Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach

Alex Gershkov and Benny Moldovanu*

Abstract

We study the revenue maximizing allocation of several heterogeneous, commonly ranked objects to impatient agents with privately known characteristics who arrive sequentially. There is a deadline after which no more objects can be allocated. We first characterize implementable allocation schemes, and compute the expected revenue for any implementable, deterministic and Markovian allocation policy. The revenue-maximizing policy is obtained by a variational argument which sheds more light on its properties than the usual dynamic programming approach. Finally, we use our main result in order to: a) derive the optimal inventory choice; b) explain empirical regularities about pricing in clearance sales.

We study the following dynamic mechanism design problem in continuous time: a designer has to allocate (or assign) a finite set of heterogeneous objects with known characteristics to a stream of randomly arriving, impatient agents with privately known characteristics. There is a deadline by which all objects must be sold.[1] The objects are substitutes, and each agent derives utility from at most one object. Moreover, all

*This manuscript supplants the previously circulated discussion paper "The Dynamic Assignment of Heterogeneous Objects: A Mechanism Design Approach". An anonymous referee made helpful remarks. We are grateful for financial support from the German Science Foundation, and from the Max Planck Society. We wish to thank Heidrun Hoppe, Paul Klemperer, and participants at the Conference on "Computational Social Systems", Dagstuhl 2007, and at seminars in Oxford, Warwick, Bonn, Berlin, Paris and Bruxelles for helpful comments. Gershkov, Moldovanu: Department of Economics, University of Bonn, Lennestr. 37, 53113 Bonn. alex.gershkov@uni-bonn.de, mold@uni-bonn.de.

[1]Simpler techniques can be applied to the setting where the horizon is infinite and where there is time discounting. We do present below the equivalent of our main result for this situation.

agents rank the available objects in the same way, and values for objects have a multiplicative structure involving the agents' and objects' types. Under the assumption that monetary transfers are feasible, we characterize the dynamic revenue-maximizing policy and the associated dynamic pricing scheme. We also show how important features of the revenue maximizing policy can be used to yield a straightforward solution to the larger optimization problem where the designer also chooses the size and quality composition of its inventory.

Dynamic pricing and assignment problems appear in numerous frameworks such as the retail of seasonal and style goods, the allocation of fixed capacities in the travel and leisure industries (e.g., airlines, trains, hotels, rental cars), the allocation of a fixed inventory of equipment in a given period of time (e.g., equipment for medical procedures), the assignment of personnel to incoming tasks.

The main trade-off in our paper is as follows: assigning an object today means that the valuable option of assigning it in the future - possibly to an agent who values it more-, is foregone; on the other hand, since the arrival process of agents is stochastic, and since there is a deadline, the "future" may never materialize. The revenue maximizing policy needs to define prices for each available objects such that, at each point in time, this trade-off is optimally taken into account.

The paper's main contributions are:

1. By allowing for heterogeneous objects, we can combine quantity/quality optimization with the pricing considerations that have been traditionally the focus of the literature (in Economics and Management Science) analyzing optimization under stochastic demand conditions[2].

2. We introduce to the above literature a new technical method which has two components: a) A focus - inspired by the "mechanism design philosophy" and the payoff/revenue equivalence principle - on implementable allocation policies rather than on prices; b) A variational method that yields somewhat more insight than the traditional dynamic programming/Bellman's equation (e.g., we obtain new results even for the much studied case of dynamic revenue-maximization for identical goods).

[2] In the Management Science literature this is called *revenue or yield management*.

3. Our analysis yields testable implications about the pattern of prices in relevant situations exhibiting fixed inventories of substitute goods that need to be sold by a deadline. An example concerning clearance sales for apparel is provided here.

Whereas the large literature on yield or revenue management has directly focused on revenue-maximizing pricing (mostly for the special case of our model where agents have linear, private values for identical objects) our approach starts by a characterization of all dynamically implementable, non-randomized allocation policies. Such policies are described by partitions of the set of possible agent types: an arriving agent gets the best available object if his type lies in the highest interval of the partition, the second best available object if his type lies in the second highest interval, and so on... These intervals may depend on the point in time of the arrival, and on the composition of the set of available objects at that point in time. For implementable allocation policies we derive the associated menus of prices (one menu for each point in time, and for each subset of remaining objects) that implement it, and show that these menus have an appealing recursive structure: each agent who is assigned an object has to pay the value he displaces in terms of the chosen allocation.

