy is, in fact, an $(s, c, \mathbf{m}, S)$-policy with inventory level-specific values $\mathbf{m} = \mathbf{e} = \{e_{s+1}, e_{s+2}, \ldots, e_{c-1}, e_c\}$. Similar to the (s, m, S)-policy, we capture the vector $\mathbf{e}$ by a single parameter m (i.e., $e_x = m$ for all $s < x \leq c$). In settings where the difference between the values of s^* and c^* in the optimal policy is small, the dimension of vector $\mathbf{e}$ will also be small. Then, the information loss from representing all the elements of $\mathbf{e}$ by the constant value m will be limited, and the (s, c, m, S)-policy will be close to optimal.

We evaluate the heuristics' performance by examining their optimality gaps in terms of the average cost per period. Since the total discounted cost depends on the initial state, the average cost offers a more interpretable and state-independent basis for comparison (Zipkin, 2008). In settings with infinitely high backlog cost b, we can derive analytical expressions for the maximum optimality gap through an upper bound on the heuristics' average cost and a lower bound on the optimal policy's cost. The intuition behind the bounds is as follows (we provide the detailed derivation in Section S.2 in the online Supplementary Material). First, observe that our three heuristics can mimic the (s, Q)-policy,[2] which only places regular order of size Q at cost K whenever the inventory level reaches the reorder point s. Its average cost thus provides an upper bound that can be expressed in closed form if $b = \infty$. Then, our upper bound is equal to: $\overline{C} = \lambda K / (\sqrt{2\lambda K/h} + 1) + \sqrt{\lambda K h/2}$. A lower bound can be obtained using the optimal can-order policy where discount opportunities are not capacity-constrained. Even though its cost does not have a closed-form expression, we can consider the extreme case where discount orders can always be placed (i.e., where $\mu = \infty$). In this case, the can-order policy again reduces to an (s, Q)-policy with a cost that corresponds to our lower bound: $\underline{C} = \sqrt{2\lambda k h} - h/2$. Combining both results allows us to define the worst-case performance of our heuristics when $b = \infty$:

[2] In the literature, this policy is often denoted as the (r, Q)-policy. Since we already use r to denote the realization of the discount opportunity size, we instead use s for the reorder point.

Proposition 1. *When the backlog cost $b = \infty$, the optimality gap of our heuristics will not exceed:*

$$\frac{\overline{C} - \underline{C}}{\underline{C}} = \frac{K/(2\sqrt{K} + \sqrt{2h/\lambda}) + \sqrt{K}/2}{\sqrt{k} - \sqrt{h/8\lambda}} - 1.$$

We note that the worst-case performance guarantee is generally not very tight. For instance, when $K = 50$, $k = 20$, $h = 1$, and $\lambda = 5$, the maximum optimality gap is around 60%. As we will demonstrate numerically, however, the heuristics' actual optimality gaps remain well below this bound in such settings. Analyzing the asymptotic behavior of the maximum optimality gap expression reveals that the gap increases in K and h, indicating that the heuristics are theoretically not guaranteed to perform well in environments with high order and inventory costs. In contrast, as k increases, the maximum optimality gap of the heuristics decreases because the optimal policy converges to an (s, Q)-policy, which the heuristics can mimic. Finally, the maximum optimality gap decreases w.r.t. λ, converging to $\sqrt{K/k} - 1$, suggesting that the worst-case performance does not depend on the buyer's demand volume.

As the bounds do not allow to draw conclusions on the individual performance differences between the heuristics, we additionally perform a numerical study. We compare the optimality gaps of the mod(s, c, S)-, the (s, m, S)-, and the (s, c, m, S)-policy in 360 numerical settings for different values of λ, μ, b, k, K, and p, with p the parameter of a delayed geometric distribution (having strictly positive support) that determines the size of the discount opportunity. To ensure a finite state space, we bound the maximum possible discount opportunity size by setting $\bar{r}$ as the smallest integer satisfying $\mathbb{P}(R \leq \bar{r}) > 0.99$. We standardize the holding cost rate $h = 1$ while varying the other parameters (see the horizontal axis of Fig. 4 for the different values considered).

We calculate the optimal replenishment policy using policy iteration with a continuous-time rate $\alpha = 0.01$. Our policy iteration implementation is based on (Puterman, 1994) and can be consulted via the following GitHub Repository. The heuristics are optimized by enumerating over all possible parameter combinations and calculating the average cost per period using Markov chain analysis (i.e., determining the steady-state distribution and taking the expected cost per period with respect to this distribution). To optimize the policies more efficiently, a gradient descent algorithm can be used that iteratively adjusts the policy parameters by checking the cost of neighboring solutions. As we know that the optimal policy can be described by s^*, c^*, S^*, and $\{e^*_{s+1}, e^*_{s+2}, \ldots, e^*_c\}$, we can directly find the best policy parameters using a gradient descent algorithm, instead of trying to find the structure from scratch. Numerical checks suggest convexity in s^*, S^*, and e^*, but not in c^*. Hence, to optimize c^*, different step sizes or brute force might be needed to escape local optima.

Fig. 4 shows the optimality gaps in boxplots for each parameter setting. The different colors correspond to the different heuristics. The (s, c, m, S)-policy (visualized in red) outperforms the other heuristics and has an average optimality gap of 1.7%. Removing either the can-order point c or the minimum order parameter m decreases performance. The mod(s, c, S)-policy (visualized in green) has an average optimality gap of 7.8%, and the (s, m, S)-policy (visualized in blue) has an average optimality gap of 5.2% in the settings considered. The better performance of the (s, m, S)-policy compared to the mod(s, c, S)-policy highlights the importance of the minimum order parameter m. This illustrates that minor, intuitive adaptations to the can-order policy are insufficient when discount opportunities are capacitated.

Although the average optimality gaps are relatively small across heuristics, Fig. 4 shows substantial variability across parameter settings. All three heuristics perform well when discount opportunities are rare (i.e., for low values of $\mu/(\lambda+\mu)$) or when the discount opportunities have a large expected capacity size (i.e., for high values of $1/p$). In the former case, the (s, Q)-policy is near-optimal, and in the latter case, the optimal policy resembles a conventional can-order policy. Because the

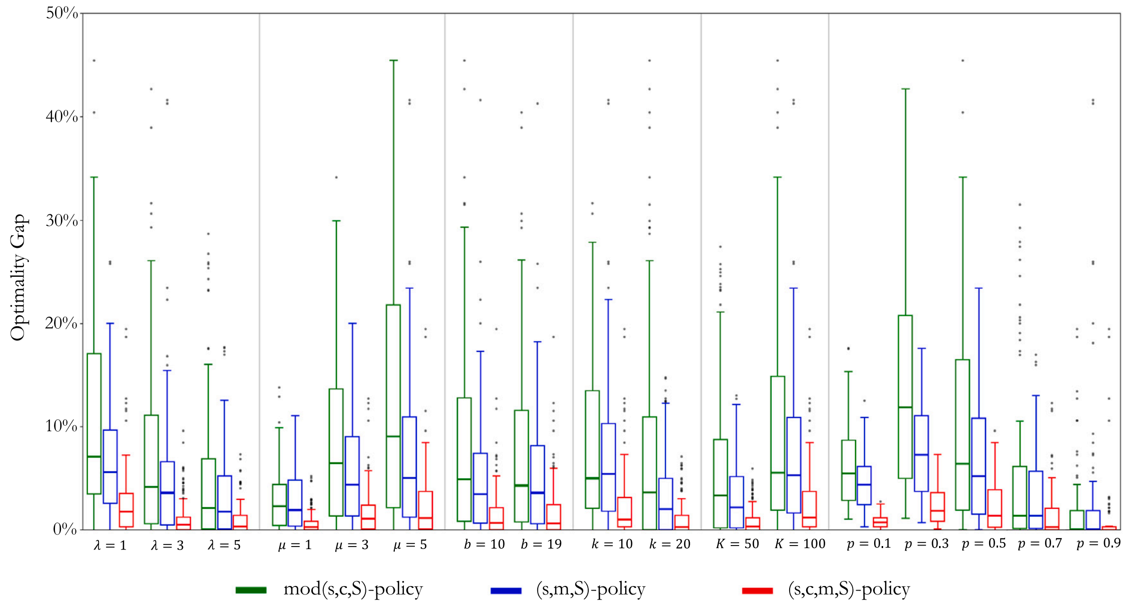

Fig. 4. Optimality gaps of the mod(s, c, S)-, (s, m, S)-, and (s, c, m, S)-policy in different settings. We observe that the (s, c, m, S)-policy dominates, followed by the (s, m, S)- and the mod(s, c, S)-policy. The holding cost rate $h = 1$ in all settings. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

