Received: 25 July 2021 | Revised: 19 July 2021 | Accepted: 24 July 2021
DOI: 10.1111/deci.12544

DECISION SCIENCES
A JOURNAL OF THE DECISION SCIENCES INSTITUTE
WILEY

ORIGINAL ARTICLE

Customers know best: Pricing policies for products with heterogeneous quality

Yalçın Akçay[1] | **Fikri Karaesmen**[2]

[1] Melbourne Business School, The University of Melbourne, Carlton, VIC, Australia

[2] Department of Industrial Engineering, Koç University, Istanbul, Turkey

Correspondence
Yalçın Akçay, Melbourne Business School, The University of Melbourne, 200 Leicester Street, Carlton, VIC 3053, Australia.
Email: y.akcay@mbs.edu

Abstract
This article studies the pricing problem of a seller given an initial inventory of products with heterogeneous quality, facing uncertain customer arrivals over a finite selling season. We consider various regimes depending on whether the seller inspects the inventory to assess the quality levels of the products, and whether customers examine the inventory themselves and pick their specific item of choice among the available products. We formulate the problem under each regime as a stochastic optimization model which maximizes the seller's expected profits, capturing the salient problem features such as stochastic customer arrivals, customers' choice behavior, and uncertain product qualities. As obtaining closed-form solutions or structural properties for the optimal prices is quite difficult, we explore the full information solution to the problem as an upper bound, as well as solution approaches that approximate some key problem characteristics. Finally, we substantiate our results through an extensive numerical study, focusing on the performance of the proposed pricing policies and approximations.

KEYWORDS
dynamic pricing, revenue management, random product quality

1 | INTRODUCTION

Product quality is an essential component of competitive advantage and a key driver of customer satisfaction. It is generally viewed as an exact variable that is measured as a function of some technical characteristics of the product (Garvin, 1984). For example, quality of timber is determined by factors such as knot size and frequency, grain uniformity, discoloration, and wane in the wood (Sunley, 1963). Even after timber products are categorized (graded) at the supply, such that products in the same category are deemed to have the same functional use, significant heterogeneity in quality still remains due to differences in the quantity of levels of such key attributes (Leffler, 1982). So, *how should a seller price an inventory of such heterogeneous quality products*? *How do key operating factors impact the seller's profit*? To address these questions, we introduce a modeling framework that takes into account whether the seller inspects the individual products for quality before pricing and/or the customers are allowed to examine the products before making purchase decisions, and develop a stochastic optimization formulation for each model. Our primary focus is to derive analytical and managerial insights. In what follows, we discuss the salient features of our problem and our modeling assumptions.

Although we use the timber industry as a motivating example, our model and subsequent analysis apply to more general settings in which a seller makes pricing decisions for a given initial inventory of products with heterogeneous quality, facing random demand over a finite selling season. We should also note that "quality," in a broader sense, refers to any product attribute for which all customers have a *common* preference for. The seller initially knows the quality distribution, but does not specifically know the true quality of each product in their inventory, which could potentially be unraveled through inspection of each individual item, at a cost. We assume that the quality levels of the individual products are *fixed* and do *not* deteriorate through time due to environmental conditions (such as meat, dairy products, fresh produce, flowers, pharmaceutical products, chemicals, and blood). In such an environment, the seller sets *one* price at each point in time, and this price holds for all products regardless of their quality. In other words, the seller does *not* pursue quality-based price discrimination to sell higher quality products at higher prices. On the other hand, a customer makes the purchase decisions based on a linear

TABLE 1 Summary of regimes

		Customers	
		Don't Examine	***Examine***
Seller	*Doesn't Inspect*	$\mathcal{R}_{\emptyset}$	$\mathcal{R}_{\mathbb{E}}$
	Inspects	$\mathcal{R}_{\mathbb{I}}$	$\mathcal{R}_{\mathbb{IE}}$

utility function that is increasing in the product's quality and decreasing in its price. As all products carry the *same* price, it is in the customers' best interest to look for (assuming that they are allowed to do so) a high-quality item among the available products, so as to maximize their surplus. To this end, we assume that the "*customers know best*"—they can basically distinguish low- versus high-quality products, and accordingly purchase (at most) one item each (we discuss the details of our choice model in Section 2). This is a stylized model that addresses the flip side of the simplest conventional models where customers may have heterogeneous valuations (preferences) for perfectly identical products; in our model, the products are similar but heterogeneous in quality and customers agree upon a common value for each product.

Based on the above description of our problem, we consider a framework with four possible regimes (summarized in Table 1) depending on whether the seller chooses to *inspect* the inventory to assess the individual quality levels of the products, and whether the seller lets customers *examine* the inventory themselves and pick their specific item of choice among the available products:

$\mathcal{R}_{\emptyset}$: Seller does *not* inspect inventory and does *not* allow customers to examine products

$\mathcal{R}_{\mathbb{E}}$: Seller does *not* inspect inventory but allows customers to examine products

$\mathcal{R}_{\mathbb{I}}$: Seller inspects inventory but does *not* allow customers to examine products

$\mathcal{R}_{\mathbb{IE}}$: Seller inspects inventory and allows customers to examine products

Under $\mathcal{R}_{\emptyset}$ and $\mathcal{R}_{\mathbb{E}}$, the seller determines the optimal price based on the distribution of product quality, whereas under $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$, he makes the pricing decision with the complete information about specific quality levels of the products. Under $\mathcal{R}_{\mathbb{E}}$ and $\mathcal{R}_{\mathbb{IE}}$, the seller sets the optimal price taking into account that customers make purchase decisions based on their *realized* utility after observing quality, whereas they would make these decisions based on their *expected* utility facing random quality under $\mathcal{R}_{\emptyset}$ and $\mathcal{R}_{\mathbb{I}}$. Specifically, we consider two pricing policies under $\mathcal{R}_{\mathbb{E}}$: (i) Under *static pricing* policy, the seller does *not* change the price over the course of the selling season; (ii) under *price updating* (PU) policy, the seller *revises* the initial price at a pre-specified time. We should also note that under $\mathcal{R}_{\emptyset}$, $\mathcal{R}_{\mathbb{I}}$, and $\mathcal{R}_{\mathbb{IE}}$, as well as under $\mathcal{R}_{\mathbb{E}}$ with static pricing, the product price remains the same throughout the selling horizon.

We introduce the stochastic optimization formulations that maximize the seller's expected profit for each regime in Sections 3 and 4.1. Our main focus in this article is on the seller's optimal pricing policy under $\mathcal{R}_{\mathbb{E}}$. We primarily use the pricing policies under the other regimes, $\mathcal{R}_{\emptyset}$, $\mathcal{R}_{\mathbb{I}}$, and $\mathcal{R}_{\mathbb{IE}}$, as benchmarks. For example, the gap between the optimal profits under $\mathcal{R}_{\mathbb{E}}$ and $\mathcal{R}_{\emptyset}$ reveals the marginal value of the customers' ability to examine products before making their purchase decisions. In fact, our numerical experiments show that the seller benefits more from this ability when inventory has overall lower and more variable quality, and customer demand is relatively scarce (we present the details of our computational study in Section 5.1 and discuss further insights regarding the impact of key problem parameters on the seller's optimal profit under $\mathcal{R}_{\emptyset}$, $\mathcal{R}_{\mathbb{I}}$, $\mathcal{R}_{\mathbb{E}}$, and $\mathcal{R}_{\mathbb{IE}}$).