We next turn to revenue-maximization. Using several basic results about the Poisson stochastic process, we first compute the revenue generated by any individual-rational, non-randomized, Markovian and implementable allocation policy. Then, we can directly use variational arguments in order to characterize the revenue-maximizing allocation policy. The associated optimal prices are of "secondary importance" since they are completely determined by the implementation conditions.

Whereas the optimal prices necessarily depend on the composition of inventory, our main result is that, at each point in time, the revenue maximizing allocation policy depends only on the size of the available inventory, but not on its exact composition. To understand the meaning of this result, consider the same model, but with identical objects. Then, for each size of available inventory, and for each point in time, the revenue maximizing allocation policy is characterized by a single cut-off type: only an arriving agent with type above that cut-off obtains one of the objects. In contrast, when objects are heterogeneous, the revenue maximizing policy is, at each point in time, and for each subset of available objects, characterized by several cut-off types which determine if the arriving agent gets the best available object, the second best, etc... Our result says that for any subset of k available heterogeneous objects, and for any point in time, the highest cut-off coincides with the optimal cut-off in a situation with one available object, the second-highest cut-off coincides with the optimal cutoff in a situation with two identical objects, and so on till the lowest cut-off which coincides with the optimal cut-off in a situation with an inventory of k identical objects.

The last part of the paper is devoted to two applications that use the characterization result described above:

1. We embed revenue maximization in the larger optimization problem where, before sales begin, the seller chooses the size and composition of the inventory. A good recent illustration is offered by the move of several US supermarket chains to reduce "shrink" - the amount of fresh food that needs to be dumped since it is not sold by expiration date. Obviously, both the amount and variety of food on display, and the temporal pricing pattern will affect shrink. Although some "intentional" waste is part of revenue maximization, the US shrink rate is twice as big as that of European retailers, suggesting that at least one of them may not be optimal.[3]

   We show how the formulae we derived for revenue maximization (together with information about marginal costs) immediately yield a set of intuitive equations that characterize the optimal number of objects and their qualities.

2. We derive some testable implications about the pattern of observed prices for different qualities in clearance sales, and confront them with available data. Compared to standard models that only consider identical objects, our analysis offers a somewhat more convincing explanation for several well-known observed regularities. For example, we explain why the average clearance mark-down (in percentage terms) is higher for the higher quality, more expensive product lines, as empirically observed in a variety of settings (see literature review below)

The rest of the paper is organized is follows: In the remainder of this Section we review the related literature.

In Section 2 we present the continuous-time model of sequential assignment of heterogeneous objects to randomly arriving, privately informed agents. Section 3 focuses

[3]The estimate is that \$20 billion worth of shrink is wasted annually in the US, about 10% of sales. See "Shrink rapped", *The Economist,* May 17, 2008, p.75.

on a characterization of implementable policies, and of the associated menus of dynamic prices that implement such policies.

In Section 4 we study the revenue maximizing policy. We obtain a general expression for expected revenue given any Markovian, non-randomized allocation policy, and we use a variational argument in order to derive functional equations that characterize the revenue maximizing allocation policy and the expected revenue generated by this policy.

In Section 5 we exhibit two applications of our main result: In Subsection 5.1 we embed our revenue maximization in a larger maximization problem where the designer also chooses the size and composition of the inventory. In Subsection 5.2 we derive several empirically testable implications and show how these can be used to explain observed phenomena.

Section 6 concludes. Proofs are relegated to an Appendix.