three heuristics can mimic both the (s, Q)- and (s, c, S)-policy in these regimes, their performance gaps are small.

When discount opportunities occur frequently (i.e., when $\mu/(\lambda+\mu)$ is high) and the average discount size is small (i.e., when $1/p$ is low), the optimal policy becomes more complex, which leads to larger optimality gaps across all heuristics. Nevertheless, the (s, c, m, S)-policy achieves optimality gaps below 10% in most instances. The instances with gaps above 10% are characterized by a high frequency of small discount opportunities and by discounted ordering costs that are low relative to regular ordering costs (i.e., low k/K). In these settings, the (s, c, m, S)-policy often fails to exploit small discount opportunities at moments when doing so would be optimal. A similar decline in performance is observed for the (s, m, S)-policy, although the effect is even more pronounced. In settings with frequent small discount opportunities, the mod(s, c, S)-policy tends to order too often at a discount when the available capacity r is small, resulting in elevated optimality gaps, particularly when the discounted order cost k and the regular order cost K are high.

3.4. Discussion of assumptions

Our analysis relies on several modeling assumptions. In this section, we discuss these assumptions and examine their implications through analytical arguments or numerical experiments. The main results are summarized below, with a detailed analysis provided in Section S.3 of the online Supplementary Material.

Lead time. If the replenishment lead time for both regular and discount orders is constant and positive, the system state can be captured by the inventory position alone (i.e., on-hand inventory - backorders + inventory in the pipeline), under our backlogging assumption (similar to Zheng (1994)). Section S.3.1 in the Supplementary Material shows that a constant lead time does not alter the structure of the optimal policy. We additionally performed a numerical experiment to analyze the sensitivity of the three heuristics to increasing lead times (for details, see Section S.3.1). We find that the performance of the mod(s, c, S)-policy is not impacted. The (s, m, S)- and (s, c, m, S)-policies, however, exhibit lower optimality gaps for increasing lead times. Longer lead times inflate both the optimal reorder point s and can-order point c, but c rises more rapidly. As a result, the region in which using a discount opportunity is considered expands, creating more opportunity to wait for an attractive discount. The (s, m, S)- and (s, c, m, S)-policy can exploit this effect, but the mod(s, c, S)-policy cannot.

Variable order costs. When the order costs have a fixed and a variable component that depends on the order size, and only the fixed component is subject to a discount, we show in Section S.3.2 that Theorem 1 remains valid under the mild condition that the per-unit variable order cost is smaller than the expected per-unit backlog cost between two decision epochs. However, when the variable order cost is also subject to a discount, this creates a different problem which further complicates the optimal replenishment policy (see, e.g., Moinzadeh (1997)).

Fixed discount price. We assume that the discount price k is fixed and does not depend on the size of the discount opportunity. This assumption is motivated by our primary example, in which K denotes the fixed cost of operating a truck and k reflects the additional handling cost incurred when another shipper utilizes the available spare capacity through a digital freight platform. Regardless of the magnitude of the spare capacity (i.e., discount opportunity), this incremental cost remains unchanged. Nevertheless, we conducted a small numerical experiment (details in Section S.3.3) where the discount depends linearly on the size of the opportunity: a shipper pays $k = \beta r$ when using a discount opportunity of size r, where β is the per-unit cost of a discount opportunity. We find that the optimal policy continues to be characterized by a reorder, can-order, and order-up-to level. However, the ordering decisions between the reorder and can-order level differ from those described in Theorem 1. In particular, there are no longer inventory level-specific thresholds that determine when to use a discount opportunity. For a given inventory level x, it may be optimal to use a discount opportunity in state (x, r) but not in state $(x, r + \epsilon)$, for some $\epsilon \in \mathbb{N}$.

Discount opportunity arrivals. We assume that discount opportunities arrive according to a Poisson process and are independent of the demand process. This assumption is standard in the joint replenishment literature and is analytically convenient (Silver, 1974; Zheng, 1994). It is further justified in settings where discount opportunities arise from the aggregation of many independent processes, such as freight-sharing platforms, where spare transportation capacity emerges from

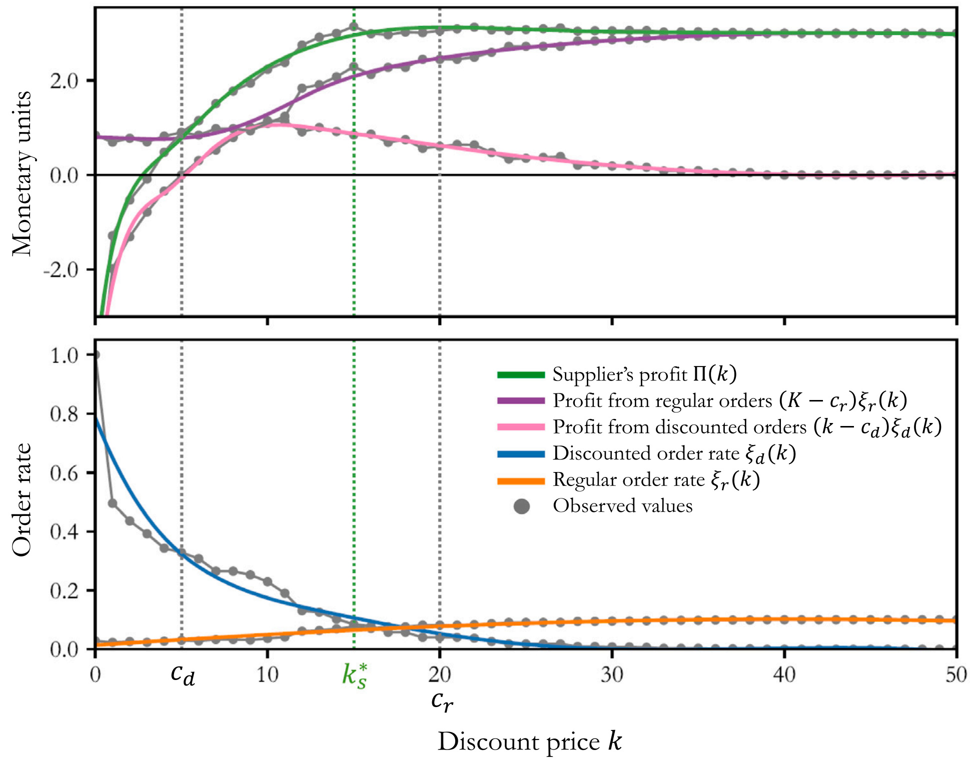

Fig. 5. The supplier's profit (green) consists of profit from regular orders (purple), i.e., $(K - c_r)\xi_r(k)$, and profit from discounted orders (pink), i.e., $(k - c_d)\xi_d(k)$. It can be shown that the rate at which discounted orders are placed $\xi_d(k)$ (blue) is monotonically decreasing in k. The rate at which regular orders are placed $\xi_r(k)$ does not always behave monotonically. The setting considered has parameters $\lambda = 1$, $\mu = 1$, $h = 1$, $b = 19$, $K = 50$, $c_d = 5$, $c_r = 20$, and the size of the discount opportunities follows a delayed geometric distribution with parameter $p = 0.5$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the replenishment activities of a large number of shippers. In some environments, the arrival of discount opportunities may be correlated with demand conditions. For example, periods of low demand may prompt suppliers to offer discounts, implying a positive dependence between demand and discount arrivals. Accordingly, our results should be interpreted with caution in settings where such dependence is pronounced.