We should mention that obtaining closed-form solutions for the optimal prices is quite an arduous task. In our problem, not only are customer arrivals stochastic but also each product in inventory is essentially different from one another through uncertain quality. This calls for the analysis to deal with multiple order statistics and manipulations of their joint distribution functions. To address these challenges, we consider three approximations—a *fixed proportions* approximation scheme that overlooks the stochastic nature of product quality and assumes that a fixed (deterministic) proportion of available inventory rather than a random proportion is above a given quality level; a *mean-value-equivalence* scheme that in addition to the fixed proportions assumption ignores demand uncertainty; and an *unlimited demand* scheme that further assumes that demand is not a constraint for the seller. We computationally assess the performance of these approximations relative to the optimal pricing strategies under regime $\mathcal{R}_{\mathbb{E}}$ in Section 5.2, and show that the fixed proportions policy is the most effective one among the three. Moreover, the performance of the approximations improves as inventory becomes more abundant in quantity and less variable in quality.

Pricing of seasonal products has been widely studied in extant literature. We refer the reader to Phillips (2005) and Talluri and Van Ryzin (2005) for overviews of the basic pricing models, and to Elmaghraby and Keskinocak (2003), Bitran and Caldentey (2003), and Chen and Simchi-Levi (2012) for reviews of this rich literature. An implicit assumption in the typical line of research is that all items in the seller's inventory for a particular product are identical (homogenous inventory), and that customers consequently have the same preference for each unit. On the other hand, our work assumes that items are variants of a single product, differentiated in their random quality levels (heterogeneous inventory). In this sense, the pricing component of our work is related to the pricing of vertically differentiated product inventories that arise in assortment planning. For instance, Pan and Honhon (2012) assume that customers differ in their valuations of quality, and consequently, develop several efficient algorithms to optimally price the products (jointly with optimal assortment decisions). Akcay et al. (2010) study the dynamic pricing problem for an inventory of vertically differentiated products over a finite selling season and derive structural properties of the optimal pricing policy. All of these models inherently assume that the quality levels of the products in inventory are predetermined (perhaps at a strategic level) and known by the firm, whereas a distinct feature of our model is the uncertainty in the quality levels of the products. In particular, we focus on how to price the products under different levels of information on their quality. To this end, our main contribution in this article is to develop and analyze a series of stylistic models that enhance our understanding of pricing an inventory that is heterogeneous in quality. The models present unique challenges both on the customer choice side and the supply side, and to perform a complete unified analysis, we simplify some of the features while respecting the essence of the main issues.

2 | CHOICE MODEL

We consider a seller selling products to price- and quality-sensitive customers over a single selling season. The seller purchases an initial stock of n units from a supplier before demand is realized and sells these products during the season. Any unsold product by the end of the selling season is valueless (no salvage value). We assume that the quality levels of products are identically and independently distributed continuous random variables denoted by Q (for any randomly selected item) with cumulative distribution $F_Q(\cdot)$ and density $f_Q(\cdot)$ over support $[0, \infty)$. For notational convenience, we index the products in decreasing order of their quality and let $i, i \in \{1, 2, \dots, n\}$, indicate the product with the ith *highest* quality product among all n products in the supplier's initial inventory. Subsequently, Q_i designate the random quality level of product i, that is, $Q_1 \geq Q_2 \geq \cdots \geq Q_n$. By definition, Q_i is also known as the ith-*order statistic* of Q. We denote by U_i a customer's utility for product i with quality level Q_i at price p, and express it as

$$U_i = Q_i - p. \tag{1}$$

Note that all customers have the same utility for product i, that is, customers are *homogenous* in their valuations of the products. Unless otherwise stated, the seller does not set different prices for different quality items and p is the same for all products in (1). We assume that the *no-purchase* option, denoted by 0 for notational convenience, is always available, and its utility is 0, that is, $U_0 = 0$. Therefore, a customer would never purchase a product with a negative utility. If the customer examines product i, they can resolve the uncertainty associated with random quality Q_i. We assume that the cost of acquiring and processing product information (search cost) is negligible, and therefore, a new customer examines *all* products available at the time of arrival, identifies the highest quality product, and purchases it if the utility is nonnegative. This choice model is commonly referred to as a *pure characteristics*demand model (Berry & Pakes, 2007) and has been widely used to describe *vertical* demand (e.g., Bhargava & Choudhary, 2001; Bresnahan, 1987; Tirole, 1988; Wauthy, 1996).

3 | PRICING WITHOUT INSPECTION

In this section, we formulate the seller's pricing problem when the seller does *not* inspect products to resolve their exact quality levels. Section 3.1 describes our model under regime $\mathcal{R}_\emptyset$, whereas Section 3.2 presents two pricing policies under $\mathcal{R}_\mathbb{E}$. Throughout the article, we focus on the revenue optimization problem of the seller. Unsold items have no salvage value or disposal cost.

3.1 | Pricing under $\mathcal{R}_\emptyset$

Customers do *not* examine the products before they make their purchase decisions under regime $\mathcal{R}_\emptyset$. This is a common situation for products sold over the Internet or through mail-order catalogs, where customers cannot experience the traditional *touch-and-feel* shopping, hence diminishing their ability to assess the quality of products considerably. As the seller also does not inspect the available inventory, we assume that he randomly assigns products to customers. Hence, it is likely that a customer ends up with an "unsatisfactory" product, that is, the customer observes a negative *ex-post* utility because the true quality of the product turns out to be less than its price. Selling with "money-back guarantee" is a widely used approach to improve customer satisfaction in today's hypercompetitive business environment. Accordingly, the seller allows customers to return products that do not meet their expectations as a result of "poor" quality back to the seller for a full refund. We assume that: (i) all returned products are discarded by the seller, that is, they do not go back to the seller's inventory; (ii) customers do not seek to *replace* an unsatisfactory product with a new product; and (iii) customers keep a product as long as their utility is nonnegative and do not return a product back to the seller hoping to *exchange* it with a higher quality item (clearly, if such exchanges were allowed by the seller, our results would be quite different). We refer the reader to Fruchter and Gerstner (1999), Su (2009), and Akcay et al. (2013) for detailed studies of pricing and inventory models under money-back guarantees.

Under $\mathcal{R}_\emptyset$ with full refunds, the customer's utility is $(q - p)$, given that the product's realized quality q exceeds p. On the other hand, the customer's utility is 0 if the product's quality is below p (because the customer returns the product for full refund). The customer's expected utility is then:

$$\mathrm{E}[U] = \mathrm{E}[(Q - p)^+], \tag{2}$$

where $z^+ \equiv \max(z, 0)$. The customer's expected utility in (2) is nonnegative. Therefore, customers would always be willing to purchase the product.[1] In turn, the seller's expected profit from a unit of product sold would be $p(1 - F_Q(p))$ because $pF_Q(p)$ of the selling price p is refunded.

Let the total number of customer *arrivals* over the course of the selling season be a discrete random variable, denoted by A with cumulative distribution $F_A(\cdot)$. We assume that A is independent of the quality of the seller's inventory n and the selling price p. On the other hand, the final sales would ultimately be a function of A, Q, p, and n.

We write the seller's expected profit under $\mathcal{R}_\emptyset$ as

$$\Pi_\emptyset(p) = p(1 - F_Q(p))\mathrm{E}[\min\{n, A\}],$$

and express the optimal expected profit as

$$\Pi^*_\emptyset = \max_{p\geq 0} \Pi_\emptyset(p). \tag{3}$$

Note that the optimal expected profit is scaled by $\mathrm{E}[\min\{n, A\}]$, but the optimal price is independent of n and A, because Q and A are independent. This leads to the following result:

Proposition 1. *If $p(1 - F_Q(p))$ is a concave function of p or if Q has an increasing generalized failure rate (IGFR), then the optimal price $p^*_\emptyset$ satisfies:*

$$p^*_\emptyset = \frac{1 - F_Q(p^*_\emptyset)}{f_Q(p^*_\emptyset)}. \tag{4}$$

We refer the reader to Lariviere and Porteus (2001) for a formal definition of IGFR. The proof of Proposition 1 is straightforward by differentiating $p(1 - F_Q(p))$ with respect to p. Concavity or the IGFR condition in Lariviere and Porteus (2001) ensures that the first-order condition guarantees optimality (see Online Appendix A for the formal proof).