Related Literature

Cyrus Derman, Gerald J. Lieberman and Sheldon M. Ross (1972) introduced an elegant discrete-time model where a set of distinct, commonly ranked objects needs to be assigned to a set of sequentially arriving agents that have a multiplicative utility function involving their type and the object's properties. They analyze the dynamic, welfare-maximizing allocation under a complete information assumption about the agents' valuations, and characterize it in terms of cutoffs. These cutoffs do not depend on the objects's properties[4]. Christian S. Albright (1974) extended the Derman-Lieberman-Ross model and characterization to the continuous-time framework with random arrivals of agents. We add incomplete information to their basic framework, and focus here on revenue maximization. Incomplete information in a similar, but static model with non-random demand is introduced in the classic paper by Michael Mussa and Sherwin Rosen (1978) who focused on a monopolist's price/quality decisions.

There are very large theoretical and applied literatures on dynamic pricing of inventories (sometimes called *revenue or yield management*) both in Economics and in the fields of Management and Operations Research. The general goal in this literature is revenue maximization. We refer the reader to the excellent literature surveys and discussions of

[4]Similar analytical methods - involving dynamic programming and optimal stopping - have been fruitfully used in the large literature on search- see for example the surveys of Lippman, Steven, and John J. McCall (1981) and Dale T. Mortensen (1986).

involved issues by Gabriel R. Bitran and René Caldentey (2003) and Wedad Elmaghraby and Pinar Keskinocak (2003), and to the comprehensive book by Kalyan T. Talluri and Garrett J. van Ryzin (2004). Preston McAfee and Vera te Velde (2007) survey the applications to the airline industry who has pioneered many of the modern practices in revenue management.

John Riley and Richard Zeckhauser (1983) considered a single object revenue-maximizing procedure where there is learning about the distribution of the agents' values. In their model, the optimal mechanism is a sequence of take-it-or-leave-it offers. Edward P. Lazear (1986) offers a theory of clearance sales that allows for several identical units. In his model all customers have the same value, and the seller learns about this value over time. It is shown that a pattern of decreasing prices is optimal. B. Peter Pashigian (1988), B. Peter Pashigian and Brian Bowen (1991), and B. Peter Pashigian, Brian Bowen and Eric Gould (1995) test some of the empirical implications of Lazear's analysis based on data obtained from several industries

In a continuous-time framework with stochastic arrivals of agents, Wilfred M. Kincaid and Donald A. Darling (1963), and Guillermo Gallego and Garrett van Ryzin (1994) use dynamic programming in order to characterize - implicitly via Bellman's equations - the revenue maximizing pricing policy for a set of identical objects that need to be sold before a deadline. A main result is that the expected revenue in the optimal policy - which is characterized for each size of inventory by a single posted price - is increasing and concave both in the number of objects and in the length of time left till the deadline. Moreover, each relevant cutoff price drops with time as long as there is no sale, but jumps up after each sale[5]. These authors were able to calculate in closed form the solution for what amounts (in our terms) to an exponential distribution of agents' values. McAfee and te Velde (2008) find an explicit solution for a Pareto distributions of agents' values, and show that it coincides with the welfare maximizing policy. Generally, a closed form solution is not available, and even the expression of expected revenue as a function of the optimal cutoff prices is not available in the literature.

Michael A. Arnold and Steven A. Lippman (2001), Gopal Das Varma and Nikolaos Vettas (2001) and Jérémie Gallien (2006) consider the same basic problem as above,

[5] Gabriel R. Bitran and Susana V. Mondschein (1997) obtain similar results in a discrete time framework.

but in a framework with an infinite horizon and discounting. In this case, the revenue maximizing posted prices - again one price for each size of inventory -, turn out to be stationary: they do not depend on time.[6] This stationarity allows Gallien to offer a reasonable sufficient condition ensuring that all sales occur immediately upon arrival, even if the arriving agents can strategically delay their purchase. Also in an infinite horizon model, Ruqu Wang (1993) compares a posted-price regime with an auction one. In his framework, the auction option is costly.

Lazear (1986) argues that delayed purchases by strategic consumers is not likely to be significant if the number of consumers is large, but a more recent, still small, strand of the literature focuses on strategic buyer behavior in models with fixed inventories.[7] Volker Nocke and Martin Peitz (2007) exhibit conditions for a decreasing price path to be optimal even if consumers are strategic about the timing of purchases in a model with identical units and fixed demand. These authors do not consider pricing policies that condition on past realized sales. Yossi Aviv and Amit Pazgal (2008) assume that prices must be declining (which is not innocuous) in a two-period model with stochastic arrivals, but derive analytic results only for a policies that do not depend on past realized sales. Xuanming Su (2007) allows more general policies, but assumes a deterministic demand flow.

