OPERATIONS RESEARCH
Vol. 55, No. 1, January–February 2007, pp. 75–97
ISSN 0030-364X | EISSN 1526-5463 | 07 | 5501 | 0075

DOI 10.1287/opre.1060.0344

Price Guarantees in Dynamic Pricing and Revenue Management

Yuri Levin, Jeff McGill, Mikhail Nediak
School of Business, Queen's University, Kingston, Ontario, Canada K7L 3N6
{ylevin@business.queensu.ca, jmcgill@business.queensu.ca, mnediak@business.queensu.ca}

We present a new model for revenue management of product sales that incorporates both dynamic pricing and a price guarantee. The guarantee provides customers with compensation if, prior to a fixed future date, the price of the product drops below a level specified at the time of purchase. We consider the problem of simultaneously determining optimal dynamic price and guarantee policies for items from a fixed stock when demand depends both on the price and on the parameters of the price guarantee. The model can be used for pricing any items with limited availability over a fixed time horizon. We formulate this model as a discrete-time optimal control problem, prove the existence of its optimal solution, explore some of the structural properties of the solution, present lower-bounding heuristics for solving the problem, and report numerical results.

Subject classifications: inventory/production: perishable/aging items, marketing/pricing, uncertainty; dynamic programming/optimal control: applications; probability: stochastic model applications; programming: nonlinear.
Area of review: Stochastic Models.
History: Received June 2005; revision received January 2006; accepted February 2006.

1. A Price Guarantee Instrument

In the 30 years since the successes of revenue management systems in airlines were first reported, applications have spread steadily into other business areas. Revenue management is now common in such service businesses as passenger railways, cruise lines, hotel and motel accommodation, and car rentals. Other applications have been proposed in such diverse areas as broadcast advertising, sports and entertainment event management, medical services, real estate, freight transportation, and manufacturing. Details of implementation across these areas differ, but most share a basic set of attributes: a finite supply of some inventory item, a fixed time span over which sales can occur, and stochastically varying demand.

Generally speaking, revenue management systems can operate in one of two ways. First, a company can offer a set of demand or product classes and assign a different, relatively stable price to each class. The classes are differentiated through restrictions on product features like time of purchase, time of use, or refundability. The company then controls the availability of inventory in the different classes over time in response to fluctuating demand patterns. (This is the traditional form of revenue management for large airlines.) Second, they can simply offer a single product and dynamically vary the price over time. In either case, the company's objective is to establish pricing and inventory control policies that maximize expected revenues.

There is an extensive literature on revenue management and related practices. For books and surveys, see Phillips (2005), Talluri and van Ryzin (2004b), McGill and van Ryzin (1999), Belobaba (1987), and Weatherford and Bodily (1992). The studies most relevant to this paper are those on dynamic pricing by Gallego and van Ryzin (1994), Feng and Xiao (2000b, a), Zhao and Zheng (2000), and Chatwin (2000).

Whichever system is employed, the customers of a revenue-managed company are confronted with prices that can vary dramatically from day to day. This presents them with a dilemma when deciding whether or not to purchase at a given point in time—should they purchase now and be sure that they acquire the item, or wait for a price drop and risk failing to acquire the item if it sells out?

This problem of customer uncertainty about future prices has long been recognized in the retail sector, and many retailers offer some form of price guarantee to encourage customers not to delay their purchase. In one type of *internal price-matching* guarantee, a retailer ensures that a customer will be reimbursed the difference between the current purchase price and any lower price the retailer might offer within a fixed future time period. An alternative *external price-matching* guarantee offers to match the price advertised by any other retailer at the time of purchase. A limited amount of research on price matching has been done by economists and marketing scientists; see Hess and Gerstner (1991), Moorthy and Winter (2005), and Srivastava and Lurie (2001). These studies focus on the economics of price matching and do not involve revenue management practices.

In this paper, we present a new dynamic pricing model for revenue management that includes an internal price guarantee instrument. This instrument provides a customer with compensation if the price of a similar product drops below a specified level (which we call the *strike price* following the custom in option pricing). The model incorporates both a fee that may be charged for the price guarantee and the option for the customer to decline the guarantee.

A price guarantee should have the effect of increasing the probability that customers will purchase at or near the time they first inquire about a product because it greatly reduces their risk of future opportunity loss. For the company, increased early purchases can reduce the uncertainty of late-purchasing "rushes" and last-minute price reductions, facilitate forecasting and capacity planning, and improve customer satisfaction and retention. Furthermore, the price guarantee itself constitutes a service provided for a fee in addition to the regular product price. Because this fee can be set by the seller so that it exceeds potential average losses from paying compensations, the collected fees provide an additional revenue stream.

There are many possible variations of a price guarantee instrument. In our view, the most reasonable way to compute compensation is to take the difference between the price of the product at some future point in time and a strike price offered at the time of purchase. There are several possibilities for the selection of the time when the price is compared to the strike price. One of them is to allow a customer to select it. Another is to let it be the time when the price is the lowest (in which case, the seller itself would monitor it). Another set of the instrument's details is related to the notion of "similar product." In the case of airlines, for example, this could be "a ticket for the exact same flight," "a ticket for the same route by the same airline on the same date," or "a ticket on the same date by any airline with flights to this destination." Each definition of similarity will result in its own price for the service. In this paper, we restrict our attention to the case of the same items and the comparison of the price with the strike price made at the time when the price is the lowest.

The practice of internal price matching is a very common one in the retail industry. In fact, such a policy is essentially forced on the many retailers who offer "free returns" policies for a fixed period after purchase because in those cases customers who witness a price drop can simply return their purchased item for refund and replace it with the same item at a lower price. However, to get the maximum benefit from price matching the customer needs to monitor the price constantly. If all customers are willing to do that, it is certainly the worst case for the company in terms of its revenues. Our model maximizes the value of expected revenues for the company in this case. This value represents a lower bound for the general case when the customers may or may not observe the price constantly. On the other hand, if a company wants to promote goodwill, it may choose to provide the benefit of monitoring the lowest price after the purchase to its customers. In this case, the value provided by our model is optimal.

We also show the relevance of the proposed price guarantee mechanism to other industry practices that to some extent implement price guarantees—for example, fully or partially refundable tickets in the travel industry.

In this paper, we formulate the dynamic pricing with guarantee model as a discrete-time optimal control problem, solve it as a large-scale nonlinear programming (NLP) problem, show that an optimal solution exists under reasonable assumptions, and that this optimal solution has natural structural properties. The NLP problem is tractable when the number of items is small; however, for larger instances we need tractable approximations, so we propose a myopic lower-bound heuristic. For an overview of discrete-time optimal control and nonlinear programming, see, for example, Bertsekas (2000) and Bertsekas (1995).

The main contributions of this paper are:

- introduction of a new technique for managing demand in revenue management systems, with a potential to boost sales by reducing customer uncertainty about future prices, and
- analysis of a model that is non-Markovian with respect to the underlying demand process.

This paper is organized as follows. We present the notation and model formulation in §2. We analyze the proposed model in §2.2, prove the existence of a solution in §2.3, and consider the special case of free price guarantees and its relevance to several current practices in §2.4. A myopic lower-bounding heuristic for solving this problem is discussed in §2.5. We report numerical results in §3. Section 4 contains conclusions and directions for future research. Proofs of the lemmas and propositions are available in an online appendix at http://or.pubs.informs.org/Pages.collect.html.

2. Model Formulation

We consider the simplest case of the price guarantee: the compensation applies to the same items only, and it is offered by a company itself. A company can then include the offer of the compensation and its details into the pricing policy. We also assume that compensation is calculated as the difference between the lowest price for the product that occurs after the time of purchase, and the strike price if the lowest price is below the strike price. The price guarantees may also have a fixed duration after the time of sale and expire prior to the end of the planning period.

Our notation and assumptions for the dynamic pricing elements of the model are based on the continuous-time formulation of Gallego and van Ryzin (1994), Feng and Xiao (2000b, a), Zhao and Zheng (2000), and Chatwin (2000). While a general continuous-time version of our model can be written, it is not analytically or computationally practical, thus we move directly to a discrete-time formulation.

Assume that customer inquiries about a quote arrive according to a discrete-time counting process $N(t)$ with at most one arrival per time period and the probability of arrival λ in each period. (The actual sales process will be nonhomogeneous in time in a manner described below.)

The company has a total of Y items in inventory available for sale during T time periods. The price offered at time t is specified by a stochastic process $p(t)$ which is a part of the company's policy. The strike price at time t is specified by a process $k(t)$, and the fee for the guarantee, $f(t)$, is also part of the policy. Thus, the quote is a three-dimensional stochastic process $(p(t), k(t), f(t))$. For brevity, we often use Π to denote the quote vector (p, k, f); thus $\Pi(t) = (p(t), k(t), f(t))$. We assume that the guarantee duration D is fixed but do not allow guarantees to extend past the end of the planning period T. In the case that all guarantees do not expire until the end of the planning period, we let $D = T$.

In this paper, we consider policies constrained by $0 \leqslant f(t) \leqslant k(t) \leqslant p(t)$. Obviously, f should be less than or equal to k because any potential price guarantee payment to a customer is less than k. The restrictions $k \leqslant p$ and $f \geqslant 0$ are not strictly necessary: An offer with $k > p$ could represent a "cash back" promotion, while an offer with $f < 0$ could represent a "sale" combined with a price guarantee. However, combining consumer response to price guarantees with such special cases as cash back and sale promotions leads to a cumbersome demand model. In this paper, we restrict our attention to the most common situation with price guarantees only, although the mathematical methodology developed here can be modified to handle the other two promotions. A special case of our model when $k(t) = f(t) = 0$ is the discrete-time analogue of the well-known model of Gallego and van Ryzin (1994).

For modeling consumer choice, we follow a general framework similar to Talluri and van Ryzin (2004a), and do not consider potential strategic behavior or history-dependent choice behavior of consumers. Customers, who are assumed to be myopic, choose between not making a purchase, making a purchase without the price guarantee at price $p(t)$, or paying $p(t) + f(t)$ for the purchase with the guarantee. The consumer choice model can be described by two probability functions. First, we assume that the probability that a customer makes either type of purchase upon arrival at time t is given by a function $u(\Pi, t)$. Thus, the effective probability of sale at time t is $\lambda(\Pi(t), t) = u(\Pi(t), t)\lambda$. This allows us to consider general nonhomogeneous demand because the nonhomogeneity can be accounted for through the appropriate choice of $u(\Pi, t)$. Second, the conditional probability that a customer also purchases the price guarantee (given that a purchase is made) is given by $v(\Pi(t), t)$. The sales process is a two-dimensional counting process $(N_1(t), N_2(t))$, where $N_1(t)$ gives the number of sales without price guarantees and $N_2(t)$ gives the number of sales with price guarantees. In each time period, a sale without the price guarantee (an event counted in $N_1(t)$) occurs with probability $\lambda u(\Pi(t), t)(1 - v(\Pi(t), t))$, a sale with the price guarantee (an event counted in $N_2(t)$) occurs with probability $\lambda u(\Pi(t), t)v(\Pi(t), t)$, and no sale occurs with probability $\lambda(1 - u(\Pi(t), t))$.

We assume that the price guarantee payments always happen at the end of the planning period (at time T) regardless of the guarantee duration D. If a customer bought an item with the price guarantee at time t, then the price guarantee payment to this customer will be equal to

$$\max\{(k(t) - p(\tau))^+ : t \leqslant \tau < \min\{T, t + D\}\}.$$

In Markovian formulations of dynamic pricing, pricing decisions depend only on the remaining inventory at time t. However, pricing decisions with price guarantees no longer possess this property because the entire price path after a guarantee sale affects the objective function and, consequently, future decisions of the company. Thus, the controls in our model are not Markovian with respect to the sales process $(N_1(t), N_2(t))$.

The class of policies used in the model should be practical and based on the information available to a company. We consider policies dependent only on the history of the sales $(N_1(t), N_2(t))$ and quote $\Pi(t)$ processes up to time t. One of the standard approaches to stochastic control problems is dynamic programming (DP). If we use a standard DP formulation, then the history of the sales and quote processes has to be included in the state representation because DP solves a *synthesis* problem of finding current optimal policy values for *all* possible states, even those arising under nonoptimal policies. Including the histories in the state representation results in a formulation whose state dimension expands in time, which is impractical. However, the use of the entire history of the quote process is redundant, as the future revenues only depend on the strike prices of the previously sold guarantees and the minimums of the price process since the times of those sales up to t. Thus, the model formulation based on policies of this form can be reduced by an appropriate state augmentation to a control problem for a Markov process with the dimension bounded by a linear function of Y. Unfortunately, the dimension of this process is still too high for most practical purposes. An alternative approach, which allows us to include only the history of the sales process in the state representation, is based on *nonrandomized policies*; that is, policies that at each time t are functions of the histories of the sales and quote processes up to time t. Moreover, an optimal nonrandomized policy exists even if we optimize over a much wider class of *randomized policies* under the assumption that the control (quote) Π belongs to a bounded set, and the probabilities $u(\cdot)$ and $v(\cdot)$ are continuous on this set; see the corollary on page 15 of Gīhman and Skorohod (1979). While such policies lead a priori to a random quote $\Pi(t)$ at time t, this randomness is only due to the randomness of the sales process $(N_1(t), N_2(t))$, and the history of the sales process alone fully determines the policy at time t. The problem is then to determine the policy values for all possible histories of the sales process. We will show that a nonlinear programming (NLP) approach to solution permits optimizing all of these policy values simultaneously.

Therefore, we do not need to explicitly include the history of the quote process in the state representation, and the state is identified with the history of the sales up to time t.

We will define the policy variables (p, k, f) of the model for all possible histories of the sales process $(N_1(t), N_2(t))$, which, in the discrete-time case, can be described as triples $(\mathcal{N}_1, \mathcal{N}_2, t)$ such that:

(a) The elements of the (possibly empty) ordered lists $\mathcal{N}_1, \mathcal{N}_2 \subseteq \{0, \dots, T-1\}$ are distinct time points. List elements have to be strictly smaller than t because the policy cannot depend on a concurrent sale.

(b) t belongs to the set $\{0, \dots, T-1\}$. Also, if the cardinality of $\mathcal{N}_i$ is $|\mathcal{N}_i|$, $i = 1, 2$, the triples for which there are still items available for sale (therefore policy variables should be defined) must satisfy $|\mathcal{N}_1| + |\mathcal{N}_2| < Y$.

We denote by $\mathfrak{N}$ the set of triples that satisfy the above conditions and refer to such triples as *internal*. The set of *boundary* triples $\mathfrak{N}^\Delta$ includes triples $(\mathcal{N}_1, \mathcal{N}_2, t)$ satisfying (a) above and either one of two boundary possibilities instead of (b):

- $|\mathcal{N}_1| + |\mathcal{N}_2| = Y$ and $t = \max\{\mathcal{N}_1, \mathcal{N}_2\} + 1$,
- $|\mathcal{N}_1| + |\mathcal{N}_2| < Y$ and $t = T$.

Finally, we let $\bar{\mathfrak{N}} = \mathfrak{N}^\Delta \cup \mathfrak{N}$ be the set of *free* triples.

The boundary triples are the elementary outcomes in the probability space for our discrete-time model. These outcomes are sample paths of the discrete-time sales process. Internal triples only specify these paths up to, but not including, a given time t and can be thought of as events in the probability space of the model.

We introduce *expected revenue variables* $J(\mathcal{N}_1, \mathcal{N}_2, t)$ for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \bar{\mathfrak{N}}$, and *policy* variables $p(\mathcal{N}_1, \mathcal{N}_2, t)$, $k(\mathcal{N}_1, \mathcal{N}_2, t)$, and $f(\mathcal{N}_1, \mathcal{N}_2, t)$ for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$. In the probability space of our model, the value of $J(\mathcal{N}_1, \mathcal{N}_2, t)$ will always be equal to the expected revenue over the interval $[t, T]$, given that the previous sales occurred at times described by $\mathcal{N}_1$, $\mathcal{N}_2$ and the policy was selected according to $p(\cdot)$, $k(\cdot)$, $f(\cdot)$. In particular, $J(\varnothing, \varnothing, 0)$ will give the expected revenues over the entire interval $[0, T]$. This is enforced by the use of the appropriate constraints. We will use the abbreviated notation

$$\Pi(\mathcal{N}_1, \mathcal{N}_2, t) = (p(\mathcal{N}_1, \mathcal{N}_2, t), k(\mathcal{N}_1, \mathcal{N}_2, t), f(\mathcal{N}_1, \mathcal{N}_2, t))$$

for the policy vector at a particular triple, and write

$$\bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t) = \lambda u\big(p(\mathcal{N}_1, \mathcal{N}_2, t), k(\mathcal{N}_1, \mathcal{N}_2, t), f(\mathcal{N}_1, \mathcal{N}_2, t), t\big) \quad (1)$$

and

$$v\big(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t\big) = v(p(\mathcal{N}_1, \mathcal{N}_2, t), k(\mathcal{N}_1, \mathcal{N}_2, t), f(\mathcal{N}_1, \mathcal{N}_2, t), t\big) \quad (2)$$

for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$. The value of $\bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)$ is the probability of one sale at time t given sales history $\mathcal{N}_1$, $\mathcal{N}_2$. The value of v is the conditional probability of the price guarantee sale given that an item was sold. The feasible policy must satisfy

$$p(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant k(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant f(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant 0 \quad \text{for all } (\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}. \quad (3)$$

The expected revenues given the history of the sales process $\mathcal{N}_1$, $\mathcal{N}_2$ at time t are computed via the expected revenues at $t+1$ as follows:

$$\begin{aligned} J(\mathcal{N}_1, \mathcal{N}_2, t) &= \bar{u}\big(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t\big) \\ &\quad \cdot \big\{p(\mathcal{N}_1, \mathcal{N}_2, t) + \big(1 - v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)\big) \\ &\quad \cdot J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) + v\big(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t\big) \\ &\quad \cdot \big[f(\mathcal{N}_1, \mathcal{N}_2, t) + J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1)\big]\big\} \\ &\quad + \big(1 - \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)\big) J(\mathcal{N}_1, \mathcal{N}_2, t+1) \\ &\quad \text{for all } (\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}. \end{aligned} \quad (4)$$

This recursion forms the *revenue decomposition* constraint in our model.

To account for the total price guarantee payments at time T, we need to introduce additional *terminal loss* variables $z_t(\mathcal{N}_1, \mathcal{N}_2)$ for all $t \in \mathcal{N}_2$, and all boundary triples $(\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta$. The variable $z_t(\mathcal{N}_1, \mathcal{N}_2)$ represents a bound on the amount the company has to pay in the future because of the price guarantee sale that occurred at time $t \in \mathcal{N}_2$. Although $z_t(\mathcal{N}_1, \mathcal{N}_2)$ is introduced for $(\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta$, we will not explicitly specify t' in the notation $z_t(\mathcal{N}_1, \mathcal{N}_2)$. Indeed, for given disjoint lists $\mathcal{N}_1$ and $\mathcal{N}_2$, there always exists a unique t' such that we can form a boundary triple $(\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta$. In particular, we have $t' = T$ for the case of $|\mathcal{N}_1| + |\mathcal{N}_2| < Y$, and $t' = \max\{\mathcal{N}_1, \mathcal{N}_2\} + 1$ when $|\mathcal{N}_1| + |\mathcal{N}_2| = Y$.

The terminal loss variables are constrained to be nonpositive:

$$z_t(\mathcal{N}_1, \mathcal{N}_2) \leqslant 0 \quad \text{for all } (\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta,\ t \in \mathcal{N}_2. \quad (5)$$

The total guarantee payments (which we treat as nonpositive values) are expressed via terminal loss variables as

$$J(\mathcal{N}_1, \mathcal{N}_2, t') = \sum_{t \in \mathcal{N}_2} z_t(\mathcal{N}_1, \mathcal{N}_2) \quad \text{for all } (\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta, \quad (6)$$

while

$$\begin{aligned} z_t(\mathcal{N}_1, \mathcal{N}_2) &\leqslant p(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau) \\ &\quad - k(\mathcal{N}_1 \cap [0, t), \mathcal{N}_2 \cap [0, t), t) \\ &\text{for all } (\mathcal{N}_1, \mathcal{N}_2, t') \in \mathfrak{N}^\Delta,\ t \in \mathcal{N}_2,\ \tau \in (t, \min\{t', t+D\}). \end{aligned} \quad (7)$$

This latter constraint captures the following fact: The amount of loss for a price guarantee sold at time $t \in \mathcal{N}_2$ will be greater than or equal to the difference between the strike price of the sale and the price at every time $\tau \in (t, \min\{t', t+D\})$.

The discrete-time expected revenue maximization problem is to find the policy and corresponding revenue and terminal loss variables which

$$\max J(\varnothing, \varnothing, 0)$$

under constraints (3)–(7).

Note that in any optimal solution only those $z_t(\mathcal{N}_1, \mathcal{N}_2)$s for which the minimum of the right-hand side of (7) is less than zero will be strictly negative.