Proposition 1 is the analog of the results in the classical unconstrained price optimization problem for customers with random valuations for homogeneous items where the optimal price only depends on the customer valuations but not on the market size (potential demand). In the classical problem, the optimal price is typically nondecreasing in the available supply but here the optimal price does not depend on n because customers have homogeneous valuations for items of same quality. On the other hand, the expected optimal revenue is nondecreasing in n.

In our case, the randomness comes from the product quality levels and the optimal price only depends on Q. To compare the impact of random quality on optimal price and expected profit, let us define Q^X and Q^Y as random variables associated with two different quality-level distributions, $F_{Q^X}(\cdot)$ and $F_{Q^Y}(\cdot)$. To understand how the optimal price depends on Q, we can remark that if a given product quality level Q^X stochastically dominates another quality level Q^Y in terms of the hazard rate order, then the optimal price for Q^X is higher. As the hazard rate order carries a similar meaning to the regular stochastic order (in fact, it implies the stochastic order), this can loosely be interpreted as stochastically higher quality in the hazard order sense, leading to higher optimal prices and higher expected revenues.

3.2 | Pricing under $\mathcal{R}_\mathbb{E}$

The seller could inspect their inventory so as to resolve the uncertainty around the quality levels of products. In some cases though, such an effort can be a significant cost factor for the seller. A study by the Juran Institute reveals that costs associated with measuring, evaluating or auditing products (or services) to assure conformance to quality standards amount to 10–30% of sales or 25–40% of operating costs for most U.S. companies (De Feo, 2005). There are also relevant situations where the seller simply does not have the expertise to assess the quality of the inventory. This may happen when the seller is an occasional trader of items such as objects of art or crafts. Further, customers may also have more expertise in artisanal objects than occasional traders. This particular behavior of customers could potentially save the seller the inspection effort and enable higher profits.

In light of the above arguments, suppose that the seller does not perform any inspection when the products are received from the supplier, hence not knowing the exact (realized) quality levels of the products in the inventory. On the other hand, the customer examines the inventory, and chooses the highest quality product available at the time of their arrival. The objective of the seller is to maximize their expected profit over the finite selling season, given a particular pricing policy and *without* full information about product quality. In Section 3.2.1, we first present our pricing policy in which the seller does *not* change the price over the course of the selling season under $\mathcal{R}_\mathbb{E}$. We then extend this model in Section 3.2.2, to discuss a policy in which the seller *updates* the initial price at a prespecified time.

3.2.1 | Static pricing policy

Under the static pricing policy in regime $\mathcal{R}_\mathbb{E}$, the seller sets the price at the start of the selling season based on the prior information about product quality and taking into account that customers will be making choices after examining the inventory. Let us define $N(p)$ as the number of available items whose quality levels are greater than or equal to the selling price p. In other words, $N(p)$ corresponds to the *maximum* number of products that can be sold at price p given *infinite* demand. Accordingly, $N(p)$ would be equal to a particular value i, if the quality level of the ith highest quality product exceeds p, that is, $Q_i \geq p$, whereas the next highest quality product does not, that is, $Q_{i+1} < p$. We can then express the

[1] We assume that if the utility of a product is 0, customers prefer purchasing the product over the no-purchase option.

probability mass function of $N(p)$ as:

$$\mathcal{P}(N(p)=i)=\mathcal{P}\left(Q_i \geq p, Q_{i+1}<p\right)=\binom{n}{i} F_Q(p)^{n-i}\left(1-F_Q(p)\right)^i. \quad (5)$$

The random variable $N(p)$ follows a *binomial* distribution with number of trials n and probability of "success" $1-F_Q(p)$.

Let us next define $S(p)$ as the random number of products sold to customers at price p. Clearly, total sales is bounded by customer arrivals as well as by $N(p)$. Then, we have $S(p)=\min\{N(p), A\}$. We write the seller's expected profit as a function of price p under the static pricing policy in $\mathcal{R}_{\mathbb{E}}$ as

$$\Pi_{\mathbb{E}}^{\mathrm{SP}}(p)=p\mathrm{E}[S(p)], \quad (6)$$

where

$$\mathrm{E}[S(p)]=\sum_{i=0}^{n}\sum_{j=0}^{\infty}\min(i,j)\mathcal{P}(N(p)=i)\mathcal{P}(A=j).$$

We then express the optimal expected profit as

$$\Pi_{\mathbb{E}}^{\mathrm{SP}*}(p)=\max_{p\geq 0}\Pi_{\mathbb{E}}^{\mathrm{SP}}(p). \quad (7)$$

Unfortunately, $S(p)$ is a complicated random variable and explicit optimization to obtain a closed-form solution appears difficult. On the other hand, we can perform some comparative statics on the expected profit. We can establish that the expected optimal revenue increases when the quality level Q stochastically increases, and that the expected optimal revenue increases as the number of arrivals stochastically increases. These results are formalized and justified in Online Appendix B.

If $\Pi_{\mathbb{E}}^{\mathrm{SP}}(p)$ is concave in p, we can express the optimal price under the static pricing policy as the solution of:

$$p_{\mathbb{E}}^{\mathrm{SP}*}=\frac{-\mathrm{E}[S(p_{\mathbb{E}}^{\mathrm{SP}*})]}{\frac{d(\mathrm{E}[S(p_{\mathbb{E}}^{\mathrm{SP}*})])}{dp}}.$$

Apart from special cases, the optimal price from the above expression has to be computed using numerical integration and derivation or simulation-based methods that use derivative estimators. To gain further insight into the optimal prices, we next explore two approximations that relax part of our model assumptions.

First, we relax the assumption that items in inventory are discrete, and instead, let the inventory be *infinitesimally* divisible. Then, the number of products whose quality ratings are above a certain level is a *fixed proportion* (FP) of the initial inventory (see Honhon & Seshadri, 2013; Hopp & Xu, 2008, as similar approximations in extant literature). Specifically, we approximate the quantity of products that can *potentially* be sold at price p as $n(1-F_Q(p))$. We can then express the optimal expected profit of the seller under the FP approximation to the SP policy as follows:

$$\Pi_{\mathbb{E},\mathrm{FP}}^{\mathrm{SP}*}=\max_p p\mathrm{E}\left[\min\left\{A, n(1-F_Q(p))\right\}\right]. \quad (8)$$

If $p(1-F_Q(p))$ is concave in p and since the minimum of two concave functions is concave and expected value preserves concavity, one can establish the concavity of $\Pi_{\mathbb{E},\mathrm{FP}}^{\mathrm{SP}}(p)$ in p. Then, the optimal price can uniquely be determined under this approximate policy. Unfortunately, there appears to be no general analytical characterization of the optimal price because both demand arrival randomness and quality randomness play a role.

To make further headway, we next introduce a second relaxation, and replace the random number of arrivals with its mean value, in addition to the assumptions of FP approximation. Effectively, this captures a situation in which the market size is fixed. Then, the optimal profit of the seller under the mean-value approximation (MVE) to the SP policy is given by

$$\Pi_{\mathbb{E},\mathrm{MVE}}^{\mathrm{SP}*}=\max_p p\left[\min\left\{\mu_A, n(1-F_Q(p))\right\}\right], \quad (9)$$

where $\mu_A := \mathrm{E}[A]$. In this case, we can establish that the optimal price is characterized by one of the two following conditions given in the next proposition.

Proposition 2. *If $p(1-F_Q(p))$ is a concave function of p or Q has an* IGFR*, the optimal price under the MVE approximation is given by:*

$$\tilde{p}_{\mathbb{E},\mathrm{mve}}^{\mathrm{sp}}=\max\{p_u, p_c\}, \quad (10)$$

where p_u is the solution of $p_u=(1-F_Q(p_u))/f_Q(p_u)$ that is the unconstrained solution and

$$p_c=F_Q^{-1}\left(1-\frac{\mu_A}{n}\right),$$

which is the constrained (market satiating price) under the MVE approximation.