To conclude, as Bitran and Caldentey (2003) noted, due to the technical complexity, the literature on dynamic revenue maximization with stochastic demand has focused on models with identical objects, contrasting the present framework.[8] Moreover, no solution is yet available to the Gallego and van Ryzin (1994) type of models with strategic customers.

[6]Das Varma and Vettas consider a model with discrete time and deterministic arrivals, whereas Gallien and Arnold and Lippman have continuous time models with stochastic arrivals. The latter authors assume (rather than derive) the stationarity of posted prices, and also compare these to reservation prices in a model where arriving agents announce bids.

[7]This should be contrasted with the earlier literature on the so-called "Coase Conjecture" where the inventory can be replenished. Preston McAfee and Thomas Wiseman (2008) show that the introduction of capacity costs counters the Coasian insight. For other papers that examine the relationship between limited inventories and pricing schemes see Milton Harris and Artur Raviv (1981) and Charles A. Wilson (1988).

[8]See Gallego and Van Ryzin (1997) for an exception. Some models assume that customers belong to several known classes which allows the use of third-degree price discrimination.

I. The Model

There are n items (or objects) that need to be allocated to arriving agents before a deadline $T > 0$. Each item i is characterized by a "quality" q_i. Agents (or buyers) arrive according to a Poisson process with intensity λ, and each can only be served upon arrival (i.e., agents are impatient). Upon arrival, each agent observes the set of available objects. After an item is assigned, it cannot be reallocated in the future. Each agent j is characterized by a "type" x_j. An agent with type x_j who obtains an item with quality q_i enjoys a utility of $q_i x_j$. The designer (or seller) seeks to maximize expected revenue.

While the items' types $0 \leq q_n \leq q_{n-1} \leq ... \leq q_1$ are assumed to be known constants, the agents' types are assumed to be represented by independent and identically distributed random variables X_i on $[0, +\infty)$ with common twice differentiable c.d.f. F (f denotes the corresponding density function). The realization of X_i is private information to agent i. We assume that each X_i has a finite mean, denoted by μ, and a finite variance. Moreover, we assume that $f(x) < \infty$ and that $x - \frac{1-F(x)}{f(x)}$ increases for any $x \in [0, \infty)$.[9]

The multiplicative structure employed above is restrictive: it is one of the simplest structures allowing a meaningful treatment of several qualities while allowing the modeling of private information as a one-dimensional variable. The present formulation is standard in the relevant literature, following the seminal static analysis in Mussa and Rosen (1978).

Another restrictive assumption pertains to the fact that agents are infinitely impatient. In particular, strategic considerations about the moment of purchase do not play a role in our present analysis.(This is not a constraint in the discounted infinite horizon model presented in Section 4.1 because of the stationarity of the optimal policies there- see also Gallien (2006).) Alternatively, we could assume that agents bear a high enough cost of waiting per unit of time.

II. Implementable Policies

Without loss of generality, we restrict attention to direct mechanisms where every agent, upon arrival, reports his characteristic x_i and where the mechanism specifies an allocation

[9] The expression $x - \frac{1-F(x)}{f(x)}$ is the so-called *virtual valuation* of the buyer of type x. For a detailed discussion see Roger B. Myerson (1981) and Jeremy Bulow and John Roberts (1989).

(which item, if any, the agent gets) and a payment. The schemes we develop also have an obvious and immediate interpretation as indirect mechanisms, where the designer sets a time-dependent menu of prices, one for each item, and where arriving agents choose out that menu.

An allocation policy is called *non-randomized* and *Markovian* if, at any time t, and for any possible type of agent arriving at t, it uses a non-random allocation rule that only depends on the arrival time t, on the declared type of the arriving agent, and on the set of items available at t, denoted by Π_t. Thus, the policy depends on past decisions only via the state variable Π_t. We also restrict attention to interim-individually rational policies, where no agent ever pays more than the utility obtained from the physical allocation.