Using $\bar{u}v$, $\bar{u}(1-v)$, and $(1-\bar{u})$ as probabilities of a purchase with a price guarantee, a purchase without a price guarantee, and no purchase at a particular state in the discrete-time model, one can recursively define the probabilities of triples $(\mathcal{N}_1, \mathcal{N}_2, t) \in \bar{\mathfrak{N}}$ as

$$P(\mathcal{N}_1, \mathcal{N}_2, t) = \begin{cases} \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t) \\ \quad \cdot P(\mathcal{N}_1, \mathcal{N}_2 \backslash \{t-1\}, t-1) \quad \text{if } t-1 \in \mathcal{N}_2, \\ \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)(1 - v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)) \\ \quad \cdot P(\mathcal{N}_1 \backslash \{t-1\}, \mathcal{N}_2, t-1) \quad \text{if } t-1 \in \mathcal{N}_1, \\ (1 - \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t))P(\mathcal{N}_1, \mathcal{N}_2, t-1) \\ \qquad \text{if } t-1 \notin \mathcal{N}_1, \mathcal{N}_2, \end{cases} \tag{8}$$

where $P(\varnothing, \varnothing, 0) = 1$. In this way, we complete the construction of the probability space of our discrete-time model because the above formula defines the probabilities for all elementary outcomes $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}^\Delta$.

We next compare our modeling approach with DP.

Remark 1. Once the time horizon T and the initial inventory Y are given, the set of all triples $\bar{\mathfrak{N}}$ defines a probability tree rooted at $(\varnothing, \varnothing, 0)$. An example of such a tree for $Y = 2$, $T = 3$ is given in Figure 1. Each node $(\mathcal{N}_1, \mathcal{N}_2, t)$ branches off into three other nodes $(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$, $(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1)$, and $(\mathcal{N}_1, \mathcal{N}_2, t+1)$ (a purchase without a price guarantee, a purchase with a price guarantee, and no purchase). The terminal nodes of the tree are elements of $\mathfrak{N}^\Delta$ (those with $t = T$ or $|\mathcal{N}_1| + |\mathcal{N}_2| = Y$). Because the policy is nonrandomized, once we have jointly specified $\Pi = (p, k, f)$ at all internal nodes in the tree, the complete history of the quote process is available at every node $(\mathcal{N}_1, \mathcal{N}_2, t)$. The optimization is performed jointly at all nodes. Thus, while the policy variables at $(\mathcal{N}_1, \mathcal{N}_2, t)$ do not explicitly depend on the history of the quote process

$$\{\Pi(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau), \ \tau = 0, \ldots, t\},$$

the dependence is embedded into the payoff (and the decision) through the terminal loss variables.

In contrast, a DP approach would require that the history of the policy up to the node $(\mathcal{N}_1, \mathcal{N}_2, t)$ be included in the state description. Even with a discrete set of feasible policies, this would result in a very large problem that is unlikely to be computationally tractable.

Figure 1. The probability tree of our model in the case of $Y = 2$ and $T = 3$.

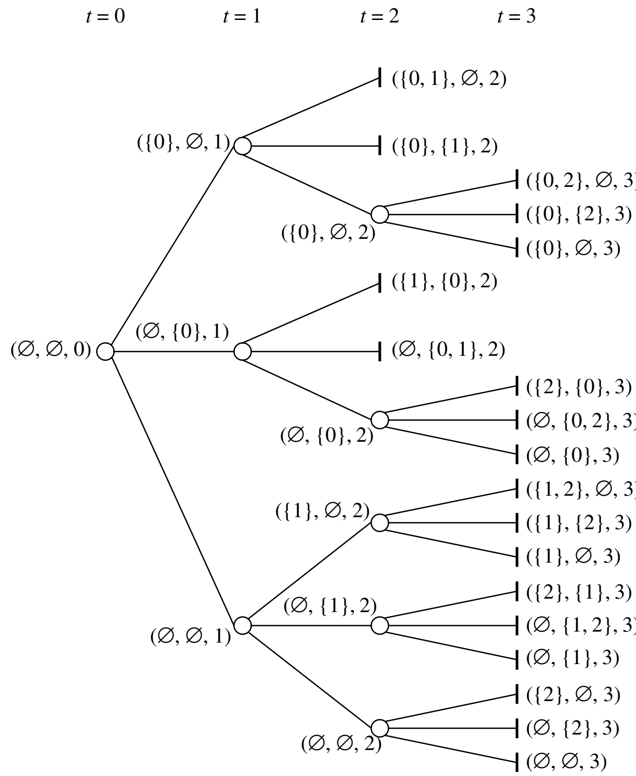

2.1. Selecting a Consumer Choice Model—A Functional Form of *u* and *v*

In this paper, we describe two ways to select a consumer choice model (functions u and v). First, we present a constructive approach to consumer choice modeling in §2.1.1. Second, in §2.1.2, we present the set of natural assumptions that a general model for u and v has to satisfy, and then give an example of a specific model that satisfies all these assumptions.

2.1.1. Constructive Approach. In this approach, we need to construct a decision model that describes how customers make the actual purchasing decisions. This model will suggest a specific form of u and v. Because the customers in our model are assumed to be myopic, it is reasonable to assume that their rationality is bounded and they cannot reconstruct the distribution of the future policy. An argument in favor of this assumption is that at the time of decision making a customer is not likely to have access to the sales history and the exact consumer choice model used by the company. If a customer wanted to reconstruct the distribution of the future policy, he/she would not only need to recompute the policy for every possible consumer choice model, but also have to employ a prior distribution on all possible history triples at time t, and possible consumer choice models. The task of future policy distribution reconstruction does not look feasible under these circumstances. Therefore, we assume that customers have some a priori anticipations about the future. In line with the bounded-rationality nature of the customers, we also assume that these anticipations do not change after the price quote is observed. One possible way to describe a customer's anticipations is in terms of "anticipated cost" of an item for a particular customer, X_c, which is a random variable. Due to variability between customers, the distribution of this random variable will also vary. Thus, we let the index c specifying the distribution of X_c be random. We further assume that a customer evaluates the offer in terms of the expected surplus $E[\texttt{Surplus} \mid c]$, where

$$\texttt{Surplus} = X_c - \texttt{overall cost for the customer}.$$

At each time t, a customer has three choices: Choice 1—do not buy, Choice 2—buy at price p without price guarantee, and Choice 3—buy at price $p+f$ with price guarantee at strike price k, $f \leqslant k \leqslant p$. The expected surpluses $E[\texttt{Surplus}_i \mid c]$ for choices $i=1$, 2, and 3 are

$$E[\texttt{Surplus}_1 \mid c] = 0, \tag{9}$$

$$\begin{aligned} E[\texttt{Surplus}_2 \mid c] &= E[X_c - p \mid c] \\ &= -p + E[X_c \mid c], \quad \text{and} \end{aligned} \tag{10}$$

$$\begin{aligned} E[\texttt{Surplus}_3 \mid c] &= E[X_c - (p + f - (k - X_c)^+) \mid c] \\ &= -p - f + E[\max\{X_c, k\} \mid c]. \end{aligned} \tag{11}$$

Because the distribution of the anticipated cost is determined by the value of c, the choice in terms of these expected values is deterministic for a particular customer.

The customers, however, are drawn at random from an infinite population of potential customers. The result of the draw is the value of c. Thus, a randomly drawn customer will make a purchase with probability

$$\begin{aligned} u &= P(\max\{E[\texttt{Surplus}_2 \mid c], E[\texttt{Surplus}_3 \mid c]\} \geqslant 0) \\ &= P(E[\texttt{Surplus}_2 \mid c] \geqslant 0 \text{ or } E[\texttt{Surplus}_3 \mid c] \geqslant 0). \end{aligned} \tag{12}$$

The probability of purchase with price guarantee is

$$\begin{aligned} uv &= P(E[\texttt{Surplus}_3 \mid c] \geqslant \max\{E[\texttt{Surplus}_2 \mid c], 0\}) \\ &= P(E[\texttt{Surplus}_3 \mid c] \geqslant E[\texttt{Surplus}_2 \mid c] \text{ and} \\ &\qquad E[\texttt{Surplus}_3 \mid c] \geqslant 0). \end{aligned} \tag{13}$$

Example 1. If the anticipated cost of an item X_c for a customer follows the uniform$(0, c)$ distribution, with parameter c following the exponential distribution with parameter $1/\mu$, then the probability of purchase

$$u = \begin{cases} \exp\left(-\dfrac{2p}{\mu}\right) & \text{if } k^2 < 4pf, \\ \exp\left(-\dfrac{1}{\mu}(p + f + \sqrt{(p+f)^2 - k^2})\right) & \\ \qquad\qquad \text{if } k^2 \geqslant 4pf, & \end{cases} \tag{14}$$

and the probability of purchase of price guarantee given that the purchase is made

$$v = \begin{cases} 0 \quad \text{if } k^2 < 4pf, \\ 1 - \exp\left(-\dfrac{1}{\mu}\left(-p - f + \dfrac{k^2}{2f} - \sqrt{(p+f)^2 - k^2}\right)\right) \\ \quad \text{if } k^2 \geqslant 4pf. \end{cases} \tag{15}$$

The functions $u(p, f, k)$ and $v(p, f, k)$ are continuous, as $u(p, k, f) = \exp(-2p/\mu)$ and $v(p, k, f) = 0$ at $k^2 = 4pf$. Expression (14) shows how the presence of price guarantees can boost the demand. Indeed, the probability of purchase u increases from $\exp(-2p/\mu)$ when $k = f > 0$ (unattractive price guarantee) to $\exp(-p/\mu)$ when $k = p$ and $f = 0$ (free price matching).

Remark 2. The risk attitude of customers can be incorporated into the above consumer choice model by replacing the expected surpluses $E[\texttt{Surplus}_i \mid c]$ with the mean-variance surpluses $E[\texttt{Surplus}_i \mid c] - \beta \operatorname{Var}[\texttt{Surplus}_i \mid c]$, $i = 1, 2, 3$, where β is a risk parameter increasing with risk aversion. The case $\beta = 0$ corresponds to risk-neutral behavior, and $\beta > 0$ and $\beta < 0$ give risk-averse and risk-seeking behavior, respectively.

The main drawback of the constructive approach to consumer choice model selection is the difficulty of obtaining closed-form expressions for u and v with the exceptions of the most simple settings. Note that such expressions are necessary for any numerical solution procedure. Moreover, it is hard to ensure that these functional expressions will be sufficiently smooth (solvers quite often require the availability of second partial derivatives for all model components). Thus, we next consider an alternative approach for selecting the functions u and v.

2.1.2. A General Model for u and v. Recall that $u(p, k, f, t)$ defines the probability that a customer accepts the quote (p, k, f) at time t, and $v(p, k, f, t)$ is the conditional probability that a customer will also purchase the price guarantee (given that an item is purchased). In selecting a model for u and v, we need to consider the following:

(1) It may often be reasonable to assume that the conditional probability of purchasing the price guarantee v depends only on time t, the *strike price ratio* $\kappa = k/p$, and the *fee ratio* $\phi = f/k$. Subsequent discussion assumes that this is true. We will also use a similar parametrization for u.

(2) The probability v should be decreasing in ϕ and increasing in κ (while the other ratio and t are fixed). Similar dependence will be assumed for u. Additionally, we will assume that u is decreasing in p for fixed κ, ϕ, t.

(3) The value of u should approach zero as $p \to \infty$.

(4) The value of v should approach zero as $\phi \to 1$ or $\kappa \to 0$.

(5) The value of v should approach one as $\phi \to 0$.

(6) The last two conditions seem to require that v is discontinuous at $\phi = 0$, $\kappa = 0$. This will have to be accounted

for in the analysis of the model. We will assume, however, that v is continuously differentiable anywhere for $0 \leqslant \phi, \kappa \leqslant 1$, $0 \leqslant t \leqslant T$, with an exception at $\phi = 0$, $\kappa = 0$.

(7) u is assumed to be continuously differentiable everywhere.

(8) Note that a reasonable model for v has to be time dependent. Indeed, because v represents the customer's interest in purchasing the price guarantee, it should approach zero at the end of the planning interval $[0, T]$ for all $\phi > 0$, while having strictly positive values around $t = 0$. It also seems reasonable to assume that v is decreasing in t for fixed ϕ, κ.

(9) Similarly, u can be time dependent so that as the value of t approaches T, the effects of price guarantee components ϕ, κ of the policy decrease. This corresponds to the assumption that the derivatives of u with respect to ϕ, κ decrease in t, and tend to zero as $t \to T$.

(10) The value of $u(0, \kappa, \phi, t)$ is uniformly bounded away from zero. That is, there exists $u_0 > 0$ such that $u(0, \kappa, \phi, t) \geqslant u_0$ for all κ, ϕ, t.

The following example shows plausible functions satisfying all of the above conditions. These specific functions are not required for the general results to follow, but are used in the numerical examples of §3.

Example 2. Consider

$$v(p, \kappa, \phi, t) = \left(\frac{\kappa}{\kappa + \delta\phi}\right)^{\beta} (1-\phi)^{\gamma} \left(1 - \frac{t}{T}\right)^{\phi/\rho}, \quad \beta, \gamma, \delta, \rho > 0. \quad (16)$$

The parameter β essentially determines how fast v decreases when κ is decreased from one. The parameter γ determines how fast v decreases as we increase ϕ from zero. The quantity $1/\delta$ is the value of ϕ that results in reduction of v by at least a factor of $(1/2)^{\beta}$ when $\kappa = 1$. Finally, ρ is a value of ϕ such that $v(p, \kappa, \phi, t)$ decreases linearly with time. For $\phi > \rho$, the decrease is superlinear in t ($v(p, \kappa, \phi, t)$ drops to near zero quickly in the beginning of the planning interval), while for $\phi < \rho$ the decrease is sublinear (it is very moderate in the beginning and fast at the end).

The dependence of u on ϕ, κ can be selected to have a functional form related to that of v. The important difference is that while it is very natural that $v = 1$ whenever $\phi = 0$ (and therefore it is constant in κ when $\phi = 0$), the value of u must still nontrivially depend on κ. One possibility is

$$u(p, \kappa, \phi, t) = \exp\{-p - \alpha[1 - \kappa^{\beta'}(1-\phi)^{\gamma'} v_0(t)]\}, \quad (17)$$

where $\alpha > 0$ specifies by how much the probability of purchase decreases when the price guarantee is unavailable, $\beta', \gamma' > 0$, and $v_0(t)$ is decreasing in t, with $v_0(0) = 1$. Note that the widely used exponential demand model $u = \exp(-p)$ is a special case of (17).

Because we have assumed in (1) that $v(\cdot)$ is of the form $v = v(k/p, f/k, t)$, it may be convenient to make the substitution $\kappa = k/p$, $\phi = f/k$ in the discrete-time model described in expressions (1) to (7). With this substitution, $v(\cdot)$ can be expressed as a function of only κ, ϕ, t, but not p ($v = v(\kappa, \phi, t)$), and our analysis will be somewhat simplified. The new policy variables must satisfy

$$0 \leqslant \kappa(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant 1, \quad (18)$$

$$0 \leqslant \phi(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant 1, \quad (19)$$

and we must also substitute

$$k(\mathcal{N}_1, \mathcal{N}_2, t) = \kappa(\mathcal{N}_1, \mathcal{N}_2, t) p(\mathcal{N}_1, \mathcal{N}_2, t), \quad (20)$$

$$f(\mathcal{N}_1, \mathcal{N}_2, t) = \phi(\mathcal{N}_1, \mathcal{N}_2, t) \kappa(\mathcal{N}_1, \mathcal{N}_2, t) p(\mathcal{N}_1, \mathcal{N}_2, t) \quad (21)$$

throughout the entire model. To emphasize the nature of the parameterization in vector notation, we will use $\mathit{\Pi} = (p, \kappa, \phi)$ versus $\Pi = (p, k, f)$. The choice of parameterization, (p, k, f) or (p, κ, ϕ), will depend on the context and convenience of notation. We emphasize that Assumption 1 is not required whenever the parameterization is (p, k, f).

2.2. The Analysis of the Model

In this section, we prove some natural monotonicity properties of the revenue variables, which must hold at the optimum. We also describe optimality conditions and provide a probabilistic interpretation for Lagrange multipliers corresponding to revenue decomposition constraints (4). The analysis is done under the assumption that $v = v(k/p, f/k, t)$ (which is Assumption 1 of the previous section), the substitution (20)–(21) is carried out, and the constraints (18)–(19) replace (3).

Our monotonicity results are intuitively obvious, but their careful verification, while straightforward, is somewhat lengthy. The intuitive summary is as follows. Given the sales history up to time t, consider expected future revenues derived from sales of remaining items and corresponding price guarantees (less all necessary price guarantee payments for sales that have already occurred).

(1) Consider an additional sale that may occur at time t without a price guarantee. The expected future revenue given that this sale has occurred cannot be larger than that which the company could earn given that no sale has happened.

(2) Also, if a sale happens at time t, the expected future revenue with a price guarantee cannot be larger than that without a price guarantee.

(3) Finally, given the same past sales history, the company cannot earn more in expected revenues when it has less time left.

Lemma 1, below, formalizes these results. Its logical consequence is Lemma 2, which shows that the price component of the policy may be bounded from below without sacrificing optimality. Both lemmas are proved in the online appendix.

LEMMA 1. *For any feasible solution, there exists a feasible solution of greater or equal expected value such that the following inequalities hold for any* $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$:

$$J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) \leqslant J(\mathcal{N}_1, \mathcal{N}_2, t+1), \tag{22}$$

$$J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) \leqslant J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1), \tag{23}$$

$$J(\mathcal{N}_1, \mathcal{N}_2, t+1) \leqslant J(\mathcal{N}_1, \mathcal{N}_2, t), \tag{24}$$

and the minimum price $\min_{(\mathcal{N}_1, \mathcal{N}_2, t)} p(\mathcal{N}_1, \mathcal{N}_2, t)$ *in the new solution is greater than or equal to the minimum price in the old solution.*

LEMMA 2. *For any feasible solution such that inequalities* (22)–(23) *are satisfied, there exists a feasible solution of greater or equal value such that*

$$\begin{aligned} p(\mathcal{N}_1, \mathcal{N}_2, t) &\geqslant p^* \\ &= \inf_{\substack{0\leqslant\kappa,\,\phi\leqslant 1\\ 0\leqslant t\leqslant T-1}} \inf\{p\colon p \text{ is a local } \max \text{ of } pu(p,\kappa,\phi,t)\} \end{aligned} \tag{25}$$

for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$. *The value of* p^* *is strictly positive.*

Lemmas 1 and 2 allow us to conclude the following proposition proved in the online appendix:

PROPOSITION 1. *Any optimal solution to the discrete-time expected revenue maximization problem satisfies* (22), (24), *and* (25) *for all triples* $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$ *such that* $P(\mathcal{N}_1, \mathcal{N}_2, t) > 0$. *For those triples where, in addition,* $v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t) < 1$, *inequality* (23) *is satisfied as well.*

Our discrete-time expected revenue maximization problem belongs to the general class of nonlinear programming problems. The optimality conditions for this class of problems are typically in the form of Karush-Kuhn-Tucker (KKT) necessary optimality conditions. To apply these conditions to our specific problem, we first introduce Lagrange multipliers for the constraints. Assuming that substitutions (20)–(21) have been carried out, the Lagrange multipliers are:

- $\pi(\mathcal{N}_1, \mathcal{N}_2, t)$—corresponding to the first-order approximation (4),
- $\pi(\mathcal{N}_1, \mathcal{N}_2, T)$—corresponding to the terminal loss costs (6),
- $\zeta(\mathcal{N}_1, \mathcal{N}_2, \tau, t)$—corresponding to the terminal loss upper bound (7) (note that $\tau > t$),
- $\zeta(\mathcal{N}_1, \mathcal{N}_2, t, t)$—corresponding to the terminal loss negativity constraint (5),
- $\eta_0(\mathcal{N}_1, \mathcal{N}_2, t)$ and $\eta_1(\mathcal{N}_1, \mathcal{N}_2, t)$—corresponding to the left and right inequalities, respectively, in the strike price ratio constraint (18), and
- $\xi_0(\mathcal{N}_1, \mathcal{N}_2, t)$ and $\xi_1(\mathcal{N}_1, \mathcal{N}_2, t)$—corresponding to the left and right inequalities in the fee ratio constraint (19).

We do not introduce a multiplier for the constraint $p(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant 0$ because it never needs to be tight in the optimal solution.