The proof follows from the (quasi)-concavity of the unconstrained problem under the IGFR condition. If the unconstrained optimal price leads to a supply that is above the average demand, then it is optimal to increase the price so that the demand constraint becomes binding. This is the reverse of a corresponding result for pricing under random customer valuations and limited supply (see Phillips, 2005, Chapter 5) that establishes that the optimal price under limited supply is greater than or equal to the optimal price under ample supply. For our case, we remark from Proposition 2 that the optimal price under limited demand must be greater than or equal to the unconstrained price that is the optimal price given in Proposition 1. Accordingly, the comparison between $\tilde{p}_{\mathbb{E},\mathrm{MVE}}^{\mathrm{SP}}$ and $p_{\emptyset}^{*}$ reveals the price *premium* when the seller lets customers examine the products before making their purchase

TABLE 2 Summary of all approximations

Policy	Quality	Demand
Optimal	**Random**	**Random**
fp	Deterministic	Random
mve	Deterministic	Deterministic
ul	Deterministic	Unlimited

decisions (assuming that the MVE approximation is relatively accurate).

In contexts where customers arrive in abundance, and hence do not impose a constraint on the seller's price problem, we can take the above approximations one step further. In the unlimited demand (UL) approximation, we first adopt the FP approximation for the available inventory, and subsequently remove the demand constraint from the problem, that is, we assume that the demand is infinite. Then, the optimal profit of the seller under the UL approximation to the SP policy is given by

$$\Pi^{\text{SP}*}_{\mathbb{E},\text{UL}} = \max_{p} p\left[n(1 - F_Q(p))\right]. \tag{11}$$

Note that under the UL assumption, FP and MVE approximations become identical. Consequently, the optimal price of the seller under the UL approximation to the SP policy satisfies

$$p^{\text{SP}}_{\mathbb{E},\text{UL}} = \tilde{p}^{\text{SP}}_{\mathbb{E},\text{MVE}} = \frac{1 - F_Q\left(p^{\text{SP}}_{\mathbb{E},\text{UL}}\right)}{f_Q\left(p^{\text{SP}}_{\mathbb{E},\text{UL}}\right)},$$

which coincides with the optimal price from Proposition 1.

Table 2 provides a conceptual summary and comparison of all approximate policies that are proposed in this section.

3.2.2 | Price updating policy

When the seller does not know the exact quality levels of the products in the inventory, they can observe the realized sales that is driven by customers' quality assessments. A smart seller should then be able to infer quality information pertaining to its remaining inventory, though acquiring and exploiting this information is not trivial. As inventory is depleted, adjusting the price dynamically enables the seller to better exploit the current supply-demand mismatch (see Gallego & van Ryzin, 1994, and Bitran & Mondschein, 1997). In this particular setting where the higher quality items are sold first, price updating has the potential to be an even stronger lever, because not only the remaining inventory level but also the remaining quality-level distribution provides valuable information. From a theoretical perspective, the seller could continuously update the product price so as to promptly adjust to the ever-changing operational conditions. Nevertheless, because of the high costs associated with frequent price changes, the number of times that the seller can adjust the price is typically limited in practice (Netessine, 2006). Moreover, differentiated and directly unobservable quality make such pricing decisions in our context more challenging. Therefore, we consider a PU policy in which the seller first sets an initial price for all products at the start of the selling season, and then adjusts this price *once*, before the end of the selling season. By doing so, the seller divides the selling season into two periods, and makes a pricing decision at the start of each period. Note that the seller *only* observes realized sales at the end of the first period and makes the price update decision for the second period *without* the full information about customer arrivals. In other words, the seller does *not* monitor the actual arrivals but rather only keeps record of depleted inventory (hence observes only sales). Let A_1 and A_2 be random variables that correspond to the number of arrivals in periods 1 and 2. We assume that A_1 and A_2 are independent (this assumption can be relaxed but at the expense of further complicating an already challenging model).

To find the optimal pricing policy, we formulate a two-period dynamic program starting with the second period. $\pi^{\text{PU}}_{2}(n-i, p_1)$ denotes the expected *optimal* profit-to-go in the beginning of period 2, given that $n-i$ items are remaining in inventory (i.e., i items were sold at price p_1 in the first period. The price p_2 is action to take in the second period. We have

$$\pi^{\text{PU}}_{\mathbb{E},2}(n-i, p_1) = \max_{p_2} \sum_{j=0}^{n-i} \mathcal{P}\{S_2(p_2) = j | S_1(p_1) = i\} j p_2 \quad \text{for } i < n$$

and the boundary condition $\pi^{\text{PU}}_{\mathbb{E},2}(0, p_1) = 0$.

For the first period, we define $\Pi^{\text{PU}}_{\mathbb{E}}(p_1)$ as the expected profit-to-go obtained by optimally pricing in period 2 with an initial inventory level of n

$$\Pi^{\text{PU}}_{\mathbb{E}}(p_1) = \sum_{i=0}^{n} \mathcal{P}\{S_1(p_1) = i)\}\left(ip_1 + \pi^{\text{PU}}_{\mathbb{E},2}(n-i, p_1)\right), \tag{12}$$

and express the optimal expected profit-to-go as follows:

$$\Pi^{\text{PU}*}_{\mathbb{E}} = \max_{p_1 \geq 0} \Pi^{\text{PU}}_{\mathbb{E}}(p_1). \tag{13}$$

Although the formulation of the dynamic program appears simple enough, the computation of the conditional probabilities $\mathcal{P}\{S_2(p_2) = j | S_1(p_1) = i)\}$ is rather involved. Note that $S_1(p_1)$ depends on the number of arrivals in period 1, A_1, which are not observable. If A_1 were observable and equal to k, then we could have expressed this conditional probability as: $\mathcal{P}\{S_2(p_2) = j | S_1(p_1) = i, A_1 = k)\}$. When A_1 is not observable, however, we need to rely on the total law of probability to write:

$$\mathcal{P}\{S_2(p_2) = j | S_1(p_1) = i)\} = \sum_{k=0}^{\infty} \mathcal{P}\{S_2(p_2) = j | S_1(p_1) = i, A_1 = k)\} P(A_1 = k),$$

which only requires information regarding the probability distribution of customer arrivals, and not the specific realizations of the random arrivals. We present the details of these conditional probability calculations, which are fairly tedious, with illustrative examples, in Online Appendix B.

The formulation in (12) clearly does not lend itself to further analysis to obtain closed-form expressions for the optimal prices under the PU policy. Similar to those in Section 3.2.1, fixed proportions and MVEs can be developed for this case as well. The most explicit approximation combines these approximations and is described in Online Appendix B.

Finally, we should note that the idea behind the PU policy is akin to price skimming and dynamic pricing in revenue management (e.g., Besanko & Winston, 1990; Talluri & Van Ryzin, 2005). Therefore, there is no reason to restrict the number of price updates to *one* from a "theoretical" perspective. Nevertheless, with each price update, the seller also needs to dynamically revise the quality distribution of the remaining inventory. Clearly, as the number of price updates increases, the required analysis is expected to become more complicated, even intractable. On the other hand, the PU policy delivers an actionable and scalable pricing policy that is relatively easy to understand, compute, and implement. Our numerical experiments indicate that if the seller updates the price of the remaining inventory more frequently, the expected revenue increases but the marginal benefit from each additional price update decreases. Although not surprising, our observations are analogous to that in Aviv and Pazgal (2008).