Discount opportunity sizes. Our analytical results assume that the sizes of discount opportunities are independent across periods. In Section S.3.4, we provide details of a numerical experiment in which we assume that the discount opportunity sizes are serially correlated. This reflects a situation with systematic demand fluctuations that influence the size of the opportunities offered. Our results reveal that the optimal order quantity then additionally depends on the last observed discount opportunity size. If opportunities are correlated, the state space thus becomes three-dimensional. If the expected size of the next discount opportunity is large, we obtain: (1) a lower reorder point s; (2) a lower can-order point c; and (3) higher threshold values $\{e_{s+1}, e_{s+2}, \ldots, e_c\}$. In summary, buyers will place fewer regular orders and will be less likely to use small discount opportunities if the expected size of the next discount opportunity is large. The opposite effect is observed when the expected size of the next discount opportunity is small.

4. How to set the discount price?

Our analysis so far shows that the buyer's response to a given discount offer depends on the discount price, its size, and the inventory level. While a higher discount on the regular order cost increases the buyer's probability of accepting the offer, it also diminishes the supplier's margin on discount sales and may lower revenue from the regular sales channel. In this section, we discuss the trade-offs that come with setting the discount price.

We express the supplier's average profit rate per period by $\Pi(k) = (k - c_d)\xi_d(k) + (K - c_r)\xi_r(k)$, where k is the discount price set by the supplier, K is the regular order cost, and c_d and c_r are the execution costs incurred by the supplier when the buyer places a discount or regular order, respectively. We let $\xi_d(k)$ and $\xi_r(k)$ denote the average rates at which the buyer places discount and regular orders, respectively, under the optimal replenishment policy at discounted order cost k. As we show in Theorem 1, for a certain k, the buyer decides to use a discount offer or place a regular order based on its current inventory level and the size of the discount opportunity. Given the complex nature of the buyer's optimal replenishment policy, it is hard, if not impossible, to analytically express the order rates $\xi_d(k)$ and $\xi_r(k)$. Still, we can numerically analyze their behavior and prove some monotonic properties to obtain insight into the optimal discount price.

Fig. 5 illustrates the dynamics resulting from setting the discount price k for a specific numerical setting. The gray dots display the observed numerical values. As these observations do not appear smooth due to the discretization of the buyer's problem, we fit functions to highlight trends (full colored lines). The supplier's profit rate (green line in Fig. 5) consists of profit from regular orders (purple line) and discounted orders (pink line). These components are driven by the respective ordering rates for discounted (blue line) and regular (orange line) orders.

The supplier's profit depends on k through three factors: (1) the margin $(k - c_d)$ on discount orders; (2) the rate at which discount orders are used $\xi_d(k)$, which decreases monotonically in k (see Appendix B for a proof); and (3) the rate at which regular orders are placed $\xi_r(k)$. A natural intuition is that $\xi_r(k)$ should increase monotonically in k, as higher discount prices make discounted purchasing less attractive

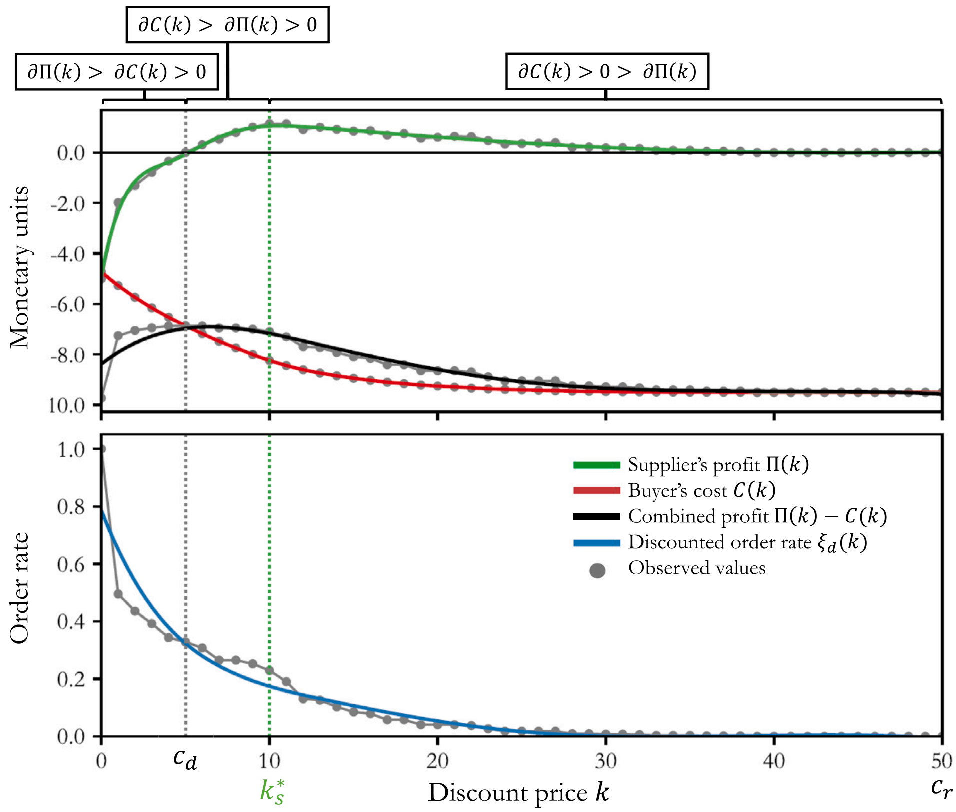

Fig. 6. A higher discount price k increases the buyer's inventory and order costs (red) but reduces the buyer's acceptance rate of discounted orders (blue). The supplier's profit-maximizing (green) discount price, k_s^*, leads to a low acceptance rate of discounted orders. When $c_r = K$, we can show that the discount price that maximizes the combined buyer's and supplier's profit (black) is equal to the execution cost of discounted orders c_d. This can only be attained through coordination or redistribution of the buyer's savings. The setting considered has parameters $\lambda = 1$, $\mu = 1$, $h = 1$, $b = 19$, $K = 50$, $c_d = 5$, $c_r = K = 50$, and the size of the discount opportunities follows a delayed geometric distribution with parameter $p = 0.5$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and should shift demand toward regular orders. In Fig. 5, however, we observe that this intuition does not always hold: increasing the discount price can also reduce the rate of regular sales, leading to a drop in profit from regular orders. The reason is that higher discount prices may induce changes in the buyer's optimal inventory policy that reduce overall ordering activity and increase inventory holdings, rather than reallocating purchases from discounted to regular orders.

The supplier's profit maximizing k, indicated by k_s^* in Fig. 5, optimally balances $(k - c_d)$, $\xi_d(k)$, and $\xi_r(k)$. Setting the discount price at k_s^* yields limited discount sales, as discounts are rarely used when the discount price k is too high relative to the regular order cost K. The discount acceptance rate could be increased by reducing the discount price. Although this reduces the supplier's profit rate, it leads to cost savings for the buyer, whose costs increase monotonically with the discount price k (see Appendix B for the proof).