The Lagrangian as a function of all primal and dual variables is

$$\begin{aligned}
\mathscr{L} &= J(\varnothing,\varnothing,0) \\
&\quad - \sum_{(\mathcal{N}_1,\mathcal{N}_2,t)\in\mathfrak{N}} \pi(\mathcal{N}_1,\mathcal{N}_2,t) \\
&\qquad \cdot \big\{J(\mathcal{N}_1,\mathcal{N}_2,t) - J(\mathcal{N}_1,\mathcal{N}_2,t+1) - \bar{u}(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t) \\
&\qquad\quad \cdot \big[p(\mathcal{N}_1,\mathcal{N}_2,t) + J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1) \\
&\qquad\qquad - J(\mathcal{N}_1,\mathcal{N}_2,t+1) + v(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t) \\
&\qquad\qquad \cdot (\phi(\mathcal{N}_1,\mathcal{N}_2,t)\kappa(\mathcal{N}_1,\mathcal{N}_2,t)p(\mathcal{N}_1,\mathcal{N}_2,t) \\
&\qquad\qquad + J(\mathcal{N}_1,\mathcal{N}_2\cup t,t+1) - J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1))\big]\big\} \\
&\quad - \sum_{(\mathcal{N}_1',\mathcal{N}_2',t')\in\mathfrak{N}^\Delta} \pi(\mathcal{N}_1',\mathcal{N}_2',t') \\
&\qquad \cdot \Big\{J(\mathcal{N}_1',\mathcal{N}_2',t') - \sum_{t\in\mathcal{N}_2'} z_t(\mathcal{N}_1',\mathcal{N}_2')\Big\} \\
&\quad - \sum_{\substack{(\mathcal{N}_1',\mathcal{N}_2',t')\in\mathfrak{N}^\Delta,\, t\in\mathcal{N}_2',\\ t<\tau<\min\{t',t+D\}}} \zeta(\mathcal{N}_1',\mathcal{N}_2',\tau,t) \\
&\qquad \cdot \big\{z_t(\mathcal{N}_1',\mathcal{N}_2') - p(\mathcal{N}_1'\cap[0,\tau),\mathcal{N}_2'\cap[0,\tau),\tau) \\
&\qquad\quad + \kappa(\mathcal{N}_1'\cap[0,t),\mathcal{N}_2'\cap[0,t),t) \\
&\qquad\qquad \cdot p(\mathcal{N}_1'\cap[0,t),\mathcal{N}_2'\cap[0,t),t)\big\} \\
&\quad - \sum_{(\mathcal{N}_1',\mathcal{N}_2',t')\in\mathfrak{N}^\Delta,\, t\in\mathcal{N}_2'} \zeta(\mathcal{N}_1',\mathcal{N}_2',t,t) z_t(\mathcal{N}_1',\mathcal{N}_2') \\
&\quad - \sum_{(\mathcal{N}_1,\mathcal{N}_2,t)\in\mathfrak{N}} \big\{-\eta_0(\mathcal{N}_1,\mathcal{N}_2,t)\kappa(\mathcal{N}_1,\mathcal{N}_2,t) \\
&\qquad + \eta_1(\mathcal{N}_1,\mathcal{N}_2,t)(\kappa(\mathcal{N}_1,\mathcal{N}_2,t)-1) - \xi_0(\mathcal{N}_1,\mathcal{N}_2,t) \\
&\qquad \cdot \phi(\mathcal{N}_1,\mathcal{N}_2,t) + \xi_1(\mathcal{N}_1,\mathcal{N}_2,t)(\phi(\mathcal{N}_1,\mathcal{N}_2,t)-1)\big\}.
\end{aligned} \tag{26}$$

The KKT conditions for a primal-dual pair of variables include: primal feasibility, which includes all of the problem's constraints; dual feasibility, which stipulates that partial derivatives of the Lagrangian with respect to all primal variables are equal to zero, and the Lagrange multipliers corresponding to inequality constraints are nonnegative; and complementary slackness, which requires that the product of a Lagrange multiplier with the corresponding constraint is zero. Due to the relatively large size of our problem, we do not state all of these conditions here. Instead, we introduce only those required for further analysis.

The optimality conditions for the discrete-time problem imply the following interpretation of multipliers $\pi(\mathcal{N}_1, \mathcal{N}_2, t)$ for the revenue decomposition constraints as probabilities of triples $(\mathcal{N}_1, \mathcal{N}_2, t)$:

PROPOSITION 2. *The Lagrange multipliers* $\pi(\mathcal{N}_1, \mathcal{N}_2, t)$, *satisfying the KKT conditions, are such that*

$$\pi(\mathcal{N}_1, \mathcal{N}_2, t) = P(\mathcal{N}_1, \mathcal{N}_2, t) \quad \textit{for all } (\mathcal{N}_1, \mathcal{N}_2, t) \in \overline{\mathfrak{N}}.$$

Differentiate $\mathcal{L}$ with respect to $J(\mathcal{N}_1, \mathcal{N}_2, t)$ and let the derivative be equal to zero. For $t > 0$, there are three possible cases:

Case 1. $t-1 > \max\{\mathcal{N}_1 \cup \mathcal{N}_2\}$ so that

$$\frac{\partial \mathcal{L}}{\partial J(\mathcal{N}_1, \mathcal{N}_2, t)} = -\pi(\mathcal{N}_1, \mathcal{N}_2, t) + (1 - \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t-1))) \cdot \pi(\mathcal{N}_1, \mathcal{N}_2, t-1) = 0.$$

Case 2. $t-1 = \max\{\mathcal{N}_1\}$ so that

$$\frac{\partial \mathcal{L}}{\partial J(\mathcal{N}_1, \mathcal{N}_2, t)} = -\pi(\mathcal{N}_1, \mathcal{N}_2, t) + \bar{u}(\Pi(\mathcal{N}_1 \backslash \{t-1\}, \mathcal{N}_2, t-1)) \cdot (1 - v(\Pi(\mathcal{N}_1 \backslash \{t-1\}, \mathcal{N}_2, t-1))) \cdot \pi(\mathcal{N}_1 \backslash \{t-1\}, \mathcal{N}_2, t-1) = 0.$$

Case 3. $t-1 = \max\{\mathcal{N}_2\}$ so that

$$\frac{\partial \mathcal{L}}{\partial J(\mathcal{N}_1, \mathcal{N}_2, t)} = -\pi(\mathcal{N}_1, \mathcal{N}_2, t) + \bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2 \backslash \{t-1\}, t), t) \cdot v(\Pi(\mathcal{N}_1, \mathcal{N}_2 \backslash \{t-1\}, t), t) \cdot \pi(\mathcal{N}_1 \backslash \{t-1\}, \mathcal{N}_2, t-1) = 0.$$

Note also that

$$\frac{\partial \mathcal{L}}{\partial J(\varnothing, \varnothing, 0)} = 1 - \pi(\varnothing, \varnothing, 0) = 0.$$

Finally, observe that this results in recursive rules for computing $\pi(\mathcal{N}_1, \mathcal{N}_2, t)$s that are equivalent to those for $P(\mathcal{N}_1, \mathcal{N}_2, t)$s given in (8).

2.3. The Existence of an Optimal Solution

In this section, we study the existence of a solution to the discrete-time expected revenue maximization problem. It can be established under either one of the following conditions:

(i) There exist $\epsilon \in (0, 1)$ and $p_{\max}$ such that

$$\frac{p}{u}\frac{\partial u}{\partial p} \leqslant -\frac{1}{1-\epsilon}$$

for all $p > p_{\max}$, all $0 \leqslant \kappa \leqslant 1$, $0 \leqslant \phi \leqslant 1$, and $0 \leqslant t \leqslant T$;

(ii) There exists a *shutdown price* $p_{\max}$ such that $u(p, \kappa, \phi, t) = 0$ for all $p > p_{\max}$, $0 \leqslant \kappa \leqslant 1$, $0 \leqslant \phi \leqslant 1$, and $0 \leqslant t \leqslant T$.

The proof of existence under condition (ii) is very easy because then the feasible policy set can be made bounded. However, in practice, the shutdown price may be difficult to specify or too large for any practical purpose. Therefore, a practitioner may be forced to use a functional form for u such that the shutdown price does not exist. On the other hand, a *regularity condition*

$$\lim_{p\to\infty} pu(p) = 0, \tag{27}$$

related to (i), is frequently used in the literature, for example, in Gallego and van Ryzin (1994). Naturally, in this classical regularity condition, u does not depend on the price guarantee parameters κ and ϕ. In the absence of these parameters, (27) is a consequence of (i). Finally, we observe that natural extensions of the standard demand models (such as the generalization (17) of the exponential demand model, or a similar generalization of the power demand model) often do not have a shutdown price, but rather satisfy (i). Thus, we focus our analysis of existence on condition (i).

Due to the presence of several policy variables, our analysis is more complex than that of Gallego and van Ryzin (1994). Observe that (i) is a restriction on elasticity of demand with respect to price. However, the form of condition (i) is very specific to the choice of parameterization (p, κ, ϕ). This issue cannot possibly arise when u only depends on p. To illustrate this point, we perform a change of variables in (i) from (p, κ, ϕ) back to (p, k, f). Consider a function $\tilde{u} = \tilde{u}(p, k, f, t)$. In the (p, κ, ϕ) parameterization we have a corresponding function $u(p, \kappa, \phi, t) = \tilde{u}(p, p\kappa, p\kappa\phi, t)$. Differentiate u with respect to p:

$$\frac{\partial u}{\partial p} = \frac{\partial \tilde{u}}{\partial p} + \kappa\frac{\partial \tilde{u}}{\partial k} + \kappa\phi\frac{\partial \tilde{u}}{\partial f} = \frac{\partial \tilde{u}}{\partial p} + \frac{k}{p}\frac{\partial \tilde{u}}{\partial k} + \frac{k}{p}\frac{f}{k}\frac{\partial \tilde{u}}{\partial f}.$$

Condition (i) then transforms to

$$\frac{p}{u}\frac{\partial u}{\partial p} = \frac{p}{\tilde{u}}\frac{\partial \tilde{u}}{\partial p} + \frac{k}{\tilde{u}}\frac{\partial \tilde{u}}{\partial k} + \frac{f}{\tilde{u}}\frac{\partial \tilde{u}}{\partial f} \leqslant -\frac{1}{1-\epsilon},$$

which is a restriction on the sum of elasticities of demand with respect to all variables of the policy triple.

The existence proof will require several steps. We first claim the existence of an upper bound for the values of $J(\mathcal{N}_1, \mathcal{N}_2, t)$s that holds under any feasible policy (the proof is presented in the online appendix):

LEMMA 3. *Under either one of conditions* (i) *or* (ii), *there exists* U^*, *such that for every feasible solution of* (3)–(7), *we have*

$$J(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant U^* \tag{28}$$

for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \bar{\mathfrak{N}}$.

Unfortunately, the boundedness of $J(\mathcal{N}_1, \mathcal{N}_2, t)$s alone cannot guarantee the existence of the optimal policy. While the supremum of $J(\varnothing, \varnothing, 0)$ is finite, it may only be attainable in the limit, as some of the prices increase to infinity. Therefore, consider a problem of maximizing $J(\varnothing, \varnothing, 0)$, subject to (3)–(7) and an additional constraint

$$p(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant p_{\max}(t) \tag{29}$$

for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$ and some positive function of time $p_{\max}(t)$, $t \in \{0, \ldots, T-1\}$. We will refer to this problem as a *modified problem* in this section. The Lagrangian of the modified problem is

$$\mathscr{L}' = \mathscr{L} - \sum_{(\mathcal{N}_1, \mathcal{N}_2, t)\in\mathfrak{N}} \sigma(\mathcal{N}_1, \mathcal{N}_2, t)[p(\mathcal{N}_1, \mathcal{N}_2, t) - p_{\max}(t)],$$

where $\sigma(\mathcal{N}_1, \mathcal{N}_2, t)$ is the multiplier corresponding to (29). The role of $p_{\max}(t)$ is to make the set of feasible policies bounded. Consider a system of inequalities

$$p_{\max}(t) \geqslant \frac{2}{\epsilon}\left(\sum_{0\leqslant\theta<t} p_{\max}(\theta) + U^*\right), \quad 1 \leqslant t \leqslant T-1, \tag{30}$$

$$p_{\max}(0) \geqslant p_{\max}, \tag{31}$$

where we use $p_{\max}$ and ϵ from condition (i). Any solution to this system is strictly increasing in t because $U^* \geqslant 0$ and $\epsilon < 1$. We let $\bar{p}_{\max}(t)$ be the solution satisfying (30)–(31) as equalities. Note that for any given bounded function of time, there is a solution of (30)–(31) that majorizes it.

Lemma 4. *Suppose that $u(\cdot)$ satisfies condition* (i) *and let $p_{\max}(t)$ be a solution of* (30)–(31). *Then, for any optimal solution of the modified problem,*

(a) *the corresponding multipliers in the KKT conditions are such that $\sigma(\mathcal{N}_1, \mathcal{N}_2, t) = 0$ for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$, and*

(b) *for all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$, $p(\mathcal{N}_1, \mathcal{N}_2, t) > \bar{p}_{\max}(t)$ implies $\pi(\mathcal{N}_1, \mathcal{N}_2, t) = 0$.*

We present the proof of Lemma 4 in the online appendix.

Remark 3. Part (a) of Lemma 4 establishes that any optimal solution to the modified problem satisfies the KKT conditions of the original problem. Part (b) of Lemma 4 essentially claims that all optimal solutions of the modified problem do not violate (29) using $\bar{p}_{\max}(t)$ except perhaps on triples of probability zero. Any such solution can be modified to satisfy (29) without any changes to the value of the objective function or Lagrange multipliers.

The following proposition, proved in the online appendix, states the existence of a solution to the discrete-time expected revenue maximization problem.

Proposition 3. *Under either one of conditions* (i) *or* (ii), *an optimal solution to the discrete-time expected revenue maximization problem exists.*

2.4. Special Case of Free Price Guarantees

In this section, we study the case of a simplified policy space where the fee for the price guarantee is always zero. Such a policy will not produce lower direct expected revenues than a policy with no price guarantees because the price guarantee problem is a relaxation of the no-guarantee problem. On the other hand, a free price guarantee policy will offer a lower bound for the value of the optimal general policy. We first show that the local effects of the price on demand are larger than the local effects of the strike price at any optimal solution to the discrete-time expected revenue maximization problem. Then, in §2.4.1, we discuss the relevance of free price guarantees to current industry practices (e.g., price matching in retail, and nonrefundable, refundable, and partially refundable tickets in the travel industry). Modeling refundable and partially refundable tickets involves incorporating no-shows into the model. Finally, we discuss an analytic lower bound based on fixed policies in §2.4.2.

A zero fee for the price guarantee immediately implies that $v = 1$ and all triples $(\mathcal{N}_1, \mathcal{N}_2, t)$ with nonempty $\mathcal{N}_1$ have probability zero. Therefore, it is convenient to drop $\mathcal{N}_1$ from all variables and write $\mathcal{N}$ instead of $\mathcal{N}_2$. The vector notation for the policy with a zero fee is $\Pi^0 = (p, k, 0)$. The revenue decomposition constraint

$$J(\mathcal{N}, t) = \bar{u}(\Pi^0(\mathcal{N}, t), t)\{p(\mathcal{N}, t) + J(\mathcal{N} \cup t, t+1)\} + (1 - \bar{u}(\Pi^0(\mathcal{N}, t), t))J(\mathcal{N}, t+1) \tag{32}$$

simplifies significantly. Here, it is convenient to use the (p, k) parameterization instead of (p, κ). In these policy variables, the Lagrangian is

$$\begin{aligned}\mathscr{L} = {} & J(\varnothing, 0)\\ & - \sum_{(\mathcal{N},t)\in\mathfrak{N}} \pi(\mathcal{N}, t)\big\{J(\mathcal{N}, t) - J(\mathcal{N}, t+1) - \bar{u}(\Pi^0(\mathcal{N}, t), t)\\ & \qquad \cdot \big[p(\mathcal{N}, t) + J(\mathcal{N} \cup t, t+1) - J(\mathcal{N}, t+1)\big]\big\}\\ & - \sum_{(\mathcal{N}',t')\in\mathfrak{N}^\Delta} \pi(\mathcal{N}', t')\left\{J(\mathcal{N}', t') - \sum_{t\in\mathcal{N}'} z_t(\mathcal{N}')\right\}\\ & - \sum_{\substack{(\mathcal{N}',t')\in\mathfrak{N}^\Delta,\, t\in\mathcal{N}',\\ t<\tau<\min\{t', t+D\}}} \zeta(\mathcal{N}', \tau, t)\\ & \qquad \cdot \big\{z_t(\mathcal{N}') - p(\mathcal{N}' \cap [0, \tau), \tau) + k(\mathcal{N}' \cap [0, t), t)\big\}\\ & - \sum_{(\mathcal{N}',t')\in\mathfrak{N}^\Delta,\, t\in\mathcal{N}'} \zeta(\mathcal{N}', t, t) z_t(\mathcal{N}')\\ & - \sum_{(\mathcal{N},t)\in\mathfrak{N}} \big\{\eta(\mathcal{N}, t)[k(\mathcal{N}, t) - p(\mathcal{N}, t)] - \xi(\mathcal{N}, t)k(\mathcal{N}, t)\big\},\end{aligned} \tag{33}$$

where we denote the Lagrange multipliers of $k(\mathcal{N}, t) \leqslant p(\mathcal{N}, t)$ and $k(\mathcal{N}, t) \geqslant 0$ as $\eta(\mathcal{N}, t)$ and $\xi(\mathcal{N}, t)$, respectively. Again, we see a significant simplification compared to the general case.

The following proposition is proved in the online appendix, using the analysis of the simplified optimality conditions:

Proposition 4. *In any optimal solution, for all $(\mathcal{N}, t)$ such that $P(\mathcal{N}, t) > 0$ and $J(\mathcal{N}, t) > J(\mathcal{N}, t+1)$, we have*

$$\frac{\partial \bar{u}}{\partial k}(\Pi^0(\mathcal{N}, t), t) + \frac{\partial \bar{u}}{\partial p}(\Pi^0(\mathcal{N}, t), t) \leqslant 0. \tag{34}$$

Note that $\partial \bar{u}/\partial k \geqslant 0$, while $\partial \bar{u}/\partial p \leqslant 0$. We then have

$$\left|\frac{\partial \bar{u}}{\partial p}(\Pi^0(\mathcal{N}, t), t)\right| \geqslant \frac{\partial \bar{u}}{\partial k}(\Pi^0(\mathcal{N}, t), t)$$

as a restatement of (34). When restated in this way, Proposition 4 affirms the intuitive perception that any effects of price guarantee on demand are secondary to the effects of price. Given optimality, one should not be able to compensate for a decrease in demand resulting from a marginal increase in price by a marginal readjustment of terms of the price guarantee. The proposition also has a potential computational implication for pricing models where the condition of the form $\partial \bar{u}/\partial k + \partial \bar{u}/\partial p \leqslant 0$ is not satisfied for all values of policy variables. In this case, we may restrict the search for the optimal solution to a subset of policy space where (34) is satisfied.

One can also ask how (34) is related to condition (i) of §2.3, which is proved to be sufficient for the existence of an optimal solution in Proposition 3. We first note that while Proposition 3 does not apply directly to the restricted problem with $f = 0$, its proof can be easily modified to cover this case. Condition (i) can be restated in the (p, k) parameterization for $\bar{u}$ as: There exist $\epsilon \in (0, 1)$ and $p_{\max}$ such that

$$\frac{p}{\bar{u}}\frac{\partial \bar{u}}{\partial p} + \frac{k}{\bar{u}}\frac{\partial \bar{u}}{\partial k} \leqslant -\frac{1}{1-\epsilon}$$

for all $p > p_{\max}$, all $0 \leqslant k \leqslant p$, and $0 \leqslant t \leqslant T$. Suppose that an inequality that is the opposite of (34) holds for all (p, k), $0 \leqslant k \leqslant p$:

$$\frac{\partial \bar{u}}{\partial k} + \frac{\partial \bar{u}}{\partial p} > 0.$$

Then, for all $p \geqslant 0$ and $k = p$, we have

$$\frac{p}{\bar{u}}\frac{\partial \bar{u}}{\partial p} + \frac{k}{\bar{u}}\frac{\partial \bar{u}}{\partial k} > 0.$$

This contradicts the existence of $p_{\max}$, which is required by our sufficient condition.

2.4.1. Price Guarantee-Type Instruments and Mechanisms Used in Industry. A most notable (free) price guarantee mechanism is price matching, which is widely used in the retail industry. In the context of a monopoly market, this is *internal* price matching. That is, a company will reimburse a customer if he/she finds that the price set by the company drops prior to some point in the future. The benefit for the customer is maximized if the price is monitored constantly. Respectively, it is also the worst case for the company in terms of its revenues. Our model maximizes the revenues in this worst case. Thus, the optimal value in our model provides a lower bound for the general case when the customers may or may not monitor the price constantly.

Other practices that, implement price guarantees to some extent are fully and partially refundable tickets in the travel industry. Even so-called nonrefundable tickets have restricted price guarantees that are difficult, but possible, to exercise. If a customer buys a nonrefundable ticket, and the price drops, he/she can go to the airport and have the ticket exchanged for the lower-price ticket for a fee, currently about $100 per ticket. The relevance of free price guarantees presented in the paper to such "nonrefundable" fares can be seen if we consider our model with $f = 0$ and $k = p - \texttt{exchange fee}$. Because the timing of an exchange is chosen by the customer, our model provides a lower bound on the expected revenues for a company.

The case of fully or partially refundable products presents an additional complexity for pricing. The following extension of our model encompasses such situations. It is reasonable to expect that customers who choose to take advantage of refundability also want to fully utilize the flexibility offered by the product. Thus, we assume that returns for a refund always occur at the end of the planning period, and returned items cannot be resold (for example, no-shows in the travel industry). Consider our model with $f = 0$ and $p = k = \texttt{item price}$ in the case of fully refundable products; and $p = \texttt{item price}$ and $k = p - \texttt{restocking fee}$ in case of partially refundable products. Let $q(t)$ be the probability that an item purchased at time t is returned for a refund at the end of the planning horizon. Consider a final state $(\mathcal{N}, t') \in \mathfrak{R}^{\Delta}$. If a customer who purchased at $t \in \mathcal{N}$ chooses to exercise the refund option, the company will have to reimburse the amount $k(\mathcal{N}, t)$. The terminal payments condition (6) assumes the form (in the case of free price guarantees)

$$J(\mathcal{N}, t') = \sum_{t \in \mathcal{N}} [(1 - q(t)) z_t(\mathcal{N}) + q(t) k(\mathcal{N}, t)] \quad \text{for all } (\mathcal{N}, t') \in \mathfrak{R}^{\Delta}. \tag{35}$$

All other components of the model remain unchanged. Again, our model provides a lower bound on the expected revenues for a company.