Denote by $Q_t : [0, +\infty) \times 2^{\Pi_0} \to \Pi_0 \cup \emptyset$ a non-randomized Markovian allocation policy for time t with an additional requirement that for any $A \subseteq 2^{\Pi_0}$ and $x \in [0, +\infty)$, $Q_t(x, A) \in A \cup \emptyset$. That is for any set of the available objects, the allocation policy will assign either one of the available objects or no object at all. Denote by $P_t : [0, +\infty) \times 2^{\Pi_0} \to \mathbb{R}$ the associated payment rule and by k_t the cardinality of set Π_t. Finally, denote by $q_{(j:\Pi_t)}$ the j'th highest element of the set Π_t.

The next Proposition shows that a non-randomized, Markovian allocation policy is implementable if and only if it is based on a partition of the agents' type space.[10] In other words, implementability reduces here to setting a menu of prices, one for which object, from which the arriving agent has to choose.

Proposition 1 *Assume that Π_t is the set of objects available at time t, and assume that $q_j \neq q_k$ for any $q_j, q_k \in \Pi_t$, $j \neq k$.*

1. *A non-randomized, Markovian policy Q_t is implementable if and only if there exist $k_t + 1$ functions $\infty = y_{0,\Pi_t}(t) \geq y_{1,\Pi_t}(t) \geq y_{2,\Pi_t}(t) \geq \cdots \geq y_{k_t,\Pi_t}(t) \geq 0$, such that $x \in [y_{j,\Pi_t}(t), y_{j-1,\Pi_t}(t)) \Rightarrow Q_t(x, \Pi_t) = q_{(j:\Pi_t)}$ and $x < y_{k_t,\Pi_t}(t) \Rightarrow Q_t(x, \Pi_t) = \emptyset$.*[11]

[10] The result holds for any deterministic policy. But, since the rest of the analysis focuses on the Markov case, and in order to save on notational complexity, we consider only this case here.

[11] Types at the boundary between two intervals can be assigned to either one of the neighboring elements of the partition. That is, if $x_i \in \{y_{k_t,\Pi_t}(t), y_{k_t-1,\Pi_t}(t), ..., y_{2,\Pi_t}(t), y_{1,\Pi_t}(t)\}$, then $Q_t(y_{i,\Pi_t}(t), \Pi_t) \in \{q_{(i:\Pi_t)}, q_{(i+1:\Pi_t)}\}$, $i = 1, 2, .., k_t$.

2. *The associated payment scheme is given by* $P_t(x, \Pi_t) = \sum_{i=j}^{k_t}(q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)})y_{i,\Pi_t}(t) + S(t)$ *if* $x \in [y_{j,\Pi_t}(t), y_{j-1,\Pi_t}(t))$ *where* $S(t)$ *is some allocation- and type-independent function.*[12]

Proof. See Appendix. ■

The payment scheme has an intuitive interpretation: assume for a moment that the analyzed setup is the static one with k_t objects and $k_t + 1$ agents where, in addition to an agent with type x, there are k_t other "dummy" agents with types $y_{1,\Pi_t}(t), y_{2,\Pi_t}(t), \cdots, y_{k_t,\Pi_t}(t)$. The payment for the object with the j-highest quality, $\sum_{i=j}^{k_t}(q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)})y_{i,\Pi_t}(t)$, represents the externality imposed by the agent with type x_i on the dummy agents in the corresponding efficient allocation.

III. Dynamic Revenue Maximization

In this section we solve the dynamic revenue maximization problem. A main feature that differentiates our analysis from previous ones is the fact that we use the *mechanism design approach* developed in Section 4, and the insight behind the celebrated *payoff/revenue equivalence* theorem. Thus, we focus on the dynamic allocation policy that underlies revenue maximization, while pricing plays only a "secondary" role since, once the allocation is fixed, it is automatically induced by the implementation requirements.