Increasing the discount price k does not lead to proportional changes in the supplier's profit and the buyer's costs. An increase in k raises the buyer's total cost, whereas the supplier's profit does not necessarily increase at the same rate, or even increase at all. This asymmetry arises because the buyer optimizes its replenishment policy for each given value of k. As a result, changes in k may alter the composition of inventory and the balance between regular and discounted ordering costs.

The way in which the supplier's profit and the buyer's cost vary with k depends on the level of k itself, and a general characterization is not possible. One exception occurs when $c_r = K$, in which case the supplier's profit simplifies to $\Pi(k) = (k - c_d)\xi_d(k)$. This profit function can be interpreted as one in which the supplier is responsible only for discounted orders, while regular orders are fulfilled by a different supplier. In Fig. 6, we plot the supplier's profit $\Pi(k)$ (green line), the buyer's cost $C(k)$ (red line), and the discounted order rate $\xi_d(k)$ (blue line) when $c_r = K$. Note that the buyer's costs are expressed as negative values. We additionally plot the combined profit of the supplier and the buyer, expressed by $\Pi(k) - C(k)$ (black line).

When $c_r = K$, we can characterize the relation between the buyer's cost and the supplier's profit in three zones, illustrated in Fig. 6. In the first zone, where $0 \le k \le c_d$, increasing the discount price improves the supplier's profit rate more than it inflates the buyer's cost, i.e., $\partial\Pi(k) > \partial C(k) > 0$. In the second zone, where $c_d < k \le k_s^*$, increasing the discount price still improves the supplier's profit rate, but the buyer's cost rate escalates more rapidly, i.e., $\partial C(k) > \partial\Pi(k) > 0$. Finally, when $k > k_s^*$, increasing the discount price reduces the supplier's profit rate and augments the buyer's cost, i.e., $\partial C(k) > 0 > \partial\Pi(k)$. We visualize these dynamics in black through the combined supplier's profit and buyer's cost rate, expressed by $\Pi(k) - C(k)$. In Appendix B, we prove that the combined profit is maximized at a discount price $k = c_d$, which we formalize as follows:

Proposition 2. *If $c_r = K$, the combined profit of the supplier and the buyer $\Pi(k) - C(k)$ is maximized when the discount price equals the execution cost of discount offers, c_d.*

These results reveal that when the buyer uses discount offers, it reduces sourcing costs (at the supplier's expense) and additionally benefits from reduced inventory costs. However, as these inventory cost savings do not transfer to the supplier, coordination between the buyer and the supplier may be required to achieve the system-wide optimal solution. For instance, if the buyer shares its inventory and sourcing cost savings due to higher use of discount orders, the supplier can be incentivized to price discount orders below k_s^*. As such, additional buyer-to-supplier compensation can increase discount transactions while benefiting both parties. Similar incentive misalignments and coordination failures have been studied in the supply chain literature, where mechanisms such as revenue-sharing contracts are employed to align decentralized decisions with system-wide objectives and restore efficiency (e.g., Cachon (2003) and Cachon and Lariviere (2005)).

5. Conclusion

We study a setting where a supplier occasionally offers a buyer a discount on the fixed order cost. The quantity that can be ordered at a discount is restricted by a random capacity constraint. While random discount opportunities and capacity constraints have been studied in isolation, we address both complexities simultaneously. We first analyze the buyer's cost-minimizing response to capacitated discount offers. Understanding this behavior allows us to analyze the effect of setting the discount price.

We prove that the buyer's optimal policy to place regular and discount orders resembles a can-order policy (characterized by a unique reorder, can-order, and order-up-to level) but additionally features inventory level-specific thresholds. When the inventory level drops to the reorder point, it is optimal to place an order to replenish up to the order-up-to level. A discount offer is used only if the inventory level is at or below the can-order level *and* if the discount opportunity's capacity exceeds the inventory level-specific threshold. Then, we conjecture it is optimal to replenish up to the order-up-to level or as close as possible given the discount's capacity constraint.

As the thresholds complicate the optimal policy — limiting its practical applicability — we introduce three heuristic decision rules inspired by the optimal policy's structure. We provide an analytical bound on their worst-case optimality gap in settings with an infinitely high backlog cost rate. In a numerical study, we compare their performance and quantify the dominance of the four-parameter (s, c, m, S)-policy, that extends the (s, c, S) can-order policy with a minimum threshold parameter m. The parameter m ensures that an order is placed only when its size exceeds m units. Across 360 numerical settings, the (s, c, m, S)-policy achieves an average optimality gap of 1.7%, though at a higher computational cost than the simpler three-parameter heuristics. The (s, m, S)- and mod(s, c, S)-policies exhibit higher average optimality gaps at 5.2% and 7.8%, respectively. While the former extends the (s, S)-policy with a minimum threshold parameter m, the latter constrains the orders of a regular can-order policy by the capacity limit. Our results underscore the critical role of the minimum order parameter m, indicating that discount offers should only be accepted when sufficient volume can be ordered.

With the buyer's optimal response to discount opportunities in mind, we study the supplier's decision to set the discount price. A low discount price increases the probability that the buyer will accept the offer, but erodes the supplier's margin on discount sales and may reduce revenue from regular sales. We find that the profit-maximizing discount that balances this trade-off leads to limited discount sales, below the level that would be optimal for the combined profit of the buyer and supplier. An additional buyer-to-supplier compensation is necessary to incentivize the supplier to set a lower discount price which would benefit both parties.

CRediT authorship contribution statement

Bram J. De Moor: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Robert N. Boute:** Writing – review & editing, Writing – original draft, Supervision, Conceptualization. **Stefan Creemers:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Joren Gijsbrechts:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Proofs of Section 3.2

Proof of Lemma 1(a). We must show that $v^*(x, r) \geq v^*(x + \epsilon, r)$ holds for every state (x, r), with $\epsilon \in \mathbb{N}$. Let q_ϵ^* be the optimal action in state $(x + \epsilon, r)$, meaning that it is the action that minimizes Eq. (4):

$$v^*(x+\epsilon, r) = C((x+\epsilon, r), q_\epsilon^*) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+\epsilon+q_\epsilon^*-1, 0) + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+\epsilon+q_\epsilon^*, r') \right],$$

which is always lower than or equal to the future discounted costs of taking any other action $q \neq q_\epsilon^*$. Let q^* be the optimal action in state (x, r), and note that $q^* - \epsilon$ is not per se optimal in state $(x + \epsilon, r)$. By considering the (sub-optimal) action $q^* - \epsilon$ in state $(x + \epsilon, r)$ and exploiting the fact that q^* is optimal in (x, r), we show that $v^*(x, r) \geq v^*(x + \epsilon, r)$ as follows:

$$\begin{aligned} v^*(x+\epsilon, r) &\overset{(i)}{\leq} C((x+\epsilon, r), q^*-\epsilon) \\ &\quad + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+\epsilon+q^*-\epsilon-1, 0) \right. \\ &\quad \left. + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+\epsilon+q^*-\epsilon, r') \right] \\ &\overset{(ii)}{=} C((x+\epsilon, r), q^*-\epsilon) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+q^*-1, 0) \right. \\ &\quad \left. + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+q^*, r') \right] \\ &\overset{(iii)}{\leq} C((x, r), q^*) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+q^*-1, 0) \right. \\ &\quad \left. + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+q^*, r') \right] \\ &\overset{(iv)}{=} v^*(x, r). \end{aligned}$$