4 | PRICING WITH INSPECTION

In this section, we formulate the seller's pricing problem when the seller inspects all items in the inventory and has *full* information about the true quality of each product. We discuss our model under regimes $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$ in Section 4.1. We then present an *upper bound* on the expected profit of the seller in Section 4.2.

4.1 | Pricing under $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$

As customers do not examine the seller's inventory upon their arrival under $\mathcal{R}_{\mathbb{I}}$, customers make their purchase decisions solely based on their *prior* information regarding quality, as in regime $\mathcal{R}_{\emptyset}$. Assuming that the seller refunds all unsatisfied customers, the expected utility of a customer is always nonnegative, as given by (2). Unlike $\mathcal{R}_{\emptyset}$, however, the seller would *not* randomly assign *any* available product to a customer under $\mathcal{R}_{\mathbb{I}}$, but instead assigns only one of those products whose true quality level is above the price. In other words, the seller would not attempt to sell a product that would eventually be returned by the customer for a refund (because of a negative postpurchase utility). In contrast, under $\mathcal{R}_{\mathbb{IE}}$, a customer chooses the product with the highest quality among those whose true quality level exceeds the price, as customers have the opportunity to examine the products before purchase. Even though the set of products that could be sold under $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$ are identical, the specific assignment of products to customers could potentially be different. For instance, the first arriving customer would purchase the highest quality product (given that its true quality exceeds p) under $\mathcal{R}_{\mathbb{IE}}$, whereas the seller (randomly) assigns any one of the qualifying products to the customer under $\mathcal{R}_{\mathbb{I}}$. Nevertheless, the total number of items sold at price p is the same under both $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$ given the same demand. Therefore, the seller faces the exact same pricing problem in both cases.

Let $\mathbf{q} = (q_1, q_2, \ldots, q_n)$ be the *true* quality levels of the n products in the seller's initial inventory. As products are indexed in decreasing order of their quality levels, we have $q_1 \geq q_2 \geq \ldots \geq q_n$. Further, define

$$\delta_i(p) = \begin{cases} 1 & \text{if } q_i \geq p, \\ 0 & \text{if } q_i < p. \end{cases} \tag{14}$$

The indicator $\delta_i(p)$ indicates whether the product with quality q_i yields a nonnegative utility at price p. Then, the seller would solve the following problem to maximize the expected profit from an initial inventory with quality $\mathbf{q}$

$$\pi_{\mathbb{I}}(\mathbf{q}) = \pi_{\mathbb{IE}}(\mathbf{q}) = \max_{p \geq 0} \sum_{i=1}^{n} p\delta_i(p)\mathcal{P}(A \geq i) - n.c,$$

where c is the seller's unit cost of inspection. Note that the optimal price $p^*_{\mathbb{I}} = p^*_{\mathbb{IE}} \in \{q_1, \ldots, q_n\}$. In other words, the seller picks the optimal price from a discrete set of values (quality levels), and essentially establishes the set of products to offer to customers.

Taking an expectation over all possible realizations of the random quality vector $\mathbf{Q} = \{Q_1, Q_2, \ldots, Q_n\}$, we calculate the expected optimal profit under regimes $\mathcal{R}_{\mathbb{I}}$ and $\mathcal{R}_{\mathbb{IE}}$ as follows:

$$\Pi^*_{\mathbb{I}} = \Pi^*_{\mathbb{IE}} = \mathrm{E}[\pi_{\mathbb{I}}(\mathbf{Q})] = \mathrm{E}[\pi_{\mathbb{IE}}(\mathbf{Q})]. \tag{15}$$

4.2 | Complete price differentiation

Let us now assume that the seller can price each individual product *separately* given full information about quality. In terms of PU, this is equivalent to the seller continuously updating the price of the inventory after each sale so as to clear the highest quality product currently available. One could alternatively imagine that the seller sets a different price for each product in advance, as he knows their qualities. In this case, no price updating is required; the seller just has different sets prices for different products. Obviously, such a policy would be difficult to implement in practice but it does provide an upper bound on the performance of all dynamic updating policies.

Note that under this policy, the seller sets the price of the ith highest quality product (with quality level q_i) as $p_i = q_i$, for $i = 1, 2, \dots, n$, so as to maximize the total profit. Accordingly, all customers would have a zero net surplus (expected utility) for all products in the seller's inventory. Then, the expected profit of the seller from an initial inventory with quality $\mathbf{q}$ follows as

$$\bar{\pi}(\mathbf{q}) = \sum_{i=1}^{n} q_i \mathcal{P}(A \geq i) - nc.$$

The optimal expected profit under this policy with *individual* prices is then given by

$$\bar{\Pi}^* = \mathrm{E}[\bar{\pi}(\mathbf{Q})] = \sum_{i=1}^{n} \mathrm{E}[Q_i]\mathcal{P}(A \geq i) - nc, \tag{16}$$

where $\mathrm{E}[Q_i]$ denotes the expected value of ith-order statistic of the random variable Q. Note that $\bar{\Pi}^*$ is an upper bound on $\Pi^*_{\mathbb{I}} = \Pi^*_{\mathbb{IE}}$.

In the next proposition, we show how two different sets of initial inventories, whose product quality levels can be stochastically ordered, impact the expected profit of the seller under complete price differentiation.

Proposition 3. *Define Q^X and Q^Y as random variables associated with two different quality-level distributions, $F_{Q^X}(\cdot)$ and $F_{Q^Y}(\cdot)$, respectively. Under complete price differentiation,*

(i) if Q^X stochastically dominates Q^Y, then for the same inventory size, the optimal expected profit of Q^X is greater than or equal to that of Q^Y;
(ii) if Q^X and Q^Y are symmetrical with the same mean, and Q^Y is more peaked than Q^X around this mean, then for the same inventory size, the optimal expected profit of Q^X is greater than or equal to that of Q^Y.

The condition $Q^X \geq_{st} Q^Y$ implies that all order statistics are also stochastically ordered: $Q_i^X \geq_{st} Q_i^Y$ for all i (see David & Nagaraja, 2003). In return, this implies that $\mathrm{E}[Q_i^X] \geq \mathrm{E}[Q_i^Y]$ for all i. The result then follows from (16). As customers' willingness to pay increases with quality, the seller clearly has a higher pricing power with X than Y for all items. Consequently, the full information solution with X as initial stock exceeds what the customers would obtain with Y, as stated in part (i) of Proposition of 3.

Perhaps, the more interesting issue arises when sets X and Y contain products with the same average quality but differ in the *variability* of the quality levels. Such a comparison is nontrivial because increased variability in quality has a downside (i.e., more low-quality items) as well as an upside (more high-quality items that can be priced accordingly). A comparison of these two depends on all order statistics as expressed in (16) and requires a variability order called the "peakedness order" (see David & Nagaraja, 2003) from reliability theory. Note that the peakedness order is a measure of concentration around the mean. In particular, if Q^Y is more peaked than Q^X, then $Var(Q^Y) \leq Var(Q^X)$.

Part (ii) of Proposition 3 suggests that an initial inventory with more variable quality provides a more lucrative sales opportunity for the seller. In the following, we give an intuitive explanation for this result. First, if Q^Y is more peaked than Q^X, then it is more likely for the seller to have a high-quality (above-average-quality) product in X compared to Y. Accordingly, with X, the seller can generate a larger profit from products with higher than average quality levels. On the other hand, the seller is also more likely to have a low-quality (below average quality) product in X compared to Y. Hence, the seller can potentially collect higher profits with Y, by asking higher prices than what he could with X for these below average-quality products. However, customer arrivals over the selling season might be less than the initial inventory. Any leftover stock contains lower quality products as customers buy high-quality products first, resulting in the seller not being able to fully realize the potential profit from low-quality products. As a result, the nature of customer behavior along with uncertain demand enables the seller to achieve a larger profit with the initial stock X than with Y.