There is also another reason why it may be beneficial to use price guarantees as a modeling tool. If customers are actively engaged in the practice of exchanging the product at the time of a lower price, then the company suffers losses in revenue. Such losses are not captured by models without price guarantees assuming myopic customer behavior. The reason is that customer behavior involving returns for refund with subsequent repurchase is strategic in its nature. Considering strategic customer behavior greatly increases the overall model complexity. The lower bound provided by our model may serve as a reasonable compromise in practice.

2.4.2. A Fixed-Policy Lower Bound. In this subsection, we consider a lower bound for $J(\varnothing, \varnothing, 0)$ that is very easy to compute. It is provided by the fixed-policy pricing strategy, when the policy vector Π is constant at all $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{R}$. A particularly simple case is that of the zero fee, $\phi = 0$; and price matching, $\kappa = 1$; because then the price guarantee purchase happens with certainty. Such a static pricing policy might generate a poor lower bound, but it can serve as an analytical benchmark (rather than a numerical one, as in the myopic heuristic case considered later) to contrast against pure dynamic pricing. Note that the fixed-policy lower bound is also a benchmark for the myopic lower bound. Next, we show that in some cases this simple static policy with price guarantee performs better than dynamic pricing without guarantee, as in Gallego and van Ryzin (1994). Because our goal in this section is providing a benchmark, we will only consider price guarantees that do not expire until the end of the planning horizon. This case is the most pessimistic one in terms of flexibility in setting the pricing policy because all price guarantees sold throughout the planning horizon continue to affect pricing decisions up to the last time period.

In this analysis we will assume the following:

(1) For each p, $\bar{u}(p,1,0,t)$ is constant with respect to t. A similar assumption will be imposed on $\bar{u}(p,1,1,t)$ (the case of no price guarantee being offered). Therefore, we omit parameter t in further discussion.

(2) The function $\bar{u}(p,\kappa,\phi)$ is continuous and strictly decreasing in p. Therefore, there exist functions $p_0(\cdot)$ and $p_1(\cdot)$ that satisfy the equations $\bar{u}(p_0(\bar{u}'),1,0)=\bar{u}'$ and $\bar{u}(p_1(\bar{u}'),1,1)=\bar{u}'$, respectively.

(3) Both $p_0(\bar{u})\bar{u}$ and $p_1(\bar{u})\bar{u}$ are concave.

(4) For all p, $\bar{u}(p,1,0)>\bar{u}(p,1,1)$.

Let

- p_1^0 and p_0^0 be such that $T\bar{u}(p_1^0,1,1)=Y$ and $T\bar{u}(p_0^0,1,0)=Y$ respectively (if a solution of either equation does not exist, we assume that $p_1^0=-\infty$ or $p_0^0=-\infty$, respectively),
- p_1^* and p_0^* be maximizers of $p\bar{u}(p,1,1)$ and $p\bar{u}(p,1,0)$, respectively, and
- $L^{\text{fixed}}(Y,T)$ be the value of the lower bound obtained by utilizing a fixed-policy triple $(\max\{p_0^0,p_0^*\},1,0)$ in the problem with time horizon T and initial inventory Y.

Without loss of generality, we restrict our further analysis to the case of $Y\leqslant T$. The following lemma (proved in the online appendix) gives a lower bound on $L^{\text{fixed}}(Y,T)$.

Lemma 5. *Let $Y\leqslant T$. Then,*

$$L^{\text{fixed}}(Y,T)\geqslant\begin{cases}Yp_0^0\left(1-\dfrac{1}{2\sqrt{Y}}\right), & p_0^0\geqslant p_0^*,\\[2ex] p_0^*T\bar{u}(p_0^*,1,0)\left(1-\dfrac{1}{2\sqrt{T\bar{u}(p_0^*,1,0)}}\right), & \\ & p_0^0<p_0^*.\end{cases}$$

We next compare the value of $L^{\text{fixed}}(Y,T)$ with $V(0,0)$, the optimal value of the discrete-time version of the model of Gallego and van Ryzin (1994). The discrete-time version of the Bellman equation for that model is

$$V(y,t)=\max_p\bigl\{\bar{u}(p,t)(p+V(y+1,t+1))+(1-\bar{u}(p,t))V(y,t+1)\bigr\}. \quad (36)$$

This equation provides the optimal expected future revenues $V(y,t)$ for $Y-y$ remaining items in the inventory when there are $T-t$ time periods left and the price guarantee tool is not used. The boundary condition is

$$V(Y,t)=V(y,T)=0. \quad (37)$$

The optimal pricing policy $p^*(y,t)$ is the one attaining the maximum in (36). The probability that the purchase happens is equal to $\bar{u}^*(y,t)=\bar{u}(p^*(y,t),t)$.

To perform the comparison, we first need an upper bound on $V(y,t)$. Consider a problem of the form

$$\max \sum_{s=t}^{T-1}p(\bar{u}_s)\bar{u}_s \quad (38)$$

$$\text{s.t.} \sum_{s=t}^{T-1}\bar{u}_s\leqslant Y-y, \quad (39)$$

$$0\leqslant\bar{u}_s\leqslant\lambda,\quad s=t,\ldots,T-1, \quad (40)$$

where $p(\bar{u})$ is either $p_0(\bar{u})$ or $p_1(\bar{u})$. The optimal value of this problem is denoted by $V_0^D(y,t)$ in case $p(\bar{u})=p_0(\bar{u})$, and $V_1^D(y,t)$ in case $p(\bar{u})=p_1(\bar{u})$. Due to the concavity assumption (3) on $p_0(\bar{u})\bar{u}$ and $p_1(\bar{u})\bar{u}$, the problem (38)–(40) is convex.

Lemma 6. (i) $V_\phi^D(0,0)=\min\{Yp_\phi^0,Tp_\phi^*\bar{u}(p_\phi^*,1,\phi)\}$, $\phi=0,1$. *The minimum is attained in Yp_ϕ^0 if $p_\phi^0\geqslant p_\phi^*$, and $Tp_\phi^*\bar{u}(p_\phi^*,1,\phi)$, otherwise.*

(ii) $V_1^D(y,t)\geqslant V(y,t)$ *for all $y<Y$, $t<T$.*

The proof of the above lemma is in the online appendix.

We can now compare the value $L^{\text{fixed}}(Y,T)$ of the lower bound obtained by utilizing a fixed policy with price guarantee with the value $V(0,0)$ of the dynamic pricing model without guarantee. The comparison is done by combining the results of Lemmas 5 and 6. While it is possible to do this comparison in general, the notation is significantly simpler if we only consider the demand models of the form $\bar{u}(p,\kappa,\phi)=a(p)b(\kappa,\phi)$ (where $a(\cdot)$ is continuous). In this case, $p_1^*=p_0^*=\arg\max\{pa(p)\}$, which we denote by p^*. We also have $Y/T=a(p_\phi^0)b(1,\phi)$ and $p_\phi^0=a^{-1}(Y/(Tb(1,\phi)))$, $\phi=0,1$. Because $b(1,1)<b(1,0)$, and $a^{-1}(\cdot)$ is decreasing, it follows that $p_0^0>p_1^0$.

In the proposition below, we compare $V(0,0)$ and $L^{\text{fixed}}(Y,T)$ in the following situations:

(1) Demand is "high," so even without the promotional effect, the expected number of potential customers who could buy at p^* is higher than the inventory.

(2) Demand is "moderate." Without the promotional effect, the expected number of customers who could buy at p^* is below the inventory. It is higher than the inventory when promotion is in effect.

(3) Demand is "low," and the expected number of customers who could buy at p^* is below the inventory even with a promotion.

The proof of the following proposition is in the online appendix:

Proposition 5. *Let $\bar{u}(\cdot)$ satisfy assumptions* (1)–(4) *and be of the form $\bar{u}(p,\kappa,\phi)=a(p)b(\kappa,\phi)$ for some continuous $a(p)$. Then,*

$$L^{\text{fixed}}(Y,T)\geqslant\begin{cases}V(0,0)\dfrac{p_0^0}{p_1^0}\left(1-\dfrac{1}{2\sqrt{Y}}\right), & p^*\leqslant p_1^0,\\[2ex] V(0,0)\dfrac{Yp_0^0}{p^*Ta(p^*)b(1,1)}\left(1-\dfrac{1}{2\sqrt{Y}}\right), & p_1^0<p^*\leqslant p_0^0,\\[2ex] V(0,0)\dfrac{b(1,0)}{b(1,1)}\left(1-\dfrac{1}{2\sqrt{Ta(p^*)b(1,0)}}\right), & p^*>p_0^0.\end{cases}$$

The above statement implies that $L^{\text{fixed}}(Y,T)$ exceeds $V(0,0)$ if

$$\frac{p_0^0}{p_1^0}\left(1-\frac{1}{2\sqrt{Y}}\right)>1 \quad \text{in case } p^*\leqslant p_1^0, \tag{41}$$

$$\frac{Yp_0^0}{p^*Ta(p^*)b(1,1)}\left(1-\frac{1}{2\sqrt{Y}}\right)>1 \quad \text{in case } p_1^0<p^*\leqslant p_0^0, \tag{42}$$

$$\frac{b(1,0)}{b(1,1)}\left(1-\frac{1}{2\sqrt{Ta(p^*)b(1,0)}}\right)>1 \quad \text{in case } p^*>p_0^0. \tag{43}$$

We now apply the obtained results to the demand model given by (17) with constant $v_0(t)$. We show that $L^{\text{fixed}}(Y,T)$ exceeds $V(0,0)$ for some values of the promotional effect α.

Example 3. Let $a(p)=\exp(-p)$, $b(1,1)=\lambda\exp(-\alpha)$, and $b(1,0)=\lambda$. Then, $p^*=1$, $p_0^0=-\ln(Y/T\lambda)=\ln(T\lambda/Y)$, and $p_1^0=-\ln(Y/(T\lambda\exp(-\alpha)))=\ln(T\lambda/Y)-\alpha$. We consider three cases as in the statement of Proposition 5.

Case 1. The "high-demand" case of $p^*\leqslant p_1^0$ requires $\alpha\leqslant\ln(T\lambda/Y)-1$. Condition (41) can be rewritten equivalently as

$$\alpha>\frac{1}{2\sqrt{Y}}\ln\left(\frac{T\lambda}{Y}\right).$$

We see that the range of α that falls in the high-demand case and ensures that the fixed-policy lower bound exceeds the value of a dynamic policy without the price guarantee is from $(1/2\sqrt{Y})\ln(T\lambda/Y)$ to $\ln(T\lambda/Y)-1$. Is this range reasonable? For concreteness, let $\ln(T\lambda/Y)=2$. Then, the range is $[1/\sqrt{Y},1]$. It is nonempty for any value of initial inventory Y. Moreover, the lowest value of α resulting in $L^{\text{fixed}}(Y,T)\geqslant V(0,0)$ is decreasing in Y.

Case 2. The "moderate-demand" case of $p_1^0<p^*\leqslant p_0^0$ requires that

$$p_1^0=\ln\left(\frac{T\lambda}{Y}\right)-\alpha<p^*=1\leqslant\ln\left(\frac{T\lambda}{Y}\right)=p_0^0.$$

From $p_1^0<p^*$, we get $\alpha>\ln(T\lambda/Y)-1$. Condition (42) is equivalent to

$$\alpha>\ln\left(\frac{T\lambda}{Y}\right)-1-\ln\left[\ln\left(\frac{T\lambda}{Y}\right)\left(1-\frac{1}{2\sqrt{Y}}\right)\right].$$

We see that there is a range of α such that this condition and $\alpha>\ln(T\lambda/Y)-1$ are simultaneously satisfied. If we choose $\ln(T\lambda/Y)=1$, this leads to $\alpha>-\ln[1-1/2\sqrt{Y}]$.

Case 3. Finally, we consider the "low-demand" case of $p^*>p_0^0$, which means that $1>\ln(T\lambda/Y)$. Condition (43) is equivalent to

$$\alpha>-\ln\left(1-\frac{1}{2\sqrt{T\lambda e^{-1}}}\right).$$

For the case of $\ln(Y/T\lambda)=0$, we have $Y=T\lambda$, and the resulting lower bound on α is $-\ln(1-1/2\sqrt{Ye^{-1}})$.

Thus, the fixed-policy lower bound $L^{\text{fixed}}(Y,T)$ often exceeds the value of the optimal revenue $V(0,0)$ under the dynamic pricing policy without price guarantees. The smallest guaranteed values of α that produce this effect for specific values of inventory and $\ln(T\lambda/Y)$ are presented in Table 1.

2.5. Myopic Lower-Bounding Heuristic

Note that the number of triples in $\overline{\mathfrak{N}}$ is exponential in Y, which makes exact solution extremely difficult in most practical situations. Therefore, it is important to find efficient heuristic solution methods that provide bounds on $J(\varnothing,\varnothing,0)$. It is particularly useful in practice to have a lower bound. Typically, lower bounds are obtained by imposing additional constraints, which make a modified (bounding) problem easier to solve. In the present case, we seek lower bounds that are greater than the objective value of the pure dynamic pricing problem with no guarantees.

We propose to restrict the pricing policy to be "myopic" in the following sense. For each history of the sales process $(\mathcal{N}_1,\mathcal{N}_2,t)$ we maximize the expected total of the current and future revenues, including price guarantee payments

Table 1. The smallest values of α such that $L^{\text{fixed}}(Y,T)\geqslant V(0,0)$ for specific values of inventory Y and $\ln(T\lambda/Y)$.

	Y		
$\ln(T\lambda/Y)$	10	50	100
2	0.316	0.141	0.100
1	0.253	0.105	0.073
0	0.460	0.180	0.124

for the current and future guarantee sales, but disregarding payments for all *previously sold* guarantees. Because the history of the quote process is ignored, such a myopic policy depends only on time t and the total number of items sold y. Moreover, we assume that the future policy is fixed, and, for each y and t, modify only current policy variables. Thus, such a myopic pricing policy can be found very efficiently using dynamic programming.

Denote the policy vector for (y, t) as $\Pi(y, t) = (p(y, t), k(y, t), f(y, t))$ and the expected total of the current and future revenues as $L(y, t)$. Suppose that a sale happens with probability $\bar{u}(\Pi(y, t), t)$. Then, the company receives the current payment $p(y, t)$ and future revenues $L(y+1, t+1)$. Given that a sale occurs, the company also receives a fee $f(y, t)$ with probability $v(\Pi(y, t), t)$ and, if necessary, pays corresponding price guarantee compensation. The expected payments associated with this price guarantee can always be computed because the future prices are fixed when we optimize $\Pi(y, t)$. We can efficiently compute these payments using the distribution of the minimum future price over the duration of the guarantee. Let $p_{\min}(y, t)$ denote the (random) minimum price occurring in realizations of the price process starting from time t and up to D periods into the future, and initial sales y under the pricing policy $p(y', t')$, $y' \geqslant y$, $t' \geqslant t$. We will first discuss the case when price guarantees do not expire before the end of the planning period $(D = T)$. Then, the probability distribution of $p_{\min}(y, t)$ can be obtained recursively from those of $p_{\min}(y, t+1)$ and $p_{\min}(y+1, t+1)$. For convenience of notation, let $p_{\min}(y, T) = p_{\min}(Y, t) = +\infty$ for all y, t. Then,

$$P[p_{\min}(y, t) \quad \text{first occurs at } (y, t)] = \bar{u}(y, t) P[p_{\min}(y+1, t+1) \geqslant p(y, t)] + (1 - \bar{u}(y, t)) P[p_{\min}(y, t+1) \geqslant p(y, t)], \tag{44}$$

and, for all $y' \geqslant y$, $t' \geqslant t$,

$$P[p_{\min}(y, t) \text{ first occurs at } (y', t')] = \begin{cases} \bar{u}(y, t) P[p_{\min}(y+1, t+1) \text{ first occurs at } (y', t')] \\ \quad + (1 - \bar{u}(y, t)) P[p_{\min}(y, t+1) \text{ first} \\ \quad \text{occurs at } (y', t')] \quad \text{if } p(y, t) > p(y', t'), \\ 0, \qquad \text{otherwise.} \end{cases} \tag{45}$$

The set of (y', t') such that $P[p_{\min}(y, t)$ first occurs at $(y', t')] > 0$ is typically quite sparse among all possible $y' \geqslant y$, $t' \geqslant t$.

When we optimize $\Pi(y, t)$, we already know the distribution of $p_{\min}(y+1, t+1)$, which is found using the future pricing policy. Given this distribution, the expected price guarantee payments for a strike price k are

$$Z(k, y, t) = -\sum_{(y', t')} (k - p(y', t'))^+ \cdot P\big(p_{\min}(y+1, t+1) \text{ first occurs at } (y', t')\big). \tag{46}$$

With this notation, we compute $L(y, t)$ using DP:

$$L(y, t) = \max_{p, k, f} \big\{ \bar{u}(p, k, f, t) \big[p + L(y+1, t+1) + v(p, k, f, t)(f + Z(k, y, t)) \big] + (1 - \bar{u}(p, k, f, t)) L(y, t+1) \big\}. \tag{47}$$

Values of p, k, f that attain the maximum in (47) provide us with an optimal myopic policy at (y, t). Once the policy for $y = 0, \ldots, Y-1$ and the current t is known, we can compute the distributions of $p_{\min}(y, t)$ for all $y = 0, \ldots, Y-1$ and proceed with computation of $L(y, t-1)$s, etc. The boundary condition for (47) used to initialize this computation is

$$L(y, T) = L(Y, t) = 0, \quad y = 0, \ldots, Y-1, \; t = 0, \ldots, T. \tag{48}$$

In case of price guarantees with an expiration date $(D < T)$, we need to change the computation of the distribution of $p_{\min}(y, t)$. We can no longer use the distributions of $p_{\min}(y, t+1)$ and $p_{\min}(y+1, t+1)$, and we have to restart the computation from scratch for every time period t. While this leads to an additional computational overhead, the actual procedure is quite similar. All we need to do is to set $p_{\min}(y', t') = +\infty$ for all y' and $t' \geqslant \min\{T, t+D\}$, and $p_{\min}(Y, t') = +\infty$ for all t'. We then redo all recursive computations for $p_{\min}(y', t')$ for all y' and $t \leqslant t' < \min\{T, t+D\}$ using expressions (44)–(45).

We have constructed a policy that is feasible for the original discrete-time expected revenue maximization problem. The expected price guarantee payments included in $Z(k(y, t), y, t)$ are precisely the expected payments associated with a sale of the price guarantee with the optimal myopic strike price $k(y, t)$ occurring at time t, and the number of sales y. All price guarantee payments for the current and future price guarantee sales are accounted for in $L(y, t)$. For $y = 0$, $t = 0$ there are no "past" sales of price guarantees. Thus, we have:

Proposition 6. *$L(0, 0)$ computed according to (47) is a lower bound for the optimal value of $J(\varnothing, \varnothing, 0)$.*

3. Numerical Results

3.1. Exact Model for Three Items and Exponential Demand

In this section, we first compare the performance of our discrete-time model for three items to the one obtained by discretization (36)–(37) of the model in Gallego and van Ryzin (1994). The value of $V(0, 0)$ computed using (36)–(37) represents the optimal expected revenues in the discrete-time model when a price guarantee tool is unavailable. For the three-item case, the dimensions of the expected revenue maximization problem are relatively small, and it can be solved using standard NLP solvers without resorting to heuristics. Larger instances (more than three items) in the exact form become computationally intractable as the number of time intervals increases. The demand model in (36) is chosen to be exponential and independent of time:

$$\bar{u}(p,t) = \bar{\lambda}\exp(-p), \quad (49)$$

while for our model we use the functional form for u and v as described in Example 2, which is a natural extension of (49). For the purpose of this comparison, we keep the quantity

$$\bar{\lambda} = \lambda\exp(-\alpha)$$

fixed. The meaning of $\bar{\lambda}$ is an effective arrival probability which is "depressed" by $\exp(-\alpha)$ due to the absence of price guarantees. This normalization of arrival probability is selected so that the optimal value $V(0,0)$ of the model without price guarantees remains constant for different values of α. If we fixed λ instead, then $V(0,0)$ would have to be recomputed for every α. This normalization does not constitute an additional assumption on the demand model and is only made to set up numerical experiments in a consistent fashion.

We show the percentage of improvement in the optimal $J(\varnothing,\varnothing,0)$ over $V(0,0)$ for various values of $\bar{\lambda}$ and α. Other parameters of the model are selected so that

$$\beta = \beta' = 2,$$

$$\gamma = \gamma' = 2,$$

$$\delta = 10,$$

$$\rho = 10,$$

$$v_0(t) = \frac{T-t-1}{T-1},$$

and are kept constant throughout the experiment.