(i) follows from the fact that action $q^* - \epsilon$ is not per se optimal in state $(x + \epsilon, r)$. With some algebra, (ii) follows. (iii) stems from the fact that $C((x + \epsilon, r), q^* - \epsilon) \leq C((x, r), q^*)$, through Eq. (1). As q^* is the optimal order in state (x, r), (iv) follows. We thus find that $v^*(x + \epsilon, r) \leq v^*(x, r)$ for each state (x, r) with $\epsilon \in \mathbb{N}$. $\square$

Proof of Lemma 1(b). Consider two states (x, r) and $(x, r+\epsilon)$ with $\epsilon \in \mathbb{N}$. We show that $v^*(x, r) \geq v^*(x, r + \epsilon)\ \forall\ x$. Let q^* be the optimal order quantity in state (x, r). Using the same logic that was used to prove Lemma 1(a), we evaluate the cost of taking action q^* in state $(x, r + \epsilon)$, which is not per se optimal. Then:

$$\begin{aligned} v^*(x, r+\epsilon) &\overset{(i)}{\leq} C((x, r+\epsilon), q^*) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+q^*-1, 0) \right. \\ &\quad \left. + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+q^*, r') \right] \\ &\overset{(ii)}{\leq} C((x, r), q^*) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x+q^*-1, 0) \right. \\ &\quad \left. + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x+q^*, r') \right]. \end{aligned}$$

(i) stems from the fact that q^* is not per se the optimal order quantity in state $(x, r+\epsilon)$. (ii) follows because $C((x, r+\epsilon), q^*) \leq C((x, r), q^*)$, through

Eq. (1). We thus find that $v^*(x, r+\epsilon) \leq v^*(x, r)$ for every state (x, r) with $\epsilon \in \mathbb{N}$. □

Proof of Lemma 2(a). $G(x)$ is strictly decreasing in the interval $x \in]-\infty, 0]$, because $L(x)$ is strictly decreasing for $x \leq 0$ as $b > 0$, and $v^*(x, r)$ is non-increasing in x for every r, per Lemma 1(a). □

Proof of Lemma 2(b). Let $S \geq 0$ be a global minimum of $G(x)$. We show that $G(x)$ is strictly increasing for $x \geq S+1$, which corresponds to showing that $\Delta_x G(x) > 0$ for $x \geq S+1$, with $\Delta_x G(x) = G(x+1) - G(x)$ the discrete derivative of $G(x)$. We analyze the discrete derivative as we assume that $x \in \mathbb{Z}$.

Because S is a global minimum of $G(x)$, we know that for $x \geq S$:

$$\begin{aligned} v^*(x, r) &= \min \left[\min_{y \leq x} G(y), \min_{x<y\leq x+r} [G(y) + k], \min_{x+r<y} [G(y) + K] \right] \\ &= G(S). \end{aligned}$$

That is, if we have an inventory level $x \geq S$, we stay put or discard inventory at zero cost to reach S, which is the inventory level with the lowest expected sum of discounted costs. By replacing $v^*(x, r) = G(S)$ for every state (x, r) with $x \geq S$, we can derive the following expression for $x \geq S+1$:

$$\begin{aligned} G(x) &= L(x) + \gamma \left[\frac{\lambda}{\lambda+\mu} v^*(x-1, 0) + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') v^*(x, r') \right] \\ &= L(x) + \gamma \left[\frac{\lambda}{\lambda+\mu} G(S) + \frac{\mu}{\lambda+\mu} \sum_{r'=1}^{r'=\bar{r}} \phi_R(r') G(S) \right] \\ &= L(x) + \gamma G(S). \end{aligned}$$

Then, for every $x \geq S+1$:

$$\begin{aligned} \Delta_x G(x) &= G(x+1) - G(x) \\ &= (L(x+1) + \gamma G(S)) - (L(x) + \gamma G(S)) \\ &= L(x+1) - L(x) \\ &= \frac{h}{\lambda+\mu+\alpha}. \end{aligned}$$

$h/(\lambda+\mu+\alpha) > 0$ because $h > 0$. Thus, $G(x)$ is strictly increasing for $x \geq S+1$, with S a global minimum. □

Proof of Lemma 2(c). Lemma 2(a) shows that global minima of $G(x)$ is found in the interval $x \in [0, \infty[$, because $G(x)$ is strictly decreasing in the interval $x \in]-\infty, 0]$. Lemma 2(b) shows that $G(x)$ will be strictly increasing in the interval $x \in [S+1, \infty[$, where $S \geq 0$ is a global minimum of $G(x)$. In Section S.1 of the online Supplementary Material, we demonstrate that if $v^*(S-1, 0) \neq \frac{h}{\lambda\alpha}(\lambda+\alpha+\lambda S)$, then $G(S) < G(S+1)$, such that $G(x)$ is increasing in the interval $x \in [S, \infty[$. In this case, we prove by contradiction that $G(x)$ has a unique global minimum in S.

Assume $G(x)$ has two minima S_1 and S_2 with $0 \leq S_1 < S_2$. Given that both points are minimizers, it must hold that $G(S_1) = G(S_2)$. Provided that $v^*(S_1-1, 0) \neq \frac{h}{\lambda\alpha}(\lambda+\alpha+\lambda S_1)$, $G(x)$ is strictly increasing in the interval $x \in [S_1, \infty[$. Therefore, for any point $S_2 > S_1$, it must hold that $G(S_2) > G(S_1)$, contradicting the statement above. Hence, there exists only one global minimum. □

Lemma A.1. *If it is optimal to place a strictly positive order $q > 0$ in state $(x, 0)$, then the optimal order in state $(x-\epsilon, 0)$ with $\epsilon \in \mathbb{N}$ is $q' \geq q$.*

Proof. Let $(x, 0)$ be a state in which it is optimal to place an order $q > 0$, meaning that $v^*(x, 0) = \min_{x+0<y}[G(y) + K] = G(x+q) + K$. We prove by contradiction that the optimal order q' in state $(x-\epsilon, 0)$, with $\epsilon \in \mathbb{N}$, will always be at least as high as the optimal order q in state $(x, 0)$, i.e., $q \leq q'$. Suppose therefore that $q' < q$. We distinguish between two cases:

1. If $q' \leq 0$, then $v^*(x-\epsilon, 0) = \min_{y \leq x-\epsilon} G(y) = G(x-\epsilon+q') < G(x-\epsilon+(q+\epsilon)) + K = G(x+q) + K = v^*(x, r)$.
2. If $0 < q' < q$, then $v^*(x-\epsilon, 0) = \min_{x-\epsilon+0<y}[G(y) + K] = G(x-\epsilon+q') + K < G(x-\epsilon+(q+\epsilon)) + K = G(x+q) + K = v^*(x, 0)$.

In both cases, $v^*(x-\epsilon, 0) < v^*(x, 0)$, which contradicts Lemma 1(a). Hence, if it is optimal to place an order $q > 0$ in state $(x, 0)$, the optimal order in state $(x-\epsilon, 0)$ with $\epsilon \in \mathbb{N}$ is $q' \geq q$. □

Lemma A.2. *If it is optimal to place a strictly positive order $q > 0$ in state (x, r) with $r \geq S - x$, where S is a global minimizer of $G(x)$, then the optimal order in state $(x-\epsilon, r')$ with $\epsilon \in \mathbb{N}$ and $r' \geq S - (x-\epsilon)$ is $q' \geq q$.*

Proof. Let (x, r), with $r \geq S - x$, be a state in which it is optimal to order $q > 0$ such that $v^*(x, r) = \min_{x<y\leq x+r}[G(y) + k] = G(S) + k$. Now consider state $(x-\epsilon, r')$, with $\epsilon \in \mathbb{N}$ and $r' \geq S - (x-\epsilon)$, such that $r' > r$, in which it is optimal to order q'. We prove by contradiction that $q' \geq q$. Suppose therefore that $q' < q$ We distinguish between two cases:

1. If $q' \leq 0$, then $v^*(x-\epsilon, r') = \min_{y \leq x-\epsilon} G(y) = G(x-\epsilon+q') < G(S) + k = v^*(x, r)$.
2. If $0 < q' \leq q$, then $v^*(x-\epsilon, r') = \min_{x-\epsilon<y\leq x+r'}[G(y) + k] = G(x-\epsilon+q') + k < G(S) + k = v^*(x, r)$.