5 | NUMERICAL RESULTS

In this section, we report the results of our extensive numerical experiments to substantiate our analytical findings and complement them with new insights. Specifically, we address how key problem parameters affect the seller's optimal profit under the four regimes $\mathcal{R}_{\emptyset}$, $\mathcal{R}_{\mathbb{I}}$, $\mathcal{R}_{\mathbb{E}}$, and $\mathcal{R}_{\mathbb{IE}}$, as well as the upper bound (complete price differentiation), and the performance of the FP and MVE approximations of the SP and PU policies.

5.1 | Impact of problem parameters

We generate problems by systematically varying key parameters of our problem. We assume that customers arrive according to a Poisson process with mean arrival rate $\lambda = 10$ (homogenous). The seller has 10 items in its initial inventory ($n = 10$) at the start of the selling season, and quality levels of these items are uniformly distributed between $\underline{q}$ and $\bar{q}$, that is, $Q \sim U[\underline{q}, \bar{q}]$[2]. We vary the $\underline{q}$ and $\bar{q}$ values to reflect different initial inventory constructs. We also assume that the seller's unit purchasing cost and unit inspection cost of a product are both zero (without loss of generality). We should mention that, as optimal price expressions under each regime and the corresponding expected profit values are not in closed form in general, we used *simulation-based* stochastic optimization to solve the problem instance. Accordingly, we generated 1000

[2] We tested the robustness of our results assuming that quality follows a Beta distribution instead of uniform, and validated that our overall insights remain unchanged.

TABLE 3 Impact of *average* quality of the starting inventory ($n = 10$)

	$\underline{q}$	$\bar{q}$	$\Pi^*_{\emptyset}$	$\Pi^{sp*}_{\mathbb{E}}$	$\Pi^{pu*}_{\mathbb{E}}$	$\Pi^*_{\mathbb{IE}}$	$\bar{\Pi}^*$	$\Delta^{\mathbb{E},sp}_{\emptyset}$	$\Delta^{\mathbb{E},pu}_{\mathbb{E},sp}$	$\Delta^{\mathbb{IE}}_{\mathbb{E},sp}$	$\bar{\Delta}_{\mathbb{IE}}$
$\lambda = 5$	0.0	0.5	0.622	1.010	1.074	1.226	1.703	62.37%	6.34%	21.33%	38.94%
	0.1	0.6	0.896	1.373	1.449	1.623	2.201	53.21%	5.54%	18.23%	35.60%
	0.2	0.7	1.220	1.763	1.858	2.040	2.699	44.53%	5.39%	15.75%	32.27%
	0.3	0.8	1.593	2.173	2.277	2.472	3.196	36.44%	4.79%	13.75%	29.29%
	0.4	0.9	2.016	2.600	2.704	2.915	3.694	28.94%	4.00%	12.13%	26.74%
	0.5	1.0	2.489	3.037	3.135	3.366	4.192	22.00%	3.25%	10.84%	24.55%
$\lambda = 10$	0.0	0.5	1.094	1.223	1.517	1.516	2.364	11.84%	23.99%	23.94%	55.91%
	0.1	0.6	1.575	1.743	2.155	2.107	3.238	10.69%	23.63%	20.87%	53.71%
	0.2	0.7	2.143	2.340	2.848	2.769	4.113	9.15%	21.71%	18.34%	48.56%
	0.3	0.8	2.800	3.004	3.587	3.488	4.988	7.28%	19.43%	16.13%	43.01%
	0.4	0.9	3.543	3.726	4.359	4.250	5.863	5.15%	16.99%	14.08%	37.95%
	0.5	1.0	4.374	4.494	5.166	5.043	6.738	2.74%	14.94%	12.21%	33.61%
$\lambda = 15$	0.0	0.5	1.233	1.244	1.646	1.551	2.484	0.92%	32.30%	24.64%	60.20%
	0.1	0.6	1.775	1.792	2.372	2.184	3.471	0.95%	32.37%	21.88%	58.89%
	0.2	0.7	2.416	2.438	3.187	2.917	4.457	0.88%	30.73%	19.68%	52.78%
	0.3	0.8	3.156	3.181	4.072	3.735	5.443	0.77%	28.02%	17.45%	45.72%
	0.4	0.9	3.995	4.016	5.000	4.614	6.430	0.54%	24.49%	14.88%	39.35%
	0.5	1.0	4.932	4.937	5.953	5.531	7.416	0.10%	20.60%	12.03%	34.09%

random sample paths for Poisson arrivals over the course of the selling season, and 1000 random draws from uniformly distributed quality levels for the initial inventory.

Before proceeding to the detailed discussion of our numerical study, we first give a high-level summary of our results. For the 36 different problem scenarios we generated in this first part of our study, we find that the upper bound solution, proposed in Section 4.2, is, on average, 84.60% higher than the seller's optimal expected profit in regime $\mathcal{R}_{\emptyset}$. This gap represents the *joint* value of perfect information and price adjustment in this context —recall that the upper bound yields the seller's profit under full information and individualized pricing (maximum possible number of price points), whereas pricing in $\mathcal{R}_{\emptyset}$ gives us the profit facing completely random quality levels (no information) and a single price.

A comparison of expected profit levels in regimes $\mathcal{R}_{\mathbb{IE}}$ and the SP policy in $\mathcal{R}_{\mathbb{E}}$ reveals an interesting phenomenon—a considerable chunk of the aforementioned 84.60% gap is due to the seller's ability to adjust the price, and inspection appears to have only a limited effect. Specifically, the seller's average profit in regime $\mathcal{R}_{\mathbb{IE}}$ is merely 15.49% more than that under the SP policy in $\mathcal{R}_{\mathbb{E}}$. We observe that by leveraging the opportunity to adjust the price under PU (even if it is through a single price update) during the selling season, the seller can overcome the challenges associated with having to make pricing decisions with incomplete information; PU yields a comparable 15.75% higher profits than SP on average under $\mathcal{R}_{\mathbb{E}}$.

To understand how product quality affects policy performance, we consider two important aspects of product quality—average quality and variability of quality. For this purpose, we adjust the upper and lower limits of the uniform distribution, characterizing the quality levels of the products, to reflect an assortment of quality-related features regarding the initial inventory of the seller. We first vary the average of product quality levels, while keeping the range (variability) of the distribution at a constant level, specifically at 0.5. (i.e., the distribution of quality levels becomes more dominant in the stochastic order). As the average quality of the initial inventory improves, so does the customers' overall willingness to pay a premium for such more desirable products. Table 3 presents the impact of average quality on the seller's optimal expected profit under regimes $\mathcal{R}_{\emptyset}$, $\mathcal{R}_{\mathbb{E}}$, and $\mathcal{R}_{\mathbb{IE}}$ $(= \mathcal{R}_{\mathbb{I}})$ and the following metrics:

$$\Delta^{\mathbb{E},SP}_{\emptyset} = \frac{\Pi^{SP*}_{\mathbb{E}} - \Pi^*_{\emptyset}}{\Pi^*_{\emptyset}} \quad \text{(marginal value of } \textit{customers' examination}\text{)},$$

$$\Delta_{\mathbb{E},\,SP}{}^{\mathbb{E},PU} = \frac{\Pi^{PU*}_{\mathbb{E}} - \Pi^{SP*}_{\mathbb{E}}}{\Pi^{SP*}_{\mathbb{E}}} \quad \text{(marginal value of } \textit{singlepriceupdate}\text{)},$$

$$\Delta^{\mathbb{IE}}_{\mathbb{E},SP} = \frac{\Pi^*_{\mathbb{IE}} - \Pi^{SP*}_{\mathbb{E}}}{\Pi^{SP*}_{\mathbb{E}}} \quad \text{(marginal value of } \textit{theseller'sinspection}\text{)},$$

$$\bar{\Delta}_{\mathbb{IE}} = \frac{\bar{\Pi}^* - \Pi^*_{\mathbb{IE}}}{\Pi^*_{\mathbb{IE}}} \quad \text{(marginal value of } \textit{priceadjustment}\text{)}.$$