Without loss of generality, we can restrict attention to Markovian, non-randomized policies where the state includes the set of available objects Π_t, the period of time t , and the type of the agent that arrives at t. The optimality of Markovian, possibly randomized, policies is standard for all models where, as is the case here, the instantaneous rewards and transition probabilities are history-independent - see for example Theorem 11.1.1 in Martin L. Puterman, (2005) which shows that, for any history-dependent policy, there is a Markovian, possibly randomized, policy with the same payoff.[13] Given a Markovian

[12] If there are some identical objects, there exist implementable policies that do not take the form of partitions. But, for each such policy, there exists another implementable policy that is based on a partition, and that generates the same expected utility for all agents and for the designer.

[13] If the optimal policy is non-Markov, there are different histories after which the state is the same but the taken action differs.For each state s and for any history H, let $p(H|s)$ be the *ex-ante* probability that history H occurs, conditional on the reached state being s. Modify now the policy in the following way: at time t the new policy uses a lottery where the action used in the original policy after history

policy, at each period t the designer's problem is equivalent to a static problem where one object out of a given inventory needs to be allocated to a privately informed agent, and where the seller has different salvage values for the remaining possible inventories (the salvage values correspond to the various continuation values in the dynamic case). Analogously to Myerson (1981), the static revenue-maximization problem has a non-randomized solution as long as the agent's virtual valuation is increasing (as assumed here): if at all, the agent should get the object for which virtual valuation plus salvage value is highest.[14] Thus, at each period t in the dynamic problem the seller has a non-randomized optimal policy as well.

We first calculate the expected revenue for any given Markovian, non-randomized allocation policy, and then we use a variational argument to derive the cut-off curves describing the revenue-maximizing dynamic policy.

Recall from Proposition 1 that, in order to implement a Markovian, non-randomized allocation which is given by $\infty = y_{0,\Pi_t}(t) \geq y_{1,\Pi_t}(t) \geq y_{2,\Pi_t}(t) \geq \cdots \geq y_{k_t,\Pi_t}(t) > 0$, $\forall t$, the price at period t for the object with the j-th highest characteristic (among the remaining objects) needs to be

$$P_t^{(j)}(\Pi_t) = \sum_{i=j}^{k_t} (q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)}) y_{i,\Pi_t}(t) + S(t). \tag{1}$$

In any interim individually rational mechanism we must have $S(t) \leq 0$ for any t, and, in order to maximize the revenue, we must clearly have $S(t) = 0$.

After using simple properties of sampling out of Poisson processes (see the Proof of Proposition 7 in the Appendix), the expected revenue at time t where $\Pi_t \neq \emptyset$ takes the form

$$R(\Pi_t, t) = \sum_{i=1}^{k_t} \int_t^T \left(P_s^{(i)}(\Pi_t) + R\left(\Pi_t \backslash \{q_{(i:\Pi_t)}\}, s\right)\right) h_{i,\Pi_t}(s)\, ds$$

where

$$h_{i,\Pi_t}(s) = \lambda \left[F\left(y_{i-1,\Pi_t}(s)\right) - F\left(y_{i,\Pi_t}(s)\right)\right] e^{-\int_t^s \lambda[1-F(y_{k_t,\Pi_t}(z))]dz}$$

H is taken with probability $p(H|s)$. If the original policy was optimal, then it need to be incentive compatible for **any** history H. By construction, the modified policy will be incentive compatible as well.

[14]The expected revenue from any incentive compatible mechanism in the static problem is given by $\int \sum_{q_i \in \Pi_t \cup \varnothing} \Pr_i(x) \left[q_i \left(x - \frac{1-F(x)}{f(x)}\right) + SV(\Pi_t \backslash q_i)\right] dx$ where $\Pr_i(x)$ is the probability that the designer assigns to type x the object of quality q_i , and where $SV(\Pi_t \backslash q_i)$ is the salvage value of the set of objects $\Pi_t \backslash q_i$.

is the density of the waiting time till the first arrival of an agent with a type in the interval $[y_{i,\Pi_t}(s), y_{i-1,\Pi_t}(s))$ given that no arrival that leads to a sale (e.g., type above $y_{k_t}(s)$) has occurred.