When the demand function is of the form (49), it is especially easy to solve (36) because $(\partial\bar{u}/\partial p)(p,t) = -\bar{u}(p,t)$. Denote the right-hand side of (36) by $G(p)$ and use the necessary first-order condition

$$G'(p) = \frac{\partial\bar{u}}{\partial p}(p,t)[p + V(y+1,t+1) - V(y,t+1)] + \bar{u}(p,t) = 0$$

to find the maximum with respect to p. This equation yields the unique solution

$$p^*(y,t) = 1 + V(y,t+1) - V(y+1,t+1),$$

$$u^*(y,t) = \bar{\lambda}\exp(-p^*(y,t)),$$

$$V(y,t) = \bar{u}^*(y,t)(p^*(y,t) + V(y+1,t+1)) + (1-\bar{u}^*(y,t))V(y,t+1).$$

We remark that the first-order condition is sufficient because $G'(p) > 0$ for $p < p^*(y,t)$ and $G'(p) < 0$ for $p > p^*(y,t)$. In the case of $T = 15$, $Y = 3$ this simple computation leads to the following values of $V(0,0)$:

$\bar{\lambda}$	0.1	0.2	0.4
$V(0,0)$	0.55005	1.08361	2.03986

.

We compare these numbers to the optimal value of $J(\varnothing,\varnothing,0)$ for the three-item model with a price guarantee and $T = 15$. One potential criticism of our model with $Y = 3$ is that the offer of the price guarantee for the last item assumes that a customer is not aware that there is only one item left. For the last item, the company can offer the best terms of the price guarantee while being sure that it will not incur any price guarantee payments on this offer. In a situation of a "small market" when there are only three items offered for sale, the assumption that a customer has no knowledge of the seller's inventory may not be realistic. Accordingly, in our comparison we use a restricted model in which the price guarantee is unavailable after the second sale.

The computations were carried out with IPOPT, which is a general-purpose nonlinear solver. The model was implemented in the AMPL modeling language and submitted to IPOPT through the NEOS server (see Czyzyk et al. 1998, Gropp and Moré 1997, Dolan 2001). The results of our comparison appear in Table 2 as the percentage of additional revenues recovered by the three-item model versus the expected revenues $V(0,0)$ in the absence of price guarantees. The table shows the total additional revenues as well as its components: the revenues from extra sales and the revenues from fees. The price guarantee payments are not shown in the table because they are near zero in all terminal states of sufficiently high (larger than 0.0001) probability. Thus, in this example, the optimal policy tends to avoid lowering prices so far that price guarantee payments must be made. Observe that we can always recover some additional revenue, even for the case of $\alpha = 0$ (the absence of promotional effect of the price guarantee), because we still receive a fee for selling guarantees. For realistic values of α in the interval from 0.02 to 0.1, the percentage recovered is in the range of 2.4% to 6.6%, with smaller values corresponding to larger values of $\bar{\lambda}$. The values of $\bar{\lambda}$ are selected so that the average number $\bar{\lambda}T$ of quotes requested takes values of 1.5, 3, and 6 (lower than, comparable to, and higher than the initial inventory of $Y = 3$). It is interesting from the managerial point of view to examine the contribution of collected fees to the optimal expected revenues. The fees constitute the entire increase in revenues in the case of no promotional effect ($\alpha = 0$). Their contribution drops to about half of the increase for moderate values of the promotional effect ($\alpha = 0.05$) and a virtually negligible fraction of the increase for the largest promotional effect considered ($\alpha = 0.5$). Moreover, because the percentages in the table are taken relative to $V(0,0)$, which remains constant for each column of the table, the absolute amount of collected fees starts to decrease for larger values of the promotional effect. An intuitive explanation is that the reduction in fees enhances the positive effect of price guarantees on demand and generates additional sales. When the promotional effect is large, the company prefers to generate revenues from additional sales rather than fees. To further examine the contribution of fees, we present the increase in the optimal expected revenues in the case of free

Table 2. Additional revenues recovered (total, from extra sales, and from fees) as a percentage of revenues in the absence of price guarantees in the three-item case with exponential demand model.

	Total			From extra sales			From fees		
	$\bar{\lambda}$ (%)								
α	0.1	0.2	0.4	0.1	0.2	0.4	0.1	0.2	0.4
0.00	2.0	1.9	1.7	0.0	0.0	0.0	2.0	1.9	1.7
0.01	2.5	2.3	2.0	0.4	0.4	0.3	2.0	1.9	1.7
0.02	2.9	2.7	2.4	0.9	0.8	0.6	2.0	1.9	1.7
0.05	4.2	4.0	3.5	2.2	2.1	1.7	2.0	1.9	1.7
0.10	6.6	6.2	5.3	4.6	4.3	3.6	2.0	1.9	1.7
0.20	11.7	10.9	9.2	9.8	9.2	7.6	1.9	1.8	1.6
0.50	29.5	27.2	22.2	28.0	25.8	20.9	1.5	1.4	1.3

price guarantees (no fees) in Table 3. The increase in revenues is essentially the same as in the general model (with possible nonzero fees for price guarantees) for larger values of the promotional effect $\alpha \in \{0.2, 0.5\}$. We may conclude that the problem with free price guarantees provides a good lower bound for the general problem when the promotional effect is large. For moderate values of $\alpha \in \{0.05, 0.1\}$, we receive approximately half of the increase of the general model.

Another interesting aspect of the policy is how the price (or other policy components) changes after sales occur. We display these changes using a treelike graph of policy components p, κ, ϕ. Figures 2 and 3 show the expected future revenues J and the policy components at nodes of the subtrees of possible sample paths for $T = 15$, $Y = 3$, $\lambda = 0.4$, and $\alpha = 0.1$. These subtrees of depth 4 are rooted at the states of the form $(\varnothing, \varnothing, t)$ with starting times $t = 0, 4$. The arcs drawn by solid lines represent sales with price guarantees; dashed lines—sales without price guarantees; and dotted lines—no sales. Numbers on the arcs represent transition probabilities. Short vertical lines represent terminal nodes. The absence of κ and ϕ values for the nodes where only one item is left reflects the fact that the price guarantee is not offered for the last item. We see that there is an overall tendency for prices to decrease in the absence of sales. The price at $(\varnothing, \varnothing, t)$ approaches one (the maximizer of the single-period revenue pe^{-p}) toward the end of the planning horizon. Also, the prices increase after sales occur. For a given node of the tree, the increase after a sale with price guarantee is no larger than after a sale without price guarantee. The strike price ratio κ is equal to one (price matching) in all states where price guarantees are offered. The fee ratio ϕ tends to increase along the paths in the tree. Finally, the expected future revenues J behave according to Proposition 1.

Table 3. Additional revenues recovered as a percentage of optimal expected revenues in the absence of price guarantees in the three-item case with free price guarantees and exponential demand model.

	$\bar{\lambda}$ (%)		
α	0.1	0.2	0.4
0.00	0.0	0.0	0.0
0.01	0.5	0.5	0.3
0.02	1.0	0.9	0.7
0.05	2.5	2.3	1.9
0.10	5.1	4.8	3.9
0.20	10.5	9.8	8.1
0.50	29.0	26.7	21.6

3.2. Customer Perception of the Policy

It is interesting to consider how the introduction of a price guarantee impacts customer perception of price changes. For this experiment, we consider a notion of *regret*, defined as the difference of the amount paid by a customer (net of price guarantee refunds, if any) and the best deal this customer could have received based on the observed history after the initial transaction. The best deal is the purchase at the lowest observable price after the initial transaction. Indeed, buying a price guarantee at the lowest price is redundant because it would not generate any guarantee payments. Thus, the best deal cannot include a price guarantee. Specifically, let the history be given by a terminal triple $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}^{\Delta}$. Denote the minimum price since sale at $\tau \in \mathcal{N}_1 \cup \mathcal{N}_2$ by

$$p_{\min}(\mathcal{N}_1, \mathcal{N}_2, t, \tau) = \min_{t' \in [\tau, t)} \{p(\mathcal{N}_1 \cap [0, t'), \mathcal{N}_2 \cap [0, t'), t')\}.$$

The regret $r(\mathcal{N}_1, \mathcal{N}_2, t, \tau)$ for a customer at τ is defined as

$$r(\mathcal{N}_1, \mathcal{N}_2, t, \tau) = \begin{cases} p(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau) - p_{\min}(\mathcal{N}_1, \mathcal{N}_2, t, \tau) \\ \qquad\qquad \text{if } \tau \in \mathcal{N}_1, \\ p(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau) + f(\mathcal{N}_1 \cap [0, \tau), \\ \quad \mathcal{N}_2 \cap [0, \tau), \tau) - \max\{p_{\min}(\mathcal{N}_1, \mathcal{N}_2, t, \tau), \\ \quad k(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau)\} \quad \text{if } \tau \in \mathcal{N}_2. \end{cases}$$

Figure 2. Expected future revenues J and policy variables p, κ, ϕ as functions of a node of the sample-path subtree rooted at $(\varnothing, \varnothing, 4)$ for $T = 15$, $Y = 3$, $\lambda = 0.4$, and $\alpha = 0.1$.

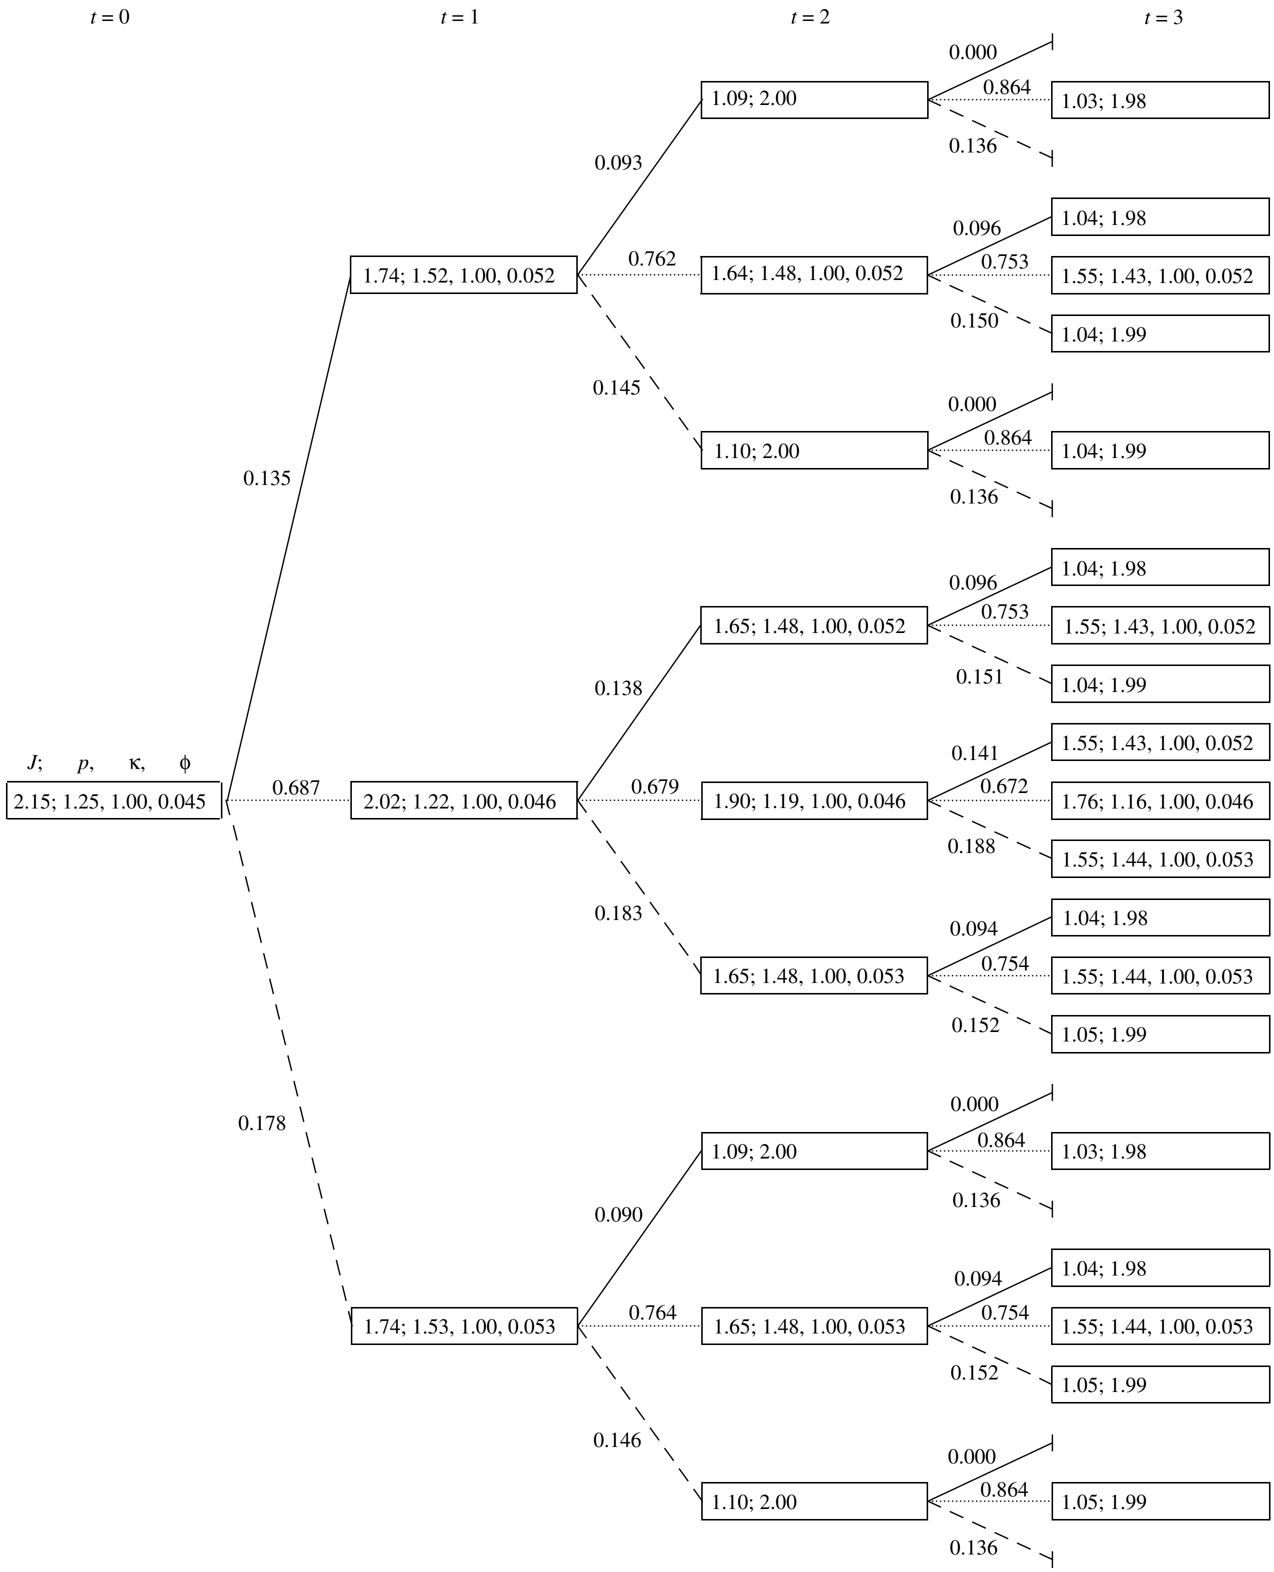

Note. Solid and dashed lines correspond to sales with and without price guarantee, respectively; transition probabilities are given next to the arcs.

Figure 3. Expected future revenues J and policy variables p, κ, ϕ as functions of a node of the sample-path subtree rooted at $(\varnothing, \varnothing, 4)$ for $T = 15$, $Y = 3$, $\lambda = 0.4$, and $\alpha = 0.1$.

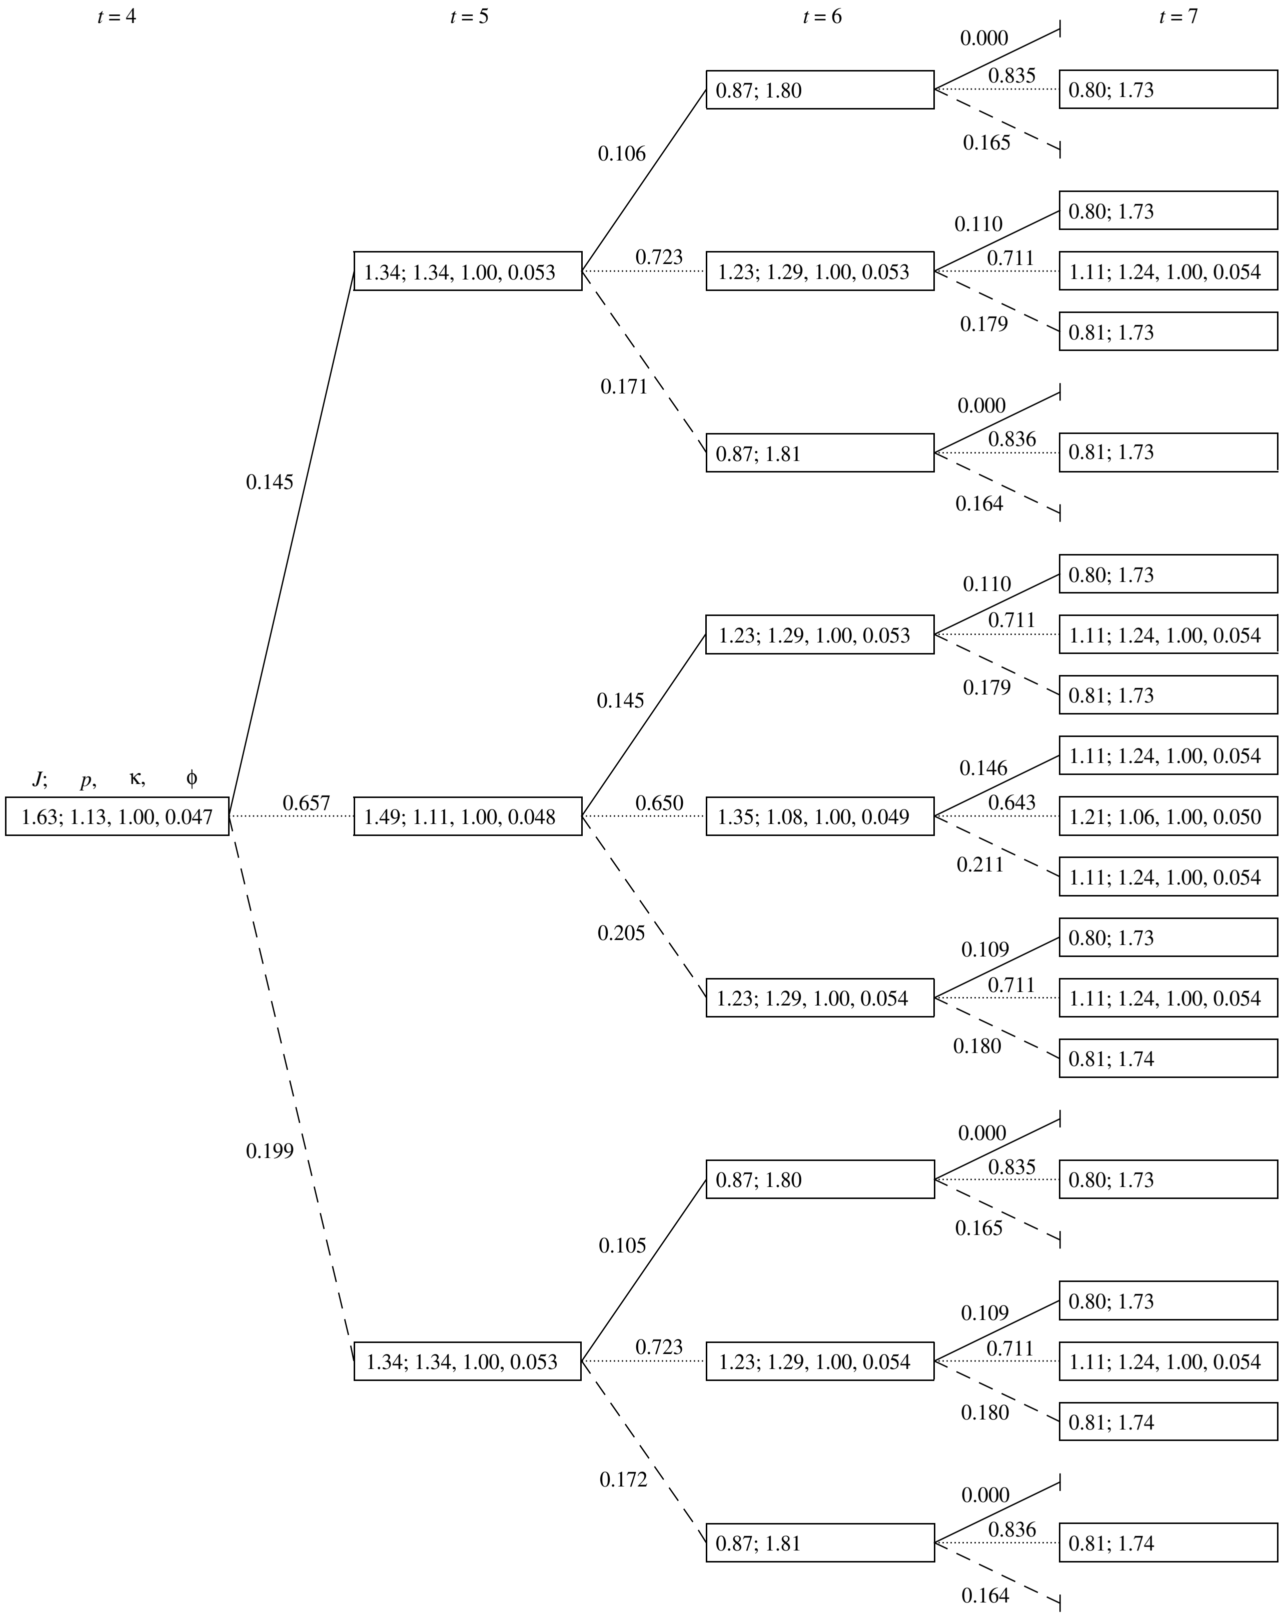

Note. Solid and dashed lines correspond to sales with and without price guarantee, respectively; transition probabilities are given next to the arcs.