As per Lemma 1, we know that $v^*(x, r) \leq v^*(x-\epsilon, r) \leq v^*(x-\epsilon, r')$ with $r' > r$. If $q' < q$, however, we find that $v^*(x-\epsilon, r') < v^*(x, r)$, with $\epsilon \in \mathbb{N}$ and $r' > r$, which contradicts Lemma 1. Therefore, the optimal order quantity q' in state $(x-\epsilon, r')$ with $\epsilon \in \mathbb{N}$ and $r' \geq S - (x-\epsilon)$ will always be equal or larger than the optimal order quantity q in state (x, r) with $r \geq S - x$. That is, if in state (x, r) the size of the discount opportunity r permits to replenish to S and it is optimal to do so, then it will also be optimal to order up to S in state (x', r') with $x' < x$ and $r' > r$ if the size of the discount opportunity r' permits to reach S. □

Lemma A.3. *If it is optimal to place an order $0 < q \leq r$ in state (x, r), then the optimal order in state $(x, r+\epsilon)$ with $\epsilon \in \mathbb{N}$ is $q' \geq q$.*

Proof. Let (x, r) be a state in which it is optimal to order $0 < q \leq r$, such that $v^*(x, r) = \min_{x<y\leq x+r} G(y) + k = G(x+q) + k$. We prove by contradiction that the optimal order q' in state $(x, r+\epsilon)$, with $\epsilon \in \mathbb{N}$, will always be at least as high as the optimal order q in state (x, r), i.e., $q \leq q'$. Suppose therefore that $q' < q$. We distinguish between two cases:

1. If $q' \leq 0$, then $v^*(x, r+\epsilon) = \min_{y \leq x} G(y) = G(x+q') < G(x+q) + k = v^*(x, r)$.
2. If $0 < q' \leq q$, then $v^*(x, r+\epsilon) = \min_{x<y\leq x+r+\epsilon} G(y) + k = G(x+q') + k < G(x+q) + k = v^*(x, r)$.

In both cases, $v^*(x, r+\epsilon) < v^*(x, r)$, which contradicts Lemma 1(b). Hence, the optimal action in state $(x, r+\epsilon)$ orders $q' \geq q$ if the optimal action in state (x, r) orders $q \leq r$. Thus, given an inventory level x, the optimal order quantity is monotonically increasing in r. □

Proof of Theorem 1. As $\lim_{x\to-\infty} G(x) = \infty > G(S) + K$, where S is a global minimizer of $G(x)$, see Lemma 2(c), there exists a point at which placing an order and incurring the order cost K is optimal. Lemma A.1 shows that this reorder point s is unique, by proving that the optimal order quantity in state $(x, 0)$ is monotonically increasing for decreasing x, i.e., if it is optimal to order at inventory level x, it is also optimal to order at inventory level $x' < x$. When $(x, 0)$ with $x \leq s$, an order is placed to replenish to S. Similarly, as $\lim_{x\to-\infty} G(x) = \infty > G(S) + k$, there exists a point at which using a discount opportunity and incurring the order cost k is optimal, provided that the size of the discount opportunity r is large enough. Lemma A.2 shows that this can-order point c is unique. In states (x, r) with $s < x \leq c$, an order is placed if r is larger than a certain inventory-level dependent threshold e_x, which follows from Lemma A.3. In states (x, r) with $x \leq c$ and $S - x \leq r$, an order is placed to replenish to S, at which $G(x)$ is minimal. □

Appendix B. Proofs of Section 4

Lemma B.1. *The buyer's optimal discount opportunity usage rate $\xi_d(k)$ is monotonically decreasing in k when all other parameters (λ, μ, h, b, p, and K) are kept constant.*

Proof. Let $C(k)$ be the buyer's minimum long-run average cost per period in function of k when all other parameters (λ, μ, h, b, p, and K) are kept constant. Define $C(k_1|\pi^*_{k_1})$ as the buyer's cost when $k = k_1$ and the optimal policy $\pi^*_{k_1}$ is followed, such that $C(k_1) = C(k_1|\pi^*_{k_1})$. Similarly, we introduce $\xi_d(\pi^*_{k_1})$ as the buyer's discount opportunity usage rate when following the optimal policy $\pi^*_{k_1}$ when $k = k_1$, such that $\xi_d(k_1) = \xi_d(\pi^*_{k_1})$. Let $C_{IK}(k_1|\pi^*_{k_1})$ be the buyer's inventory and regular order costs when following policy $\pi^*_{k_1}$ such that $C(k_1) = C(k_1|\pi^*_{k_1}) = C_{IK}(k_1|\pi^*_{k_1}) + k_1\xi_d(\pi^*_{k_1})$.

We consider two settings with different discount prices k_1 and k_2 with $k_1 < k_2$, ceteris paribus. Assume that $\xi_d(k_1) = \xi_d(\pi^*_{k_1}) < \xi_d(\pi^*_{k_2}) = \xi_d(k_2)$. We show that such a situation can never occur, such that $\xi_d(k)$ is monotonically decreasing in k.

$\pi^*_{k_2}$ is the buyer's optimal policy when $k = k_2$ which implies that the average cost per period when following policy $\pi^*_{k_2}$ when $k = k_2$ can never be higher than when following another policy, say $\pi^*_{k_1}$:

$$
\begin{aligned}
& C(k_2) = C(k_2|\pi^*_{k_2}) \le C(k_2|\pi^*_{k_1}) \\
\Leftrightarrow\quad & C_{IK}(k_2|\pi^*_{k_2}) + k_2\xi_d(\pi^*_{k_2}) \le C_{IK}(k_2|\pi^*_{k_1}) + k_2\xi_d(\pi^*_{k_1}) \\
\Leftrightarrow\quad & k_2\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big) \le C_{IK}(k_2|\pi^*_{k_1}) - C_{IK}(k_2|\pi^*_{k_2}),
\end{aligned}
\tag{10}
$$

which means that it can only be optimal to increase the discount opportunity usage rate from $\xi_d(\pi^*_{k_1})$ to $\xi_d(\pi^*_{k_2})$ when k increases from k_1 to k_2 if the resulting increase in discount opportunity costs is less than or equal to the savings in inventory and regular order costs.