Recall that a Markovian, non-randomized policy must specify an allocation decision for each possible state, i.e., for each possible subset of object $\Pi_t \neq \emptyset$ available at time t. Moreover, for each state, the policy consists of $k_t = |\Pi_t|$ cut-off curves that describe the partition of the set of agents' types - generally these curves depend on the precise composition of the set Π_t. The number of needed curves if there are n objects is $\sum_{k=1}^{n} k\binom{n}{k} = n2^{n-1}$. This yields 4 cut-off curves for two objects, 12 curves for three objects, 32 curves for 4 objects, and so on... In order to save on notation and to keep the somewhat involved proofs more transparent, we assume below that there are only two objects with characteristics $q_1 \geq q_2$. But we will describe the completely analogous solution to the revenue maximization problem for the general case with any number of distinct objects. A main result is that the dynamic revenue maximizing policy for n (possibly distinct) objects is, in fact, completely described by only n cutoff curves. In particular, it shows that this policy is *independent* of the characteristics of the available objects.

With slight abuse of notation, we write "2" instead of $\Pi_t = \{q_1, q_2\}$ at the second subscript of the allocation functions $y_{i,\Pi_t}(t)$ whenever $k_t = 2$. This should not lead here to any confusion.

Proposition 2 *The dynamic revenue maximizing allocation policy is independent of the qualities of available objects q_1 and q_2. In particular, we have:*

1. $y_{1,q_1}(t) = y_{1,q_2}(t) = y_{1,2}(t) := y_1(t)$ *where* $y_1(t)$ *is a solution of*
$$y_1(t) = \frac{1 - F(y_1(t))}{f(y_1(t))} + \lambda \int_t^T \frac{[1 - F(y_1(s))]^2}{f(y_1(s))} ds$$

2. $y_{2,2}(t) := y_2(t)$ *is a solution of*
$$y_2(t) = \frac{1 - F(y_2(t))}{f(y_2(t))} + \lambda \int_t^T \frac{[1 - F(y_2(s))]^2}{f(y_2(s))} ds - R(1,t)$$
where
$$R(1,t) = \lambda \int_t^T \frac{[1 - F(y_1(s))]^2}{f(y_1(s))} ds$$
is the expected revenue at time t if there is one available object with $q = 1$ and the optimal policy will be followed from time t on.

Proof. See Appendix. ■

Remark 1 *Let us explore in some detail the intuition for the result whereby the optimal cutoffs do not depend on qualities. Assume that there are two available objects* $q_1 > q_2$*, and that at time* t *the cutoffs are* $y_1^o > y_2^o$*. Consider the effect of small shift in the highest cut-off from* y_1^o *to* $y_1^o + \epsilon$*. This shift has any effect only if some agent arrives at* t*. Moreover, the shift has no effect on the expected revenue if the arriving agent has value below* y_1^o*. If at time* t *an agent with value* y_1^o *arrives, then the shift switches the object he gets from* q_1 *to* q_2 *- which implies that he has to pay* $P_t^{(2)}(\{q_1, q_2\})$ *instead of* $P_t^{(1)}(\{q_1, q_2\})$ *- and also switches the object that remains available for the future allocation from* q_2 *to* q_1*. The infinitesimal effect is*

$$\begin{aligned} & f(y_1^o)\left(P_t^{(2)}(\{q_1,q_2\}) + q_1 R(1,t) - P_t^{(1)}(\{q_1,q_2\}) - q_2 R(1,t)\right) \\ = \; & (q_1 - q_2) f(y_1^o)(R(1,t) - y_1^o). \end{aligned}$$

The equality in the above equation follows here from the fact that in any incentive compatible mechanism we must have $P_t^{(2)}(\{q_1, q_2\}) - P_t^{(1)}(\{q_1, q_2\}) = (q_2 - q_1) y_1^o$*. Recall also that* $P_t^{(1)}(\{q_1, q_2\}) = (q_1 - q_2) y_1^o + q_2 y_2^o$*. Therefore, the shift increases the price that will be charged to all agents with type above* $y_1^o + \epsilon$*, since supporting a more conservative allocation of the best available object requires charging a higher price to all types that should get this object. Therefore, increasing the cut-off* y_1^o *also yields a higher revenue if an agent with value above* $y_1^o + \epsilon$ *arrives. This effect is*

$$\begin{aligned} & (1 - F(y_1^o + \epsilon))((q_1 - q_2