We use two complementary ways to aggregate customer regret. The first one takes an a posteriori view. It tries to evaluate how a customer who made a purchase perceives (at the end of the planning horizon) the initial transaction after observing all subsequent prices. An evaluation of a price path is feasible in any market where prices are public information. The average regret then corresponds to a long-term time-averaged perception of policy realizations by the customers.

In contrast, an a priori view considers the expected regret of a customer who has just made a purchase and is observing the entire history of the sales process up to this time. This approach assumes that the customers can compute the probabilities of future outcomes of the system. Such an assumption is unnecessary if we take the a posteriori view because the resulting average regret values are interpreted as long-term averages.

Next, we formally describe these two approaches and report numerical experience gaging customer perception under different values of model parameters.

A Posteriori View. To evaluate the overall perception of the pricing policy, we average the regrets of customers sampled according to the following procedure:

(1) Sample a terminal history triple $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{R}^\Delta$ with respective probability $P(\mathcal{N}_1, \mathcal{N}_2, t)$.

(2) Sample one of the customers in $\mathcal{N}_1 \cup \mathcal{N}_2$ with probability $1/(|\mathcal{N}_1| + |\mathcal{N}_2|)$.

Taking the expected value of regret for customers results in the *average regret* of

$$\bar{r} = \frac{\sum_{(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{R}^\Delta} P(\mathcal{N}_1, \mathcal{N}_2, t) \sum_{\tau \in \mathcal{N}_1 \cup \mathcal{N}_2} \frac{r(\mathcal{N}_1, \mathcal{N}_2, t, \tau)}{|\mathcal{N}_1| + |\mathcal{N}_2|}}{1 - P(\varnothing, \varnothing, T)}.$$

This quantity corresponds to a long-term average perception of policy realizations by the customers.

In the case of $T = 15$, $Y = 3$, we compute the average regret under the optimal policy for the model of Gallego and van Ryzin (1994), $\bar{r}_0$, and compare it to the average regret under the optimal policy utilizing price guarantees. For each value of $\bar{\lambda} \in \{0.1, 0.2, 0.4\}$, the average regret values for different values of α and a single value $\bar{r}_0$ appear in Table 4. Note that there is an overall tendency for regret to decrease with an increase in the promotional effect α. For low and average demand ($\bar{\lambda} \in \{0.1, 0.2\}$), the average regret with the policy utilizing price guarantees is usually higher than without price guarantees. For the high-demand case of $\bar{\lambda} = 0.4$, the average regret with the price guarantees is lower.

We also compare average regret of customers making a purchase with and without price guarantee in all states where both options were available. Given a triple $(\mathcal{N}_1, \mathcal{N}_2, t)$, the availability of both options is defined as $0.0001 < v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t) < 0.9999$ for the purpose of this comparison. The results appear in Table 5. We observe that regrets for sales without price guarantees (referred to as sale Type 1 in the table) are relatively stable in the entire range of α for every $\bar{\lambda}$. On the other hand, for sales with price guarantees (referred to as Type 2), the regrets decrease significantly for larger values of α. Comparison of different sale types for the same α and $\bar{\lambda}$ shows that regret is generally higher for sales with price guarantees when $\bar{\lambda}$ is 0.1 or 0.2. For $\bar{\lambda} = 0.4$, the regrets for sales with price guarantees become smaller. This result can be explained by the fact that a customer will observe higher price fluctuations after making a purchase when the overall demand level is higher and the company has more flexibility in setting the price. A price guarantee limits the amount of perceived fluctuations either by providing guarantee payments or by limiting the range in which the company varies the price.

Table 4. Average a posteriori regret $\bar{r}$ for the case of $T = 15$, $Y = 3$.

	$\bar{\lambda}$		
α	0.1	0.2	0.4
0.00	0.0235	0.0361	0.0757
0.01	0.0239	0.0364	0.0749
0.02	0.0239	0.0361	0.0738
0.05	0.0235	0.0351	0.0707
0.10	0.0226	0.0334	0.0663
0.20	0.0204	0.0298	0.0588
0.50	0.0144	0.0206	0.0396
$\bar{r}_0$	0.0045	0.0234	0.0852

A Priori View. The expected regret of a customer who made a purchase at time t and is currently observing a particular history $(\mathcal{N}_1, \mathcal{N}_2, t+1) \in \mathfrak{R}$, $t \in \mathcal{N}_1 \cup \mathcal{N}_2$ at time $t+1$ is

$$\tilde{r}(\mathcal{N}_1, \mathcal{N}_2, t+1) = \sum_{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{R}^\Delta(\mathcal{N}_1, \mathcal{N}_2, t+1)} r(\mathcal{N}_1', \mathcal{N}_2', t', t) \frac{P(\mathcal{N}_1', \mathcal{N}_2', t')}{P(\mathcal{N}_1, \mathcal{N}_2, t+1)},$$

where $\mathfrak{R}^\Delta(\mathcal{N}_1, \mathcal{N}_2, t+1)$ is the set of sample paths consistent with a history $(\mathcal{N}_1, \mathcal{N}_2, t+1)$ up to time t (defined

Table 5. Average a posteriori regret for sales without (Type 1) and with (Type 2) price guarantees in the states with both options available.

	$\bar{\lambda}$					
	0.1		0.2		0.4	
	Sale type					
α	1	2	1	2	1	2
0.00	0.0040	0.0727	0.0225	0.0746	0.0834	0.0836
0.01	0.0047	0.0694	0.0233	0.0712	0.0835	0.0805
0.02	0.0047	0.0663	0.0234	0.0681	0.0832	0.0774
0.05	0.0049	0.0587	0.0235	0.0605	0.0825	0.0696
0.10	0.0051	0.0491	0.0239	0.0508	0.0820	0.0596
0.20	0.0056	0.0359	0.0247	0.0377	0.0822	0.0460
0.50	0.0077	0.0179	0.0279	0.0187	0.0833	0.0242

Figure 4. Expected a priori regret $\hat{r}(t)$ as a function of time.

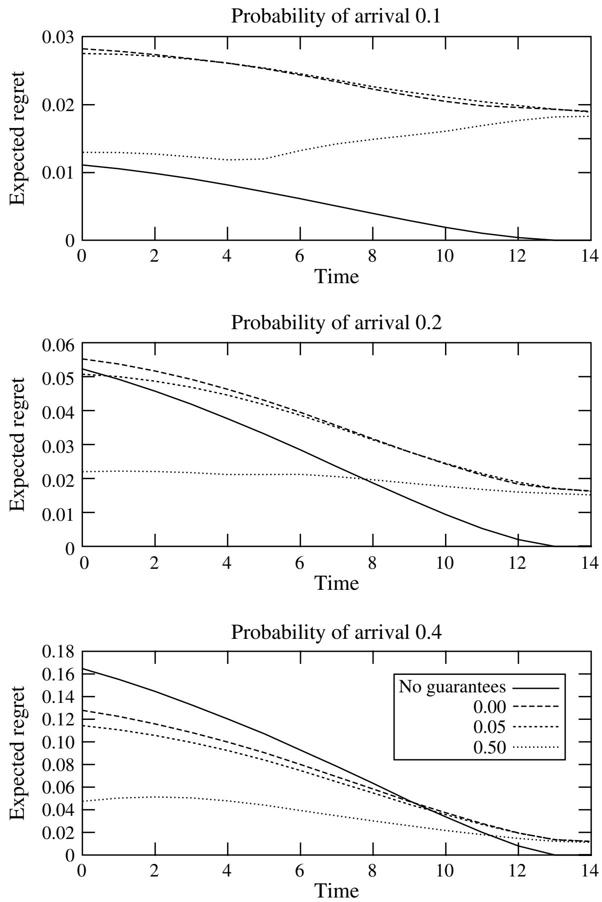

Note. Solid lines correspond to the case of no price guarantees; and others to the case with price guarantees and promotional effect $\alpha = 0, 0.05, 0.5$.

by Equation (50) in the online appendix). Here, we assume that customers can compute $P(\mathcal{N}_1', \mathcal{N}_2', t')$.

The average expected regret at time t conditional on a sale occurrence at time $t \in \mathcal{N}_1 \cup \mathcal{N}_2$ can be defined as

$$\hat{r}(t) = \frac{\sum_{(\mathcal{N}_1, \mathcal{N}_2, t+1) \in \mathfrak{N}:\, t \in \mathcal{N}_1 \cup \mathcal{N}_2} \tilde{r}(\mathcal{N}_1, \mathcal{N}_2, t+1) P(\mathcal{N}_1, \mathcal{N}_2, t+1)}{\sum_{(\mathcal{N}_1, \mathcal{N}_2, t+1) \in \mathfrak{N}:\, t \in \mathcal{N}_1 \cup \mathcal{N}_2} P(\mathcal{N}_1, \mathcal{N}_2, t+1)}$$

$$= \Bigg(\sum_{(\mathcal{N}_1, \mathcal{N}_2, t+1) \in \mathfrak{N}:\, t \in \mathcal{N}_1 \cup \mathcal{N}_2} \; \sum_{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}^{\Delta}(\mathcal{N}_1, \mathcal{N}_2, t+1)} r(\mathcal{N}_1', \mathcal{N}_2', t', t) \cdot P(\mathcal{N}_1', \mathcal{N}_2', t') \Bigg) \cdot \Bigg(\sum_{(\mathcal{N}_1, \mathcal{N}_2, t+1) \in \mathfrak{N}:\, t \in \mathcal{N}_1 \cup \mathcal{N}_2} P(\mathcal{N}_1, \mathcal{N}_2, t+1) \Bigg)^{-1}.$$

The results of these calculations for $Y = 3$, $T = 15$ are presented in Figure 4. It contains three subfigures corresponding to probabilities of arrival $\bar{\lambda} = 0.1$, 0.2, and 0.4. A solid line in each subfigure shows $\hat{r}(t)$ when price guarantees are not offered. Three remaining graphs in each subfigure represent $\hat{r}(t)$ when price guarantees are offered in case of promotional effect $\alpha = 0, 0.05, 0.5$.

For the smaller value of $\bar{\lambda} = 0.1$, an introduction of price guarantees increases the expected regret to a level above the one found in the absence of price guarantees. However, this increase gets smaller for larger values of α. Also, for larger values of $\bar{\lambda}$ the expected regret becomes smaller compared to the case of no price guarantees. The change of $\hat{r}(t)$ in time suggests that the introduction of price guarantees tends to increase the regret of customers making a purchase closer to the end of the planning horizon. This makes the graph of regret look flatter, and may even reverse the slope of $\hat{r}(t)$. The slope of $\hat{r}(t)$ is negative in the case of no price guarantees.

To conclude, both a priori and a posteriori views of consumer perception of a policy show that customers perceive a policy with price guarantees more positively if the promotional effect and/or arrival probabilities are higher.

3.3. Exact Model for Other Forms of Demand Function

We also performed the experiments for the exact three-item model with the demand models other than exponential. In particular, we ran the experiment for the extended power

$$\bar{u}(p, \kappa, \phi, t) = \frac{\lambda(1 + \alpha\kappa^{\beta'}(1 - \phi)^{\gamma'} v_0(t))}{(1 + \alpha)(1 + p)^{\psi}}$$

and the extended linear

$$\bar{u}(p, \kappa, \phi, t) = \frac{\lambda(1 + \alpha\kappa^{\beta'}(1 - \phi)^{\gamma'} v_0(t))(1 - p/p_{\max})^+}{1 + \alpha}$$

demand models. These are modifications of standard power and linear models

$$\bar{u}(p, t) = \bar{\lambda}(1 + p)^{-\psi},$$
$$\bar{u}(p, t) = \bar{\lambda}(1 - p/p_{\max})^+.$$

Again, we assume that in the absence of the price guarantee the demand is depressed: $\bar{\lambda} = \lambda/(1 + \alpha)$. The form of v and the values for all parameters also used in the experiments for the exponential model are the same. New parameters have values $\psi = 2$, $p_{\max} = 2$. The improvements in revenue previously found for the exponential demand model persist for different demand functions. Moreover, we found that qualitative features of policies and results with respect to average regrets are quite similar.

3.4. Myopic Lower-Bound Heuristic

We now discuss the performance of the myopic lower-bounding heuristic of §2.5 as defined by (47). As a benchmark, we use the example of Gallego and van Ryzin (1994). In this example, Y ranges between 1 and 20, the demand is exponential, and $\bar{\lambda} e^{-1} T = 10$. We use a discretization with $T = 50$, which results in $\bar{\lambda} = e/5$. All parameters, except for α, are exactly as in the three-item experiment with exponential demand model. As with the three-item case, we assumed that the company will not offer the price guarantee for the last item. The price guarantee is also not offered in the last time period $t = 49$. Figure 5 presents the lower

Figure 5. Percentage of additional revenues recovered by the lower-bounding heuristic compared to the example of Gallego and van Ryzin (1994) for various promotional effects α.

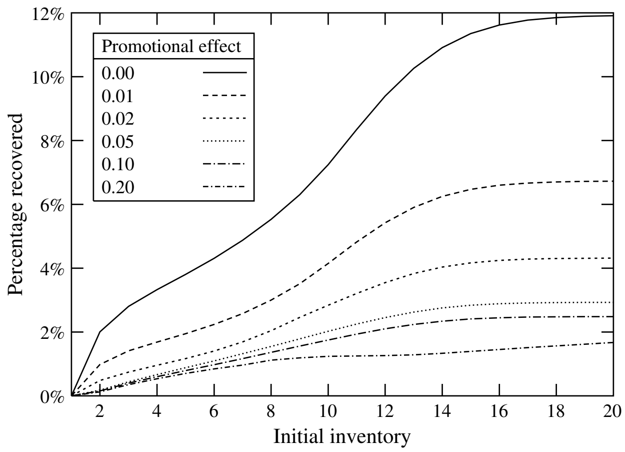

bound on the percentage of additional revenues recovered for different values of α and initial inventory Y. Due to the dynamic programming formulation of the myopic heuristic, the solution is obtained in a single run for all values of Y in the range from 1 to 20. Note that for values of α in the range $[0.05, 0.2]$ and initial inventories between 10 and 20, the percentage of additional revenues recovered is between 2.84% and 11.91%. We also present the profile of the optimal myopic policy for $\alpha = 0.2$ in terms of time and the number of remaining items in Figure 6 (the strike price and the fee are shown for inventory levels from 2 to 20 and time periods from 0 to 48). The strike price matches the price in the beginning of the planning period for smaller inventories. It subsequently changes to a value very close to the minimal observed price, which is near one. The price drops shortly after the strike price drop and stays near one until the end of the planning period. The fee charged is larger in the beginning, when we have price matching for prices greater than one. The fee decreases after the strike price drops to one.

The myopic heuristic can be compared to the exact model on some small problem instances. Table 6 provides the gap as the percentage ratio of the form

$$\frac{J(\varnothing, \varnothing, 0) - L(0, 0)}{J(\varnothing, \varnothing, 0) - V(0, 0)} \times 100\%$$

for $T = 15$, $Y = 2, 3$, $\bar{\lambda} = 0.1$, 0.2, 0.4, and a range of α. The gap generally increases for larger values of $\bar{\lambda}$, but decreases in α and Y.

Finally, we show that customers may receive nonzero price guarantee payments under some scenarios when the myopic policy is utilized. We identify those (y, t)-pairs for which the absolute value of the expected price guarantee payment (46) for the current sale of the price guarantee is greater than or equal to $10^{-4}L(y, t)$. These pairs are shaded in Figure 7 for $\alpha = 0.2$. The vertical axis represents the remaining inventory $Y - y$. These are the same pairs as those for which the strike price is greater than one. This pattern is intuitively clear because the price tends to increase for lower values of the remaining inventory. The strike price for the same pairs is equal to the sale price to maximize the effect of the price guarantee on demand. This is beneficial for revenue maximization because the price is unlikely to drop starting from these states. Certainly, there is a small chance that only a few more items are sold, and the price decreases significantly. In the latter case, the company incurs price guarantee payments.

Figure 6. Price, strike price, and fee components of myopic policy for different times and values of remaining inventory.

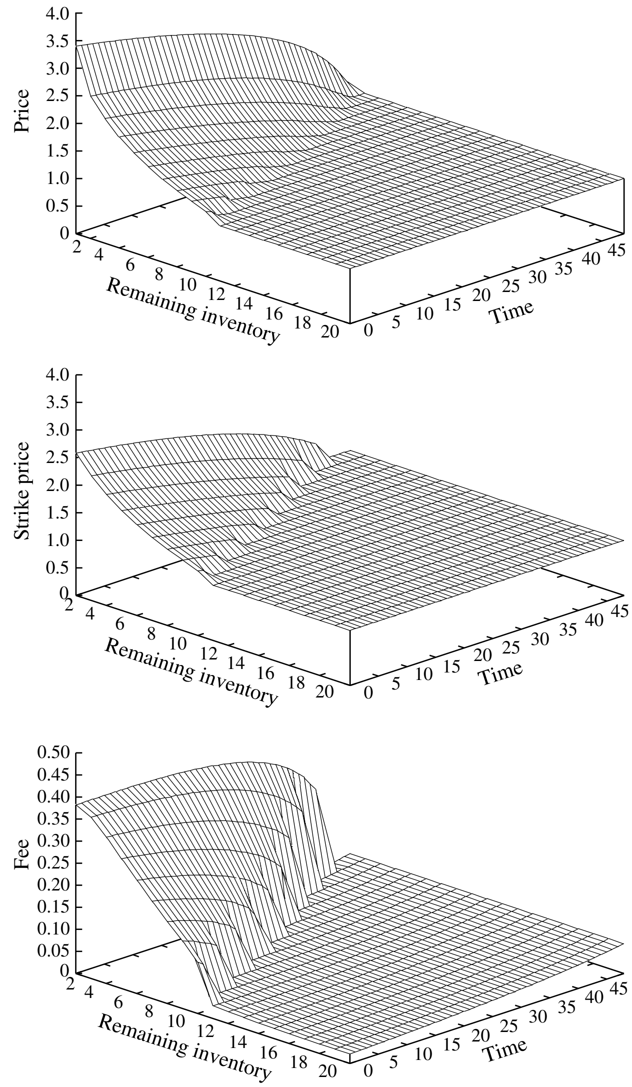

4. Conclusions and Future Research

This paper presents an extension of a classical dynamic pricing model by explicitly including a price guarantee tool

Table 6. The difference between the value of the exact model and the myopic lower bound as a percentage of the difference between the value of the exact model and the optimal expected revenues in the absence of price guarantees.

	$Y=2$			$Y=3$		
	$\bar{\lambda}$ (%)					
α	0.1	0.2	0.4	0.1	0.2	0.4
0.00	1.0	10.5	36.5	6.8	9.7	22.3
0.01	8.4	19.3	37.5	5.6	8.6	22.6
0.02	7.4	19.3	38.3	4.8	7.8	22.7
0.05	5.4	18.9	39.8	3.3	6.1	23.0
0.10	3.9	18.0	35.1	2.2	4.7	22.2
0.20	2.9	14.2	28.8	1.3	3.4	18.4
0.50	2.6	12.2	22.4	0.6	2.9	15.2

with the aim to reduce consumer uncertainty about future prices. The resulting boost in consumer demand has the potential to significantly increase revenues. The analysis of our model, which is formulated as a discrete-time optimal control problem and is solved using a nonlinear programming approach, shows that an optimal solution exists under reasonable assumptions, and that this optimal solution has natural structural properties. The main technical contribution is the analysis of a model that is non-Markovian in terms of the underlying sales process. We show that small instances of the problem can be solved using a standard nonlinear programming solver, and provide a computationally tractable heuristic for larger instances.

The future research related to this model may include the following topics:

(1) Embedding this model into a game-theoretic framework that models the effects of price guarantee tools in a competitive setting;

(2) Analysis of the effects of risk exchange between customers and companies resulting from introduction of price guarantee tools;

(3) Extensions of our model to include such factors as multiple product types, returns, exchanges, and cancellations;

(4) Price guarantee tools in the context of general inventory management systems including variable and stochastic costs, and stochastic interest rates; and

(5) Solutions within smaller classes of easily implementable policies such as price-matching structures identified in the numerical experiments.

Figure 7. $(Y-y, t)$ pairs for which the absolute value of the expected price guarantee payment (46) is greater than or equal to $10^{-4}L(y,t)$ for $\alpha=0.2$.