Similarly, $\pi^*_{k_1}$ is the buyer's optimal policy when $k = k_1$ which implies that the average cost per period when following policy $\pi^*_{k_1}$ when $k = k_1$ can never be higher than when following another policy, say $\pi^*_{k_2}$:

$$
\begin{aligned}
& C(k_1) = C(k_1|\pi^*_{k_1}) \le C(k_1|\pi^*_{k_2}) \\
\Leftrightarrow\quad & C_{IK}(k_1|\pi^*_{k_1}) + k_1\xi_d(\pi^*_{k_1}) \le C_{IK}(k_1|\pi^*_{k_2}) + k_1\xi_d(\pi^*_{k_2}) \\
\Leftrightarrow\quad & C_{IK}(k_1|\pi^*_{k_1}) - C_{IK}(k_1|\pi^*_{k_2}) \le k_1\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big).
\end{aligned}
$$

Because $k_1 < k_2$, we find:

$$
C_{IK}(k_1|\pi^*_{k_1}) - C_{IK}(k_1|\pi^*_{k_2}) \le k_1\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big) < k_2\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big). \tag{11}
$$

Note that the inventory and regular order costs do not depend on k, they only depend on the specific policy that the buyer follows. This means that:

$$
C_{IK}(k_1|\pi^*_{k_1}) - C_{IK}(k_1|\pi^*_{k_2}) = C_{IK}(k_2|\pi^*_{k_1}) - C_{IK}(k_2|\pi^*_{k_2}). \tag{12}
$$

Combining Eqs. (10), (11), and (12), we thus find that:

$$
\begin{aligned}
& k_2\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big) \le C_{IK}(k_2|\pi^*_{k_1}) - C_{IK}(k_2|\pi^*_{k_2}) \\
& = C_{IK}(k_1|\pi^*_{k_1}) - C_{IK}(k_1|\pi^*_{k_2}) < k_2\big(\xi_d(\pi^*_{k_2}) - \xi_d(\pi^*_{k_1})\big),
\end{aligned}
$$

which is impossible. Therefore, there exists no $k_1 < k_2$ for which $\xi_d(k_1) < \xi_d(k_2)$. Thus, $\xi_d(k)$ is monotonically decreasing in k. □

Lemma B.2. *The buyer's optimal long-run cost rate $C(k)$ is monotonically increasing in k when all other parameters (λ, μ, h, b, p, and K) are kept constant.*

Proof. Let $C(k)$ be the buyer's minimum long-run average cost per period in function of k when all other parameters (λ, μ, h, b, p, and K) are kept constant. Define $C(k_1|\pi^*_{k_1})$ as the buyer's cost when $k = k_1$ and the optimal policy $\pi^*_{k_1}$ is followed, such that $C(k_1) = C(k_1|\pi^*_{k_1})$. We consider two settings with different discount prices k_1 and k_2 with $k_1 < k_2$, ceteris paribus. Firstly, note that $C(k_1|\pi^*_{k_2}) \le C(k_2|\pi^*_{k_2})$, because in both cases we follow the exact same policy $\pi^*_{k_2}$, but the cost of discount opportunities is lower in the setting with $k = k_1$. Secondly, $C(k_1|\pi^*_{k_1}) \le C(k_1|\pi^*_{k_2})$, because the optimal policy yields the lowest average cost per period by definition. Combining both inequalities, we find $C(k_1) = C(k_1|\pi^*_{k_1}) \le C(k_1|\pi^*_{k_2}) \le C(k_2|\pi^*_{k_2}) = C(k_2)$. Thus, $C(k_1) \le C(k_2)$ for $k_1 < k_2$ which means that $C(k)$ is monotonically increasing in k. □

Proof of Proposition 2. Let $C(k)$ be the buyer's minimum long-run average cost per period in function of k when all other parameters (λ, μ, h, b, p, and K) are kept constant. Define $C(k_1|\pi^*_{k_1})$ as the buyer's cost when $k = k_1$ and the optimal policy $\pi^*_{k_1}$ is followed, such that $C(k_1) = C(k_1|\pi^*_{k_1})$. Similarly, we introduce $\xi_d(\pi^*_{k_1})$ as the buyer's discount opportunity usage rate when following the optimal policy $\pi^*_{k_1}$ when $k = k_1$, such that $\xi_d(k_1) = \xi_d(\pi^*_{k_1})$. Let $C_{IK}(k_1|\pi^*_{k_1})$ be the buyer's inventory and regular order costs when following policy $\pi^*_{k_1}$ such that $C(k_1) = C(k_1|\pi^*_{k_1}) = C_{IK}(k_1|\pi^*_{k_1}) + k_1\xi_d(\pi^*_{k_1})$.

Let $\Pi(k)$ be the supplier's profit when $c_r = K$, such that $\Pi(k) = (k - c_d)\xi_d(k)$. We show that $\Pi(k) - C(k)$ is maximized at $k = c_d$. To do so, we demonstrate that $\Pi(k) - C(k)$ is monotonically increasing for $k \le c_d$ and monotonically decreasing for $k > c_d$. We start by showing that $\Pi(k) - C(k)$ is monotonically increasing in $k \le c_d$. For this purpose, consider two settings with different discount prices k_1 and k_2 with $k_1 < k_2 \le c_d$, ceteris paribus. We show that $\big(\Pi(k_2) - \Pi(k_1)\big) - \big(C(k_2) - C(k_1)\big) \ge 0$.

Consider:

$$
\begin{aligned}
& C(k_2) - C(k_1) \overset{(i)}{=} C(k_2|\pi^*_{k_2}) - C(k_1|\pi^*_{k_1}) \\
\Leftrightarrow\quad & C(k_2) - C(k_1) \overset{(ii)}{\le} C(k_2|\pi^*_{k_1}) - C(k_1|\pi^*_{k_1}) \\
\Leftrightarrow\quad & C(k_2) - C(k_1) \overset{(iii)}{\le} C_{IK}(k_2|\pi^*_{k_1}) + k_2\xi_d(\pi^*_{k_1}) - C_{IK}(k_1|\pi^*_{k_1}) - k_1\xi_d(\pi^*_{k_1}) \\
\Leftrightarrow\quad & C(k_2) - C(k_1) \overset{(iv)}{\le} (k_2 - k_1)\xi_d(\pi^*_{k_1}).
\end{aligned}
$$

(i) follows from the definition of $C(k)$; (ii) follows because following a sub-optimal policy $\pi^*_{k_1}$ when $k = k_2$ can never yield a lower cost than the optimal policy $\pi^*_{k_2}$; (iii) follows from the definition of $C(k|\pi^*_k)$; (iv) follows because the inventory and regular order costs do not change in k when following the same policy $\pi^*_{k_1}$, i.e., $C_{IK}(k_2|\pi^*_{k_1}) = C_{IK}(k_1|\pi^*_{k_1})$.

Next, we analyze:

$$
\begin{aligned}
& \Pi(k_2) - \Pi(k_1) \overset{(i)}{=} (k_2 - c_d)\xi_d(\pi^*_{k_2}) - (k_1 - c_d)\xi_d(\pi^*_{k_1}) \\
\Leftrightarrow\quad & \Pi(k_2) - \Pi(k_1) \overset{(ii)}{=} k_2\xi_d(\pi^*_{k_2}) - k_1\xi_d(\pi^*_{k_1}) + c_d\big(\xi_d(\pi^*_{k_1}) - \xi_d(\pi^*_{k_2})\big) \\
\Leftrightarrow\quad & \Pi(k_2) - \Pi(k_1) \overset{(iii)}{\ge} k_2\xi_d(\pi^*_{k_2}) - k_1\xi_d(\pi^*_{k_1}) + k_2\big(\xi_d(\pi^*_{k_1}) - \xi_d(\pi^*_{k_2})\big) \\
\Leftrightarrow\quad & \Pi(k_2) - \Pi(k_1) \overset{(iv)}{\ge} (k_2 - k_1)\xi_d(\pi^*_{k_1}).
\end{aligned}
$$

(i) follows from the definition of $\Pi(k)$; (ii) follows directly by algebraic manipulation; (iii) follows because $k_2 \le c_d$ and $\xi_d(\pi^*_{k_1}) - \xi_d(\pi^*_{k_2}) \ge 0$, per Lemma B.1; (iv) follows directly by algebraic manipulation.