As expected, as average quality increases, so does the seller's expected profits (under all policies). More interestingly, the customers' ability to examine products before making the purchase decision is more valuable when the inventory is of lower quality and the customers arrive rarely. In contrast, the seller's ability to update the price is more beneficial when inventory is scarce relative to demand. The seller also

TABLE 4 Impact of *variability* in quality of the initial inventory ($n = 10$)

	$\underline{q}$	$\bar{q}$	$\Pi^*_{\emptyset}$	$\Pi^{sp*}_{\mathbb{E}}$	$\Pi^{pu*}_{\mathbb{E}}$	$\Pi^*_{\mathbb{IE}}$	$\bar{\Pi}^*$	$\Delta^{\mathbb{E},sp}_{\emptyset}$	$\Delta^{\mathbb{E},pu}_{\mathbb{E},sp}$	$\Delta^{\mathbb{IE}}_{\mathbb{E},sp}$	$\bar{\Delta}_{\mathbb{IE}}$
$\lambda = 5$	0.0	1.0	1.244	2.019	2.149	2.452	3.406	62.25%	6.41%	21.41%	38.94%
	0.1	0.9	1.260	1.975	2.081	2.354	3.223	56.77%	5.35%	19.18%	36.90%
	0.2	0.8	1.327	1.956	2.064	2.279	3.039	47.33%	5.52%	16.54%	33.35%
	0.3	0.7	1.524	1.993	2.077	2.243	2.856	30.76%	4.20%	12.51%	27.34%
	0.4	0.6	1.991	2.132	2.142	2.277	2.672	7.09%	0.46%	6.79%	17.36%
	0.5	0.5	2.488	2.488	2.488	2.488	2.488	0.00%	0.00%	0.00%	0.00%
$\lambda = 10$	0.0	1.0	2.187	2.448	3.033	3.032	4.727	11.9%	23.92%	23.88%	55.91%
	0.1	0.9	2.215	2.462	3.052	3.002	4.657	11.18%	23.98%	21.93%	55.12%
	0.2	0.8	2.333	2.560	3.132	3.049	4.586	9.73%	22.35%	19.12%	50.40%
	0.3	0.7	2.679	2.833	3.33	3.245	4.516	5.73%	17.55%	14.56%	39.15%
	0.4	0.6	3.500	3.500	3.754	3.692	4.445	0.00%	7.27%	5.51%	20.39%
	0.5	0.5	4.374	4.374	4.374	4.374	4.374	0.00%	0.00%	0.00%	0.00%
$\lambda = 15$	0.0	1.0	2.466	2.489	3.292	3.102	4.969	0.92%	32.3%	24.63%	60.20%
	0.1	0.9	2.497	2.521	3.345	3.096	4.961	0.97%	32.71%	22.80%	60.27%
	0.2	0.8	2.630	2.655	3.486	3.195	4.954	0.94%	31.31%	20.33%	55.06%
	0.3	0.7	3.021	3.039	3.812	3.512	4.946	0.62%	25.44%	15.56%	40.85%
	0.4	0.6	3.945	3.945	4.325	4.136	4.939	0.00%	9.63%	4.83%	19.43%
	0.5	0.5	4.931	4.931	4.931	4.931	4.931	0.00%	0.00%	0.00%	0.00%

benefits more from inspection when selling poor-quality products. We finally observe that the quality of the upper bound deteriorates as the average product quality decreases.

Next, in Table 4, we vary the standard deviation of the product quality levels in the base scenario while keeping the average quality for the initial inventory at a constant value of 0.5. W. Edward Deming, the eminent American statistician, famously claimed that "Variation is the enemy of quality."

In our particular problem context, variability in quality levels can intuitively be considered as a measure of the consistency of the quality *grade*. Intuitively, as the standard deviation of quality increases, it is more likely for a randomly selected item in the seller's initial inventory to have a relatively high-quality level. Consequently, the pool of high-quality products, which are clearly more attractive for customers, is expected to be larger for the more variable quality case compared to the less variable quality case for any given initial inventory level. Furthermore, the size of this pool also increases as the initial inventory of the seller increases. Therefore, although somewhat counterintuitive, an abundant initial inventory (relative to demand) consisting of highly variable quality levels presents a more desirable assortment of products for which customers potentially have a higher willingness to pay. At this point, the particular pricing policy adopted by the seller is ultimately accountable in translating this prospect into better returns. We observe that it might be possible for the seller to take advantage of the "enemy," as coined by Deming, and cash in on opportunities bestowed by variability through an intelligent pricing policy. Our results in Table 4 reveal that the seller's expected profits under $\mathcal{R}_{\emptyset}$ decrease with variability, whereas they increase under full information. The marginal value of customers' ability to examine the seller's inventory is more emphasized when quality variability is high and customer arrivals are infrequent. On the other hand, when customers arrive in abundance, we observe that the price update proves to be more beneficial for the seller. As with the customers' examining ability, the seller's inspection is more valuable when variability is high.

5.2 | Performance of approximate policies

In our next set of computations, we assess the performance of the FP, MVE, and UL approximations relative to the PUoptimal pricing strategies in regime $\mathcal{R}_{\mathbb{E}}$. We generate a wide range of problem instances from the base scenario, mainly focusing on the level of demand-supply mismatch and the variability in product quality. For this purpose, we fix $n = 40$ and let $\lambda \in \{5, 10, \ldots, 75, 80\}$. Moreover, we also vary $Q \sim U[\delta, 1-\delta]$ such that $\delta \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$. In all, we generate 96 problem instances.

Table 5 shows the overall average and the maximum optimality gap for each of the approximate policies. As one would expect, FP approximation proves to be the most effective. It is indeed near-optimal for the SP policy (even the maximum percentage gap between $\Pi^{SP*}_{\mathbb{E}}$ and $\Pi^{SP*}_{\mathbb{E},FP}$ is less than 1% over all the considered problem scenarios). Similarly, FP approximation to PU also provides the best performance with an average gap of 1.07%. However, the optimality gap can possibly exceed 10% as seen in Table 5. We observe that FP performance is followed by the MVE approximation, and

DECISION SCIENCES
A JOURNAL OF THE DECISION SCIENCES INSTITUTE

TABLE 5 Overall performance of approximate policies

	Average Gap		Maximum Gap	
	sp	**pu**	**sp**	**pu**
fp	0.13%	1.07%	0.83%	10.76%
mve	0.78%	1.78%	7.65%	17.26%
ul	5.79%	5.94%	40.19%	25.52%

TABLE 6 Average performance gap for the approximate policies under different levels of demand

	sp			**pu**		
Demand	**fp**	**mve**	**ul**	**fp**	**mve**	**ul**
Low						
$\lambda_t \in \{5, 10, 15, 20, 25\}$	0.22%	2.12%	16.19%	1.90%	3.94%	12.08%
Medium						
$\lambda_t \in \{30, 35, 40, 45, 50\}$	0.10%	0.27%	1.34%	1.12%	1.35%	4.06%
High						
$\lambda_t \in \{55, 60, 65, 70, 75, 80\}$	0.09%	0.09%	0.84%	0.33%	0.33%	2.37%

that completely ignoring demand as a constraint turns out to be an inadequate approach to the seller's problem.