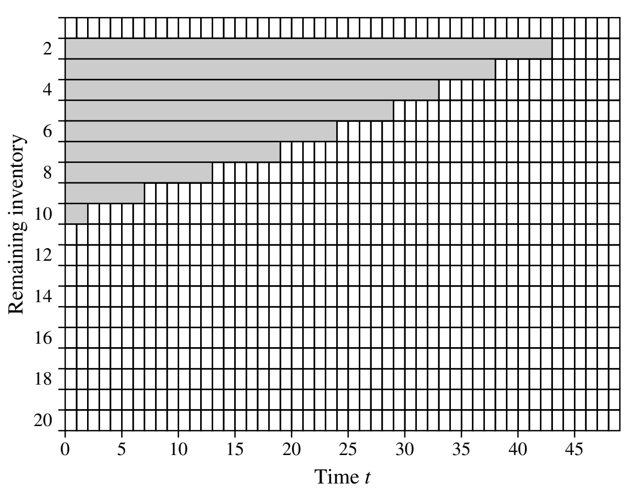

5. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Acknowledgments

The research of the first and the second authors was supported by the Natural Sciences and Engineering Research Council of Canada (grants 261512-04 and 138093-2001). The authors thank two referees for their many constructive suggestions that helped to better explain the model, simplify the analysis, and improve the results. The authors are also grateful to Peter Dacin and the attendees of their presentation at the INFORMS annual meeting in Denver, Colorado in 2004 for valuable comments and helpful suggestions.

References

Belobaba, Peter P. 1987. Airline yield management: An overview of seat inventory control. *Transportation Sci.* **21** 63–73.

Bertsekas, Dimitri. 1995. *Nonlinear Programming*. Athena Scientific, Belmont, MA.

Bertsekas, Dimitri. 2000. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA.

Chatwin, Richard E. 2000. Optimal dynamic pricing of perishable products with stochastic demand and a finite set of prices. *Eur. J. Oper. Res.* **125** 149–174.

Czyzyk, Joseph, Michael Mesnier, Jorge Moré. 1998. The NEOS server. *IEEE J. Computational Sci. Engrg.* **5** 68–75.

Dolan, Elizabeth. 2001. The NEOS server 4.0 administrative guide. Technical Report ANL/MCS-TM-250, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL.

Feng, Youyi, Baichun Xiao. 2000a. A continuous time yield management model with multiple prices and reversible price changes. *Management Sci.* **46**(5) 644–657.

Feng, Youyi, Baichun Xiao. 2000b. Optimal policies of yield management with multiple predetermined prices. *Oper. Res.* **48**(2) 332–343.

Gallego, Guillermo. 1992. A min-max distribution free procedure for the (q, r) inventory model. *Oper. Res. Lett.* **11** 55–60.

Gallego, Guillermo, Garrett van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.

Gīhman, Iosif I., Anatoli V. Skorohod. 1979. *Controlled Stochastic Processes*. Springer-Verlag, New York.

Gropp, William, Jorge Moré. 1997. Optimization environments and the NEOS server. M. D. Buhmann, A. Iserles, eds. *Approximation Theory and Optimization.* Cambridge University Press, Cambridge, UK, 167–182.

Hess, James, Eitan Gerstner. 1991. Price matching policies: An empirical case. *Managerial Decision Econom.* **12** 305–315.

McGill, Jeffrey I., Garrett J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.

Moorthy, Sridhar, Ralph Winter. 2005. Price-matching guarantees. *RAND J. Econom.* Forthcoming.

Phillips, Robert. 2005. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA.

Srivastava, Joydeep, Nicholas Lurie. 2001. A consumer perspective on price-matching refund policies: Effect on price perceptions and search behaviour. *J. Consumer Res.* **28**(2) 296–307.

Talluri, Kalyan T., Garrett J. van Ryzin. 2004a. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* **50**(1) 15–33.

Talluri, Kalyan T., Garrett J. van Ryzin. 2004b. *The Theory and Practice of Revenue Management*. Kluwer, Boston, MA.

Weatherford, Lawrence R., Samuel E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Oper. Res.* **40** 831–844.

Zhao, Wen, Yu-Sheng Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Sci.* **46**(3) 375–388.

OPERATIONS RESEARCH
DOI 10.1287/opre.1060.0344ec
pp. ec1–ec7

e-companion

ONLY AVAILABLE IN ELECTRONIC FORM

E-Companion—"Price Guarantees in Dynamic Pricing and Revenue Management" by Yuri Levin, Jeff McGill, and Mikhail Nediak, *Operations Research* 2007, 55(1) 75–97.

Appendix. Proofs of Mathematical Statements

In the proofs, we will frequently use the notion of the set of triples *extending* a given triple $(\mathcal{N}_1, \mathcal{N}_2, t)$ in $\mathfrak{N}$:

$$\mathfrak{N}(\mathcal{N}_1, \mathcal{N}_2, t) = \big\{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}\colon\ \mathcal{N}_1 = \mathcal{N}_1' \cap [0, t),\ \mathcal{N}_2 = \mathcal{N}_2' \cap [0, t),\ t \leqslant t'\big\}. \tag{EC1}$$

A triple included in this extension is any internal triple which coincides with $(\mathcal{N}_1, \mathcal{N}_2, t)$ for each time strictly before t. Note that $\mathfrak{N}(\mathcal{N}_1, \mathcal{N}_2, t)$ includes the triple $(\mathcal{N}_1, \mathcal{N}_2, t)$. The sets $\bar{\mathfrak{N}}(\mathcal{N}_1, \mathcal{N}_2, t)$ and $\mathfrak{N}^{\triangle}(\mathcal{N}_1, \mathcal{N}_2, t)$ are similarly defined:

$$\bar{\mathfrak{N}}(\mathcal{N}_1, \mathcal{N}_2, t) = \{(\mathcal{N}_1', \mathcal{N}_2', t') \in \bar{\mathfrak{N}}\colon\ \mathcal{N}_1 = \mathcal{N}_1' \cap [0, t),\ \mathcal{N}_2 = \mathcal{N}_2' \cap [0, t),\ t \leqslant t'\}, \tag{EC2}$$

$$\mathfrak{N}^{\triangle}(\mathcal{N}_1, \mathcal{N}_2, t) = \{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}^{\triangle}\colon\ \mathcal{N}_1 = \mathcal{N}_1' \cap [0, t),\ \mathcal{N}_2 = \mathcal{N}_2' \cap [0, t),\ t \leqslant t'\}. \tag{EC3}$$

EC.1. Proof of Lemma 1

Note that using probabilities of triples $P(\mathcal{N}_1, \mathcal{N}_2, t)$, one can express total expected revenues, for any fixed $t' \in [0, T]$, as

$$J(\varnothing, \varnothing, 0) = \sum_{\{(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}\colon t = t'\} \cup \{(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}^{\triangle}\colon t \leqslant t'\}} \Bigg\{ \sum_{\tau \in \mathcal{N}_1 \cup \mathcal{N}_2} p(\mathcal{N}_1 \cap [0, \tau),\ \mathcal{N}_2 \cap [0, \tau), \tau) + \sum_{\tau \in \mathcal{N}_2} f(\mathcal{N}_1 \cap [0, \tau),\ \mathcal{N}_2 \cap [0, \tau), \tau) + J(\mathcal{N}_1, \mathcal{N}_2, t) \Bigg\} P(\mathcal{N}_1, \mathcal{N}_2, t). \tag{EC4}$$

Therefore, we see that for any triple $(\mathcal{N}_1, \mathcal{N}_2, t) \in \bar{\mathfrak{N}}$, $J(\mathcal{N}_1, \mathcal{N}_2, t)$ participates in $J(\varnothing, \varnothing, 0)$ with a nonnegative coefficient $P(\mathcal{N}_1, \mathcal{N}_2, t)$. Observe that maximizing $J(\varnothing, \varnothing, 0)$ implies maximizing each of the $J(\mathcal{N}_1, \mathcal{N}_2, t)$ over the variables defined on triples extending $(\mathcal{N}_1, \mathcal{N}_2, t)$ while the policy on triples leading to $(\mathcal{N}_1, \mathcal{N}_2, t)$ is fixed.

We consider inequality (22) first. Suppose that it does not hold for some feasible solution corresponding to a policy (p, κ, ϕ) and a triple $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$, where t is the largest for which it does not hold. Construct a new policy (p', κ', ϕ') so that it differs from (p, κ, ϕ) only for triples $(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}(\mathcal{N}_1, \mathcal{N}_2, t+1)$ (note that $t \notin \mathcal{N}_1', \mathcal{N}_2'$) as follows:

$$(p', \kappa', \phi')(\mathcal{N}_1', \mathcal{N}_2', t') = \begin{cases} (p, \kappa, \phi)(\mathcal{N}_1' \cup t, \mathcal{N}_2', t') & \text{if } |\mathcal{N}_1'| + |\mathcal{N}_2'| \leqslant Y - 2, \\ (\tilde{p}, 0, 1) & \text{if } |\mathcal{N}_1'| + |\mathcal{N}_2'| = Y - 1, \end{cases}$$

where $\tilde{p}$ is equal to the maximum of $p(\mathcal{N}_1', \mathcal{N}_2', t')$ over all $(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$. Then:

1. The probabilities of purchasing any of the next $Y - 1 - (|\mathcal{N}_1| + |\mathcal{N}_2|)$ items and corresponding price guarantees in a given state $(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}(\mathcal{N}_1, \mathcal{N}_2, t+1)$ are precisely the same as in the state $(\mathcal{N}_1' \cup t, \mathcal{N}_2', t')$ for $|\mathcal{N}_1'| + |\mathcal{N}_2'| \leqslant Y - 2$.
2. The resulting price guarantee payments will be precisely the same both under the scenarios extending $(\mathcal{N}_1, \mathcal{N}_2, t+1)$ and the ones extending $(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$ because once the next $Y - 1 - (|\mathcal{N}_1| + |\mathcal{N}_2|)$ items are sold, all subsequent prices are fixed at the level $\tilde{p}$, which cannot result in higher guarantee payments.
3. It follows that the new expected revenues $J'(\mathcal{N}_1', \mathcal{N}_2', t') \geqslant J(\mathcal{N}_1' \cup t, \mathcal{N}_2', t')$ for all $(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}(\mathcal{N}_1, \mathcal{N}_2, t+1)$ and, therefore, $J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) \leqslant J'(\mathcal{N}_1, \mathcal{N}_2, t+1)$.

4. We also have $J'(\varnothing,\varnothing,0)\geqslant J(\varnothing,\varnothing,0)$ because $J'(\mathcal{N}_1,\mathcal{N}_2,t+1)\geqslant J(\mathcal{N}_1,\mathcal{N}_2,t+1)$, and it participates in $J'(\varnothing,\varnothing,0)$ with a nonnegative coefficient.

The construction is repeated until (22) is satisfied for all triples in $\mathfrak{N}$.

The Proof of (23) is similar, with the new policy being defined by

$$(p',\kappa',\phi')(\mathcal{N}_1'\cup t,\mathcal{N}_2'\backslash t,t')=(p,\kappa,\phi)(\mathcal{N}_1',\mathcal{N}_2',t')$$

for all $(\mathcal{N}_1',\mathcal{N}_2',t')\in\mathfrak{N}(\mathcal{N}_1,\mathcal{N}_2\cup t,t+1)$.

To prove (24), we rewrite (4) as

$$\begin{aligned}&J(\mathcal{N}_1,\mathcal{N}_2,t)-J(\mathcal{N}_1,\mathcal{N}_2,t+1)\\&\quad=\bar{u}(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t)\{p(\mathcal{N}_1,\mathcal{N}_2,t)+(1-v(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t))J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1)\\&\qquad+v(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t)\big[f(\mathcal{N}_1,\mathcal{N}_2,t)+J(\mathcal{N}_1,\mathcal{N}_2\cup t,t+1)\big]-J(\mathcal{N}_1,\mathcal{N}_2,t+1)\},\end{aligned}$$

and select $\kappa(\mathcal{N}_1,\mathcal{N}_2,t)=0$, $\phi(\mathcal{N}_1,\mathcal{N}_2,t)=1$. Then, $v(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t)=0$, and

$$J(\mathcal{N}_1,\mathcal{N}_2,t)-J(\mathcal{N}_1,\mathcal{N}_2,t+1)=\bar{u}(\Pi(\mathcal{N}_1,\mathcal{N}_2,t),t)\{p(\mathcal{N}_1,\mathcal{N}_2,t)+J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1)-J(\mathcal{N}_1,\mathcal{N}_2,t+1)\}.$$

We can always select $p(\mathcal{N}_1,\mathcal{N}_2,t)$ sufficiently large for the right-hand side to be nonnegative, and, therefore, for (24) to hold. Because this improves $J(\mathcal{N}_1,\mathcal{N}_2,t)$, if (24) did not hold initially, $J(\varnothing,\varnothing,0)$ cannot decrease due to this modification of the policy.

Finally, note that in all policy modifications used in this proof, we can ensure that the minimum price does not decrease.

EC.2. Proof of Lemma 2

We first show that p^* is positive. Due to continuous differentiability of $u(p,\kappa,\phi,t)$, there exists $\epsilon>0$ such that $u(p,\kappa,\phi,t)$ is Lipschitz with some constant $K>0$ and its partial derivative in p is bounded below by $-M<0$ for all $p\in[0,\epsilon]$ and all κ, ϕ, t. Consider any local maximum of $pu(p,\kappa,\phi,t)$ with respect to p for any fixed κ, ϕ, t. This point is either greater than ϵ or, by local optimality,

$$0=p\frac{\partial u}{\partial p}+u>-pM+(u_0-Kp),$$

where u_0 is a constant bounding $u(p,\kappa,\phi,t)$ away from zero by Assumption 10 of §2.1.2. It follows that $p\geqslant\min\{u_0/(K+M),\epsilon\}$. Because we considered an arbitrary local maximum of $pu(p,\kappa,\phi,t)$ for arbitrary values of κ, ϕ, t, it also follows that $p^*\geqslant\min\{u_0/(K+M),\epsilon\}>0$.

Suppose that $p(\mathcal{N}_1,\mathcal{N}_2,t)<p^*$ and t is the largest for which such a triple $(\mathcal{N}_1,\mathcal{N}_2,t)$ exists. The right-hand side of (4), after the substitution (20)–(21) is carried out, can be written as the function of the policy at $(\mathcal{N}_1,\mathcal{N}_2,t)$:

$$\begin{aligned}F(p)=\;&\bar{u}(p,\kappa,\phi,t)p[1+v(\kappa,\phi,t)\kappa\phi]+\bar{u}(p,\kappa,\phi,t)\{J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1)-J(\mathcal{N}_1,\mathcal{N}_2,t+1)+v(\kappa,\phi,t)\\&\cdot[J(\mathcal{N}_1,\mathcal{N}_2\cup t,t+1)-J(\mathcal{N}_1\cup t,\mathcal{N}_2,t+1)]\}+J(\mathcal{N}_1,\mathcal{N}_2,t+1),\end{aligned}$$

where only $\bar{u}$ depends on p but not v. If we increase p to p^*, the value of $F(p)$ can only increase because the term inside $\{\}$ is nonpositive by (22)–(23), $u(p,\kappa,\phi,t)$ is decreasing in p, and $pu(p,\kappa,\phi,t)$ is nondecreasing for all $p<p^*$. The price guarantee payments will not be affected because the future price never drops below p^* by our choice of the triple $(\mathcal{N}_1,\mathcal{N}_2,t)$. It follows that one can increase the value of $J(\mathcal{N}_1,\mathcal{N}_2,t)$ by increasing $p(\mathcal{N}_1,\mathcal{N}_2,t)$ to p^*. The value of $J(\varnothing,\varnothing,0)$ will not decrease with the increase in $J(\mathcal{N}_1,\mathcal{N}_2,t)$. Repeat this construction until there are no triples for which $p(\mathcal{N}_1,\mathcal{N}_2,t)<p^*$.

EC.3. Proof of Proposition 1

In the proofs of Lemmas 1 and 2, we relied on the fact that $P(\mathcal{N}_1,\mathcal{N}_2,t)\geqslant 0$ to argue that our "improved" policies did not decrease $J(\varnothing,\varnothing,0)$ because they could only increase $J(\mathcal{N}_1,\mathcal{N}_2,t)$ (see (EC4)). In fact, these improved policies produced a strict increase in $J(\mathcal{N}_1,\mathcal{N}_2,t)$'s whenever they were applicable. This completes the proof.

EC.4. Proof of Lemma 3

For $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}^{\triangle}$, $J(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant 0$. Suppose that, for some $(\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}$, $J(\mathcal{N}_1, \mathcal{N}_2, t)$ is not bounded. Consider the largest t such that this is the case. Because all of $J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$, $J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1)$, and $J(\mathcal{N}_1, \mathcal{N}_2, t+1)$ are bounded, it must be that the expression

$$\bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)p(\mathcal{N}_1, \mathcal{N}_2, t)(1 + v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)\kappa(\mathcal{N}_1, \mathcal{N}_2, t)\phi(\mathcal{N}_1, \mathcal{N}_2, t))$$

is unbounded. Because

$$v(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)\kappa(\mathcal{N}_1, \mathcal{N}_2, t)\phi(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant 1,$$

we conclude that $\bar{u}(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)p(\mathcal{N}_1, \mathcal{N}_2, t)$ is unbounded. This leads to a contradiction because the value of $h(p, \kappa, \phi, t) = \bar{u}(p, \kappa, \phi, t)p$ is bounded.

Indeed, in case (i), consider arbitrary κ, ϕ, t. Observe that

$$\frac{\partial h}{\partial p} = \bar{u}\left(1 + \frac{p}{\bar{u}}\frac{\partial \bar{u}}{\partial p}\right) \leqslant \bar{u}\left(1 - \frac{1}{1-\epsilon}\right) = \frac{-\bar{u}\epsilon}{1-\epsilon} < 0$$

for $p \geqslant p_{\max}$. Therefore, $h(p, \kappa, \phi, t)$ is decreasing in p for $p \geqslant p_{\max}$. In case (ii), $h(p, \kappa, \phi, t) = 0$ for $p \geqslant p_{\max}$, and, consequently, is nonincreasing. Thus, the value of $h(p, \kappa, \phi, t)$ cannot be increased by considering $p > p_{\max}$. Under either condition, the supremum of $h(p, \kappa, \phi, t)$ over the set of p, κ, ϕ such that $0 \leqslant p \leqslant p_{\max}$, $0 \leqslant \kappa$, $\phi \leqslant 1$, is attained because $h(p, \kappa, \phi, t)$ is continuous, and the set is bounded.

EC.5. Proof of Lemma 4

We first prove part (a).

It follows from the KKT conditions for the modified problem that

$$\sigma(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant 0 \qquad \text{and} \qquad \sigma(\mathcal{N}_1, \mathcal{N}_2, t)[p(\mathcal{N}_1, \mathcal{N}_2, t) - p_{\max}(t)] = 0.$$

Therefore, it is possible for $\sigma(\mathcal{N}_1, \mathcal{N}_2, t)$ to be strictly positive only if $p(\mathcal{N}_1, \mathcal{N}_2, t) = p_{\max}(t)$. Because we must also have

$$\frac{\partial \mathcal{L}'}{\partial p(\mathcal{N}_1, \mathcal{N}_2, t)} = \frac{\partial \mathcal{L}}{\partial p(\mathcal{N}_1, \mathcal{N}_2, t)} - \sigma(\mathcal{N}_1, \mathcal{N}_2, t) = 0,$$

the inequality $\sigma(\mathcal{N}_1, \mathcal{N}_2, t) > 0$ would imply that

$$\frac{\partial \mathcal{L}}{\partial p(\mathcal{N}_1, \mathcal{N}_2, t)} > 0.$$

Consider the expression for this partial derivative where, for brevity, we omit the argument $(\Pi(\mathcal{N}_1, \mathcal{N}_2, t), t)$ as long as such omission is unambiguous:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial p} &= \pi(\mathcal{N}_1, \mathcal{N}_2, t)\bigg\{\frac{\partial \bar{u}}{\partial p}[p + J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) - J(\mathcal{N}_1, \mathcal{N}_2, t+1) \\ &\qquad + v(\phi\kappa p + J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) - J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1))] + \bar{u}[1 + v\phi\kappa]\bigg\} \\ &\quad + \sum_{\substack{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}^{\triangle}(\mathcal{N}_1, \mathcal{N}_2, t), \\ t'' \in \mathcal{N}_2,\, t'' > t - D}} \zeta(\mathcal{N}_1', \mathcal{N}_2', t, t'') - \sum_{\substack{(\mathcal{N}_1', \mathcal{N}_2', t') \in \mathfrak{N}^{\triangle}(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1), \\ t < \tau < \min\{t', t+D\}}} \zeta(\mathcal{N}_1', \mathcal{N}_2', \tau, t)\kappa(\mathcal{N}_1, \mathcal{N}_2, t).\end{aligned} \tag{EC5}$$

This expression can be positive only if both $\pi(\mathcal{N}_1, \mathcal{N}_2, t)$ and the expression in $\{\}$ is positive, or one of the $\zeta(\mathcal{N}_1', \mathcal{N}_2', t, t'')$'s is positive.