Thus, because $C(k_2) - C(k_1) \le (k_2 - k_1)\xi_d(\pi^*_{k_1})$ and $\Pi(k_2) - \Pi(k_1) \ge (k_2 - k_1)\xi_d(\pi^*_{k_1})$ for $k_1 < k_2 \le c_d$, we find that $\big(\Pi(k_2) - \Pi(k_1)\big) - \big(C(k_2) - C(k_1)\big) \ge 0$ for $k_1 < k_2 \le c_d$. The combined profit of the buyer and the supplier is therefore monotonically increasing in $k \le c_d$.

We proceed with showing that $\Pi(k) - C(k)$ is monotonically decreasing for $k > c_d$ in a similar fashion. For this purpose, consider two settings with different discount prices k_3 and k_4 with $c_d < k_3 < k_4$. We show that $\big(\Pi(k_4) - \Pi(k_3)\big) - \big(C(k_4) - C(k_3)\big) \le 0$.

Consider:

$$
C(k_4) - C(k_3) \overset{(i)}{=} C(k_4|\pi^*_{k_4}) - C(k_3|\pi^*_{k_3})
$$

$$\Leftrightarrow \quad C(k_4) - C(k_3) \overset{(ii)}{\geq} C(k_4|\pi^*_{k_4}) - C(k_3|\pi^*_{k_4})$$

$$\Leftrightarrow \quad C(k_4) - C(k_3) \overset{(iii)}{\geq} C_{IK}(k_4|\pi^*_{k_4}) + k_4\xi_d(\pi^*_{k_4}) - C_{IK}(k_3|\pi^*_{k_4}) - k_3\xi_d(\pi^*_{k_4})$$

$$\Leftrightarrow \quad C(k_4) - C(k_3) \overset{(iv)}{\geq} (k_4 - k_3)\xi_d(\pi^*_{k_4}).$$

(i) follows from the definition of $C(k)$; (ii) follows because $C(k_3|\pi^*_{k_3}) \leq C(k_3|\pi^*_{k_4})$; (iii) follows from the definition of $C(k|\pi^*_k)$; (iv) follows because the inventory and regular order costs do not change in k when following the same policy $\pi^*_{k_4}$, i.e., $C_{IK}(k_4|\pi^*_{k_4}) = C_{IK}(k_3|\pi^*_{k_4})$.

Next, we analyze:

$$\Pi(k_4) - \Pi(k_3) \overset{(i)}{=} (k_4 - c_d)\xi_d(\pi^*_{k_4}) - (k_3 - c_d)\xi_d(\pi^*_{k_3})$$

$$\Leftrightarrow \quad \Pi(k_4) - \Pi(k_3) \overset{(ii)}{\leq} (k_4 - c_d)\xi_d(\pi^*_{k_4}) - (k_3 - c_d)\xi_d(\pi^*_{k_4})$$

$$\Leftrightarrow \quad \Pi(k_4) - \Pi(k_3) \overset{(iii)}{\leq} (k_4 - k_3)\xi_d(\pi^*_{k_4}).$$

(i) follows from the definition of $\Pi(k)$; (ii) follows because $k_3 - c_d > 0$ and $\xi_d(\pi^*_{k_3}) \geq \xi_d(\pi^*_{k_4})$, per Lemma B.1; and (iii) follows directly.

Thus, because $C(k_4) - C(k_3) \geq (k_4 - k_3)\xi_d(\pi^*_{k_4})$ and $\Pi(k_4) - \Pi(k_3) \leq (k_4 - k_3)\xi_d(\pi^*_{k_4})$ for $c_d < k_3 < k_4$, we find that $\big(\Pi(k_4) - \Pi(k_3)\big) - \big(C(k_4) - C(k_3)\big) \leq 0$ for $c_d < k_3 < k_4$. The combined profit of the buyer and the supplier is therefore monotonically decreasing for $k > c_d$.

As we have shown that the combined profit of the buyer and the supplier is monotonically increasing for $k \leq c_d$ and monotonically decreasing for $k > c_d$ when $c_r = K$, we find that the combined profit $\Pi(k) - C(k)$ is maximized at $k = c_d$. $\square$

Appendix C. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ejor.2026.04.016.

References

Balintfy, J. L. (1964). On a basic class of multi-item inventory problems. *Management Science*, *10*(2), 287–297.

Beckmann, M. (1961). An inventory model for arbitrary interval and quantity distributions of demand. *Management Science*, *8*(1), 35–57.

Cachon, G. P. (2003). Supply chain coordination with contracts. *Handbooks in Operations Research and Management Science*, *11*, 227–339.

Cachon, G. P., & Lariviere, M. A. (2005). Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Science*, *51*(1), 30–44.

Chaouch, B. A. (2007). Inventory control and periodic price discounting campaigns. *Naval Research Logistics*, *54*(1), 94–108.

Federgruen, A., & Zipkin, P. (1986a). An inventory model with limited production capacity and uncertain demands I. The average-cost criterion. *Mathematics of Operations Research*, *11*(2), 193–207.

Federgruen, A., & Zipkin, P. (1986b). An inventory model with limited production capacity and uncertain demands II. The discounted-cost criterion. *Mathematics of Operations Research*, *11*(2), 208–215.

Friend, J. (1960). Stock control with random opportunities for replenishment. *Journal of the Operational Research Society*, *11*(3), 130–136.

Gallego, G., & Scheller-Wolf, A. (2000). Capacitated inventory problems with fixed order costs: Some optimal policy structure. *European Journal of Operational Research*, *126*(3), 603–613.

Goh, M., & Sharafali, M. (2002). Price-dependent inventory models with discount offers at random times. *Production and Operations Management*, *11*(2), 139–156.

Moinzadeh, K. (1997). Replenishment and stocking policies for inventory systems with random deal offerings. *Management Science*, *43*(3), 334–342.

Perera, S. C., & Sethi, S. P. (2023a). A survey of stochastic inventory models with fixed costs: Optimality of (s, S) and (s, S)-type policies—Continuous-time case. *Production and Operations Management*, *32*(1), 154–169.

Perera, S. C., & Sethi, S. P. (2023b). A survey of stochastic inventory models with fixed costs: Optimality of (s, S) and (s, S)-type policies—Discrete-time case. *Production and Operations Management*, *32*(1), 131–153.

Puterman, M. L. (1994). *Markov decision processes: Discrete stochastic dynamic programming* (1st ed.). John Wiley & Sons Inc..

Rossi, R., Chen, Z., & Tarim, S. A. (2024). On the stochastic inventory problem under order capacity constraints. *European Journal of Operational Research*, *312*(2), 541–555.

Scarf, H. (1960). The optimality of (S, s) policies in the dynamic inventory problem. In K. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences* (pp. 196–202). Stanford, CA: Stanford University Press Stanford.

Shaoxiang, C. (2004). 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.

Shaoxiang, C., & Lambrecht, M. (1996). XY band and modified (s, S) policy. *Operations Research*, *44*(6), 1013–1019.

Silver, E. A. (1974). A control system for coordinated inventory replenishment. *International Journal of Production Research*, *12*, 647–671.

Tajbakhsh, M. M., Lee, C.-G., & Zolfaghari, S. (2011). An inventory model with random discount offerings. *Omega*, *39*(6), 710–718.

Wijngaard, J. (1972). *An inventory problem with constrained order capacity*: *TH-Report 72-WSK-63*, Eindhoven University of Technology.

Zheng, Y. S. (1994). Optimal control policy for stochastic inventory systems with Markovian discount opportunities. *Operations Research*, *42*(4), 721–738.

Zipkin, P. (2008). Old and new methods for lost-sales inventory systems. *Operations Research*, *56*(5), 1256–1263.