To achieve a better understanding of the key operational factors that might contribute to the performance of the approximations, we next take a closer look at our numerical results and provide average optimality gaps as functions of demand level and quality variability in Tables 6 and 7. First, we observe that the degree of supply-demand mismatch plays a significant role in the performance of our approximations. When there is abundant demand relative to the on-hand inventory, all three approximations perform well. However, if there is an oversupply of products, the seller needs to capture as much of the intricate effects of demand as possible—the performance of the approximations deteriorate drastically as demand becomes more scarce over the selling season (see Table 6). This deterioration is more emphasized for MVE,

TABLE 7 Average performance gap for the approximate policies under different levels of quality variability

	sp			**pu**		
Quality variability	**fp**	**mve**	**ul**	**fp**	**mve**	**ul**
Low						
$\delta \in \{0.3, 0.4, 0.5\}$	0.08%	0.82%	5.22%	0.98%	1.73%	5.02%
High						
$\delta \in \{0.0, 0.1, 0.2\}$	0.18%	0.74%	6.37%	1.16%	1.82%	6.85%

which ignores demand stochasticity, and even more so for UL, which ignores demand completely.

On the other hand, the approximations generally perform better when there is decreased variability in the quality levels of the products. We should mention that all three approximations treat the amount of products that can be sold at a particular price as a deterministic quantity. Therefore, in Table 7, we see that the effect of quality variability on their respective performances are comparable. The significantly poor performance under the UL approximation is due to its omission of the effect of demand on prices.

Based on these observations, we conclude that the FP and MVE approximations can be effectively used by the seller when the demand (relative to supply) is large and the inventory does not exhibit significant variability in quality; otherwise, it is critical for the seller to adopt the exact optimal pricing policy.

6 | CONCLUDING REMARKS

In this article, we considered a seller with a given initial inventory of products with heterogeneous quality, and formulated the pricing problem when facing uncertain customer arrivals over a finite selling season. We studied the problem under various operating regimes and revealed the following managerial insights:

- The expected profit of the seller improves as the average quality of inventory increases. The seller's ability to adjust the price is more valuable when having to sell lower quality inventory.
- The impact of quality variability on revenues is non-monotone and depends on the direction/degree of demand-supply mismatch. If inventory is scarce, variability tends to hurt revenues, whereas if there is oversupply, it might potentially end up improving the revenues.
  The price update is more beneficial when selling scarce and highly variable quality inventory.
- The seller might be better off being able to adjust the price (even once) without full information of quality rather than setting a unique price (with no adjustments) with complete information.
- The fixed proportions policy is the most effective among all approximate policies, followed by the mean-value equivalent policy. Performance of these policies improves as inventory becomes more abundant and less variable in quality. However, under adverse operational conditions (scarce inventory and highly variable quality), approximations of optimal pricing strategies might lead to significant revenue losses.

Heterogeneous quality presents a unique challenge in terms of modeling. To circumvent some of the difficulties in the analysis, we assumed that customers are homogeneous in their preferences. It would be useful to extend the models here to the case where both inventory and customer choices are heterogeneous. We also assumed that the search cost to look for the best available item is negligible. It would be interesting to take into account nonnegligible search costs but this appears to be a challenging extension. Finally, there might be extensions to consider on the inventory side such as replenishment policies and salvage markets for unsold products.

REFERENCES

Akcay, Y., Boyacı, T. & Zhang, D. (2013) Selling with money-back guarantees: the impact on prices, quantities, and retail profitability. *Production and Operations Management*, 22(4), 777–791.

Akcay, Y., Natarajan, H. & Xu, S.H. (2010) Joint dynamic pricing of multiple perishable products. *Management Science*, 56(8), 1345–1361.

Aviv, Y. & Pazgal, A. (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management*, 10(3), 339–359.

Berry, S. & Pakes, A. (2007) The pure characteristics demand model. *International Economic Review*, 48(4), 1193–1225.

Besanko, D.& Winston, W. (1990) Optimal price skimming by a monopolist facing rational consumers. *Management Science*, 36, 555–567.

Bhargava, H. & Choudhary, V. (2001) Information goods and vertical differentiation. *Journal of Management Information Systems*, 18(2), 89–106.

Bitran, G. & Caldentey, R. (2003) An overview of pricing models for revenue management. *Manufacturing & Service Operations Management*, 5(3), 203–229.

Bitran, G.& Mondschein, S.V. (1997) Periodic pricing of seasonal products in retailing. *Management Science*, 43, 63–79.

Bresnahan, T. (1987) Competition and Collusion in the American Automobile Industry: the 1955 Price War. *Journal of Industrial Economics*, 35(4), 457–482.

Chen, X.& Simchi-Levi, D. (2012) Pricing and inventory management. *The Oxford Handbook of Pricing Management*, 1, 784–824.

David, H. & Nagaraja, H. (2003) *Order statistics*, 3rd edition. New York, NY: Wiley.

De Feo, M. (2005) *The Juran Institute Research on Cost of Poor Quality*. Southington, Connecticut: The Juran Institute.

Elmaghraby, W. & Keskinocak, P. (2003) Dynamic pricing in the presence of inventory considerations: research overview, current practices, and future directions. *Management Science*, 49(10), 1287–1309.

Fruchter, G. & Gerstner, E. (1999) Selling with "Satisfaction Guaranteed". *Journal of Service Research*, 1(4), 313–323.

Gallego, G. & van Ryzin, G.J. (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40, 999–1020.

Garvin, D.A. (1984) What does quality really mean? *Sloan Management Review*, 25, 25–43.

Honhon, D. & Seshadri, S. (2013) Fixed vs. random proportions demand models for the assortment planning problem under stockout-based substitution. *Manufacturing & Service Operations Management*, 15(3), 378–386.

Hopp, W.J. & Xu, X. (2008) A static approximation for dynamic demand substitution with applications in a competitive market. *Operations Research*, 56(3), 630–645.

Lariviere, M.A. & Porteus, E.L. (2001) Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & Service Operations Management*, 3(4), 293–305.

Leffler, K.B. (1982) Ambiguous changes in product quality. *The American Economic Review*, 72(5), 956–967.

Netessine, S. (2006) Dynamic pricing of inventory/capacity with infrequent price changes. *European Journal of Operational Research*, 174(1), 553–580.

Pan, X. & Honhon, D. (2012) Assortment planning for vertically differentiated products. *Production and Operations Management*, 21(2), 253–275.

Phillips, R.L. (2005) *Pricing and revenue management*. Stanford, CA: Stanford University Press.

Su, X. (2009) Consumer return policies and supply chain performance. *Manufacturing & Service Operations Management*, 11(4), 595–612.

Sunley, J. (1963) Factors affecting timber quality. *The Commonwealth Forestry Review*, 42, 129–136.

Talluri, K. & Van Ryzin, G. (2005) *The theory and practice of revenue management*, volume 68. Boston, MA: Springer Verlag.

Tirole, J. (1988) *The theory of industrial organization*. Cambridge, MA: MIT Press.

Wauthy, X. (1996) Quality choice in models of vertical differentiation. *The Journal of Industrial Economics*, 44, 345–353.

AUTHOR BIOGRAPHIES

Yalçın Akçay is a professor of operations management at the Melbourne Business School, University of Melbourne. He received his dual-title Ph.D. in Operations Research and Business Administration from the Penn State University's Smeal College of Business in 2002. Prior to joining the Melbourne Business School, he was an associate professor at Koç University. His research focuses on revenue management, dynamic pricing, inventory management, retail operations, stochastic modelling of service, and manufacturing systems.

Fikri Karaesmen is a professor in the Department of Industrial Engineering at Koç University. He received his B.S. degree from METU (Ankara, Turkey) in 1990 and M.S. and Ph.D. degrees from Northeastern University (Boston, MA, USA) in 1992 and 1996. Prior to joining Koç University, he was a faculty member at Ecole Centrale Paris (France). He spent the 2008-2009 academic year as a visiting professor at Northwestern University. His main research interests are in operations research and stochastic modelling with applications to production, inventory, and service systems.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Akçay Y, Karaesmen F. Customers know best: Pricing policies for products with heterogeneous quality. *Decision Sciences*. 2021;1–12. https://doi.org/10.1111/deci.12544