First, if $\pi(\mathcal{N}_1, \mathcal{N}_2, t) = 0$, then we can easily show that all of the $\zeta(\mathcal{N}_1', \mathcal{N}_2', t, t'')$'s are zeros, and $\partial \mathcal{L}/\partial p$ cannot be positive. Thus, we will consider $\pi(\mathcal{N}_1, \mathcal{N}_2, t) > 0$, which means that triple $(\mathcal{N}_1, \mathcal{N}_2, t)$ has positive probability.

Due to the complementary slackness condition of the KKT system, $\zeta(\mathcal{N}_1', \mathcal{N}_2', t, t'') > 0$ is possible only if

$$z_{t''}(\mathcal{N}_1', \mathcal{N}_2') = p(\mathcal{N}_1, \mathcal{N}_2, t) - \kappa(\mathcal{N}_1 \cap [0, t''), \mathcal{N}_2 \cap [0, t''), t'')p(\mathcal{N}_1 \cap [0, t''), \mathcal{N}_2 \cap [0, t''), t''), \tag{EC6}$$

and, because $z_{t''}(\mathcal{N}_1', \mathcal{N}_2') \leqslant 0$, this implies that

$$p_{\max}(t) = p(\mathcal{N}_1, \mathcal{N}_2, t) \leqslant \kappa(\mathcal{N}_1 \cap [0, t''), \mathcal{N}_2 \cap [0, t''), t'')p(\mathcal{N}_1 \cap [0, t''), \mathcal{N}_2 \cap [0, t''), t'') \\ \leqslant p(\mathcal{N}_1 \cap [0, t''), \mathcal{N}_2 \cap [0, t''), t'') \leqslant p_{\max}(t''). \tag{EC7}$$

However, this is impossible because $p_{\max}(t)$ is strictly increasing and $t'' < t$.

Observe now that for any optimal solution to the modified problem, the following inequalities are true:

$$J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) \geqslant -\sum_{\theta \in \mathcal{N}_2} p_{\max}(\theta), \tag{EC8}$$

$$J(\mathcal{N}_1 \cup t,\ \mathcal{N}_2, t+1) \geqslant -\sum_{\theta \in \mathcal{N}_2} p_{\max}(\theta), \tag{EC9}$$

$$-J(\mathcal{N}_1,\ \mathcal{N}_2,\ t+1) \geqslant -U^*, \tag{EC10}$$

where U^* is an upper bound on the $J(\mathcal{N}_1, \mathcal{N}_2, t)$'s (see Lemma 3). The last two inequalities are immediate. The first one is true because we can always ensure that we do not need to make any price guarantee payments on the last sale. If this inequality did not hold, we could improve the solution simply by keeping subsequent prices at the same level as for the sale at time t (and this would contradict the solution's optimality). Then, if $p(\mathcal{N}_1, \mathcal{N}_2, t) = p_{\max}(t)$, we have

$$\frac{vJ(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) + (1-v)J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) - J(\mathcal{N}_1, \mathcal{N}_2, t+1)}{p(1 + v\kappa\phi)} \geqslant -\frac{\epsilon}{2}, \tag{EC11}$$

and it is straightforward to show that the term in {} in (EC5) is strictly negative. This implies that

$$\frac{\partial \mathcal{L}}{\partial p(\mathcal{N}_1, \mathcal{N}_2, t)} \leqslant 0,$$

and, therefore, $\sigma(\mathcal{N}_1, \mathcal{N}_2, t)$ cannot be strictly positive.

To prove part (b), consider any triple $(\mathcal{N}_1, \mathcal{N}_2, t)$ such that $p(\mathcal{N}_1, \mathcal{N}_2, t) > \bar{p}_{\max}(t)$ and $(\mathcal{N}_1, \mathcal{N}_2, t)$ is not an extension of any other triple with that property, that is

$$p(\mathcal{N}_1 \cap [0, \tau), \mathcal{N}_2 \cap [0, \tau), \tau) \leqslant \bar{p}_{\max}(\tau), \quad 0 \leqslant \tau < t.$$

Because $\bar{p}_{\max}(t)$ is strictly increasing and $p(\mathcal{N}_1, \mathcal{N}_2, t) > \bar{p}_{\max}(t)$, we conclude that the $\zeta(\mathcal{N}_1', \mathcal{N}_2', t, t'')$'s in (EC5) are all equal to zero using reasoning similar to (EC6)–(EC7). We also observe that (EC8) and (EC9) hold with $\bar{p}_{\max}(t)$ used for $p_{\max}(t)$ or we would get a contradiction with the optimality of a given solution. Then again, (EC11) holds and we conclude that {} in (EC5) is strictly negative. If $\pi(\mathcal{N}_1, \mathcal{N}_2, t) > 0$, it must be that

$$\frac{\partial \mathcal{L}}{\partial p(\mathcal{N}_1, \mathcal{N}_2, t)} < 0,$$

and the KKT conditions cannot be satisfied—a contradiction. Therefore, $\pi(\mathcal{N}_1, \mathcal{N}_2, t) = 0$ and, also, $\pi(\mathcal{N}_1', \mathcal{N}_2', t') = 0$, for all triples extending $(\mathcal{N}_1, \mathcal{N}_2, t)$. Thus, part (b) holds.

EC.6. Proof of Proposition 3

Case (ii) is trivial. Consider case (i).

We examine the modified problem with $\bar{p}_{\max}(t)$ used in (29) and an additional technical constraint

$$\kappa(\mathcal{N}_1, \mathcal{N}_2, t) + \phi(\mathcal{N}_1, \mathcal{N}_2, t) \geqslant \epsilon' \quad \text{for all } (\mathcal{N}_1, \mathcal{N}_2, t) \in \mathfrak{N}, \tag{EC12}$$

where $0 < \epsilon' \leqslant p^*/\max_t \bar{p}_{\max}(t)$. We recall that due to our general requirements on the form of v, it is discontinuous at the point $\kappa = 0$, $\phi = 0$. Thus, the constraint (EC12) is needed to ensure the continuity of all functions involved in the discrete-time model. First, we observe that this constraint does not change the proof of Lemma 4. Second, we show that by adding the constraint (EC12) we do not sacrifice optimality in the modified problem. Recall Lemma 2, which claims that we do not sacrifice optimality by considering only solutions satisfying (25). If there exists a triple $(\mathcal{N}_1, \mathcal{N}_2, t)$ for which $\kappa(\mathcal{N}_1, \mathcal{N}_2, t) + \phi(\mathcal{N}_1, \mathcal{N}_2, t) < \epsilon'$ in such a feasible solution, then the price guarantee sold at $(\mathcal{N}_1, \mathcal{N}_2, t)$ cannot result in price guarantee payments. Thus, we can ensure that

$J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) = J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$. We now show how to modify the policy at $(\mathcal{N}_1, \mathcal{N}_2, t)$ to make sure that (EC12) is satisfied but the corresponding expected value is not smaller. The right-hand side of (4), as long as $J(\mathcal{N}_1, \mathcal{N}_2 \cup t, t+1) = J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1)$, can be written as the following function of the policy variables:

$$F(p, \kappa, \phi) = \bar{u}(p, \kappa, \phi, t)\{p[1 + v(\kappa, \phi, t)\kappa\phi] + J(\mathcal{N}_1 \cup t, \mathcal{N}_2, t+1) - J(\mathcal{N}_1, \mathcal{N}_2, t+1)\} + J(\mathcal{N}_1, \mathcal{N}_2, t+1).$$

If the term in $\{\}$ is nonnegative, we can increase the value of κ to ϵ' because both $\bar{u}(p, \kappa, \phi, t)$ and $v(\kappa, \phi, t)\kappa$ are nondecreasing in κ. If the term in $\{\}$ is negative, we can increase ϕ until either $\phi = \epsilon'$ or the term in $\{\}$ becomes nonnegative. In the latter case, we can increase κ as before. In both cases, the value of $J(\mathcal{N}_1, \mathcal{N}_2, t)$ and therefore $J(\varnothing, \varnothing, 0)$ will not decrease.

An optimal solution to the modified problem exists because its feasible set is compact (as a closed and bounded set in finite dimensions), and all functions involved are continuous. Let the value of this solution be w^*.

We claim that this is also an optimal solution to the original problem. Indeed, suppose that this is not the case. Then, because the original problem is less restricted, there exists a solution of value $w > w^*$. This solution must violate (29) and, possibly, (EC12) or it would not have a higher value. Select $p_{\max}(t)$ satisfying (30)–(31) so that it majorizes the price component of the policy in this solution. (The value of ϵ' may have to be decreased as well. However, the above reasoning holds for smaller values of ϵ'.) We solve the modified problem using this new bound $p_{\max}(t)$ for the price and get an optimal solution of value $w' \geqslant w > w^*$. By part (b) of Lemma 4, this solution cannot violate constraint (29) with the initial bound $\bar{p}_{\max}(t)$ except perhaps on triples of probability zero. By Remark 3, we can modify this solution and obtain another feasible solution which *does* satisfy (29) with $\bar{p}_{\max}(t)$ everywhere while having a larger objective value w'. This contradicts the optimality of w^*.

EC.7 Proof of Proposition 4

Write down and set equal to zero the derivatives of the Lagrangian with respect to $p(\mathcal{N}, t)$ and $k(\mathcal{N}, t)$:

$$\frac{\partial \mathcal{L}}{\partial p(\mathcal{N}, t)} = \pi(\mathcal{N}, t)\left\{\frac{\partial \bar{u}}{\partial p}(\Pi^0(\mathcal{N}, t), t)[p(\mathcal{N}, t) + J(\mathcal{N} \cup t, t+1) - J(\mathcal{N}, t+1)] + \bar{u}(\Pi^0(\mathcal{N}, t), t)\right\} + \sum_{\substack{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N}, t),\\ t'' \in \mathcal{N}, t'' > t-D}} \zeta(\mathcal{N}', t, t'') + \eta(\mathcal{N}, t) = 0, \tag{EC13}$$

$$\frac{\partial \mathcal{L}}{\partial k(\mathcal{N}, t)} = \pi(\mathcal{N}, t)\frac{\partial \bar{u}}{\partial k}(\Pi^0(\mathcal{N}, t), t)[p(\mathcal{N}, t) + J(\mathcal{N} \cup t, t+1) - J(\mathcal{N}, t+1)] - \sum_{\substack{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N} \cup t, t+1),\\ t < \tau < \min\{t', t+D\}}} \zeta(\mathcal{N}', \tau, t) - \eta(\mathcal{N}, t) + \xi(\mathcal{N}, t) = 0. \tag{EC14}$$

For a boundary triple $(\mathcal{N}', t')$ and $t \in \mathcal{N}'$, consider the derivative of $\mathcal{L}$ with respect to $z_t(\mathcal{N}')$ which is equal to zero due to the KKT conditions:

$$\frac{\partial \mathcal{L}}{\partial z_t(\mathcal{N}')} = \pi(\mathcal{N}', t') - \sum_{t \leqslant \tau < \min\{t', t+D\}} \zeta(\mathcal{N}', \tau, t) = 0.$$

It follows that

$$\sum_{t < \tau < \min\{t', t+D\}} \zeta(\mathcal{N}', \tau, t) = \pi(\mathcal{N}', t') - \zeta(\mathcal{N}', t, t),$$

and we can write

$$\begin{aligned}\sum_{\substack{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N} \cup t, t+1),\\ t < \tau < \min\{t', t+D\}}} \zeta(\mathcal{N}', \tau, t) &= \sum_{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N} \cup t, t+1)} [\pi(\mathcal{N}', t') - \zeta(\mathcal{N}', t, t)] \\ &= \pi(\mathcal{N} \cup t, t+1) - \sum_{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N} \cup t, t+1)} \zeta(\mathcal{N}', t, t) \\ &= \pi(\mathcal{N}, t)\bar{u}(\Pi^0(\mathcal{N}, t), t) - \sum_{(\mathcal{N}', t') \in \Re^{\triangle}(\mathcal{N} \cup t, t+1)} \zeta(\mathcal{N}', t, t).\end{aligned}$$

Then, by summing (EC13) with (EC14), we get

$$\pi(\mathcal{N},t)\left(\frac{\partial\bar{u}}{\partial k}(\Pi^0(\mathcal{N},t),t)+\frac{\partial\bar{u}}{\partial p}(\Pi^0(\mathcal{N},t),t)\right)(p(\mathcal{N},t)+J(\mathcal{N}\cup t,t+1)-J(\mathcal{N},t+1))$$
$$+\left[\sum_{\substack{(\mathcal{N}',t')\in\mathfrak{R}^\triangle(\mathcal{N},t),\\ t''\in\mathcal{N},\,t''>t-D}}\zeta(\mathcal{N}',t,t'')+\sum_{(\mathcal{N}',t')\in\mathfrak{R}^\triangle(\mathcal{N}\cup t,t+1)}\zeta(\mathcal{N}',t,t)+\xi(\mathcal{N},t)\right]=0.$$

Because all terms enclosed in [ ] are nonnegative, $\pi(\mathcal{N},t)=P(\mathcal{N},t)>0$ and

$$p(\mathcal{N},t)+J(\mathcal{N}\cup t,t+1)-J(\mathcal{N},t+1)=\frac{J(\mathcal{N},t)-J(\mathcal{N},t+1)}{\bar{u}(\Pi^0(\mathcal{N},t),t)}>0,$$

we conclude that inequality (34) holds.

EC.8. Proof of Lemma 5

For any fixed price p, the expected revenue is given by $pE[N_p-(N_p-Y)^+]$, where N_p is a binomial random variable with T trials and probability of success $\bar{u}(p,1,0)$. Next, we use the result of Gallego (1992) that for any random variable N, with a finite mean μ, standard deviation σ, and any real number n,

$$E[(N-n)^+]\leqslant\frac{\sqrt{\sigma^2+(n-\mu)^2}-(n-\mu)}{2}.$$

If $p_0^0\geqslant p_0^*$, the random variable $N_{p_0^0}$ has $\mu=T\bar{u}(p_0^0,1,0)=Y$, and $\sigma^2=Y(1-Y/T)$. Thus, with $n=Y$,

$$L^{\text{fixed}}(Y,T)\geqslant p_0^0\left(Y-\frac{\sqrt{Y(1-Y/T)}}{2}\right)\geqslant Yp_0^0\left(1-\frac{1}{2\sqrt{Y}}\right).$$

If $p_0^0<p_0^*$, which is equivalent to $T\bar{u}(p_0^*,1,0)<Y$, then $\mu=T\bar{u}(p_0^*,1,0)$, $\sigma^2=T\bar{u}(p_0^*,1,0)(1-\bar{u}(p_0^*,1,0))$, and we have

$$L^{\text{fixed}}(Y,T)\geqslant p_0^*\left(T\bar{u}(p_0^*,1,0)-\frac{\sqrt{T\bar{u}(p_0^*,1,0)(1-\bar{u}(p_0^*,1,0))+(Y-T\bar{u}(p_0^*,1,0))^2}-(Y-T\bar{u}(p_0^*,1,0))}{2}\right)$$
$$\geqslant p_0^*\left(T\bar{u}(p_0^*,1,0)-\frac{\sqrt{T\bar{u}(p_0^*,1,0)(1-\bar{u}(p_0^*,1,0))}}{2}\right)$$
$$\geqslant p_0^*T\bar{u}(p_0^*,1,0)\left(1-\frac{1}{2\sqrt{T\bar{u}(p_0^*,1,0)}}\right).$$

This completes the proof.

EC.9. Proof of Lemma 6

Part (I): We will suppress the subscript $\phi=0,1$ because the proof is identical in both cases. Consider an arbitrary feasible solution $\tilde{u}_s$, $s=0,\ldots,T-1$, and its T cyclical permutations, including itself:

$$\begin{aligned}\tilde{\mathbf{u}}^1&=(\tilde{u}_1,\tilde{u}_2,\ldots,\tilde{u}_{T-2},\tilde{u}_{T-1}),\\ \tilde{\mathbf{u}}^2&=(\tilde{u}_2,\tilde{u}_3,\ldots,\tilde{u}_{T-1},\tilde{u}_1),\\ &\vdots\\ \tilde{\mathbf{u}}^T&=(\tilde{u}_{T-1},\tilde{u}_1,\ldots,\tilde{u}_{T-3},\tilde{u}_{T-2}).\end{aligned}$$

All of these vectors are feasible solutions with the same objective function value as in $\tilde{\mathbf{u}}^1$. The convex combination $\tilde{\mathbf{u}}'$ of $\tilde{\mathbf{u}}^s$, $s=1,\ldots,T$, with coefficients $1/T$ is a vector with all identical elements

$$\frac{1}{T}\sum_{s=0}^{T}\tilde{u}_s.$$

This convex combination is a feasible solution to (38)–(40) due to the convexity of the feasible region. Also, its objective $V^D(\tilde{\mathbf{u}}')$ is concave and at $\tilde{\mathbf{u}}'$ satisfies

$$V^D(\tilde{\mathbf{u}}') \geqslant \sum_{s=0}^{T} \frac{1}{T} V^D(\tilde{\mathbf{u}}^s) = V^D(\tilde{\mathbf{u}}^1).$$

Thus, to find the optimum, it is sufficient to consider solution vectors with constant components. The problem reduces to

$$\begin{aligned} \max\quad & Tp(\bar{u})\bar{u}, \\ \text{s.t.}\quad & T\bar{u} \leqslant Y, \\ & \bar{u} \leqslant \lambda. \end{aligned}$$

It is easy to verify the objective value claimed in the lemma.

Part (II): Define $V_1^D(Y, t) = V_1^D(y, T) = 0$ for all y, t. Then, the claim of the lemma can be established by inverse induction on t. Suppose that the claim is true for $t+1$. Then,

$$\begin{aligned} V(y,t) &= \bar{u}(p(y,t),1,1)[p(y,t)+V(y+1,t+1)] + (1-\bar{u}(p(y,t),1,1))V(y,t+1) \\ &\leqslant \bar{u}(p(y,t),1,1)[p(y,t)+V_1^D(y+1,t+1)] + (1-\bar{u}(p(y,t),1,1))V_1^D(y,t+1). \end{aligned} \tag{EC15}$$

The value of $w = \bar{u}(p(y,t),1,1)$ is such that $p(y,t) = p_1(w)$, according to the definition of $p_1(\cdot)$. Consider the optimal solutions $\bar{\mathbf{u}}^1$ and $\bar{\mathbf{u}}^2$ of $V_1^D(y+1, t+1)$ and $V_1^D(y, t+1)$, respectively. The vector of the form $\bar{\mathbf{u}}^0 = (w, w\bar{\mathbf{u}}^1 + (1-w)\bar{\mathbf{u}}^2)$ is a feasible solution of (38)–(40). Indeed, it is seen immediately that (40) is satisfied. The constraint (39) is satisfied as well since

$$w + w\sum_{s=t+1}^{T-1} \bar{u}_s^1 + (1-w)\sum_{s=t+1}^{T-1} \bar{u}_s^2 \leqslant w + w(Y-y-1) + (1-w)(Y-y) = Y-y.$$

Finally, the right-hand side of inequality (EC15) is less than or equal to the value of the objective function (38) at $\bar{\mathbf{u}}^0$ because, due to its concavity, the revenue rate at each $s = t+1, \ldots, T-1$ is such that

$$p_1(\bar{u}_s^0)\bar{u}_s^0 \geqslant wp_1(\bar{u}_s^1)\bar{u}_s^1 + (1-w)p_1(\bar{u}_s^2)\bar{u}_s^2.$$

The optimum value $V^D(y, t)$ is greater than or equal to the value of (38) at this feasible solution. This completes the proof.

EC.10. Proof of Proposition 5

In case $p^* \leqslant p_1^0$,

$$L^{\text{fixed}}(Y,T) \geqslant Yp_0^0\left(1-\frac{1}{2\sqrt{Y}}\right) = V_1^D(0,0)\frac{p_0^0}{p_1^0}\left(1-\frac{1}{2\sqrt{Y}}\right) \geqslant V(0,0)\frac{p_0^0}{p_1^0}\left(1-\frac{1}{2\sqrt{Y}}\right).$$

In case $p_1^0 < p^* \leqslant p_0^0$,

$$\begin{aligned} L^{\text{fixed}}(Y,T) &\geqslant Yp_0^0\left(1-\frac{1}{2\sqrt{Y}}\right) = V_1^D(0,0)\frac{Yp_0^0}{p^*T\bar{u}(p^*,1,1)}\left(1-\frac{1}{2\sqrt{Y}}\right) \\ &\geqslant V(0,0)\frac{Yp_0^0}{p^*Ta(p^*)b(1,1)}\left(1-\frac{1}{2\sqrt{Y}}\right). \end{aligned}$$

Finally, in case $p^* > p_0^0$,

$$\begin{aligned} L^{\text{fixed}}(Y,T) &\geqslant p^*T\bar{u}(p^*,1,0)\left(1-\frac{1}{2\sqrt{T\bar{u}(p^*,1,0)}}\right) = V_1^D(0,0)\frac{\bar{u}(p^*,1,0)}{\bar{u}(p^*,1,1)}\left(1-\frac{1}{2\sqrt{T\bar{u}(p^*,1,0)}}\right) \\ &\geqslant V(0,0)\frac{b(1,0)}{b(1,1)}\left(1-\frac{1}{2\sqrt{Ta(p^*)b(1,0)}}\right). \end{aligned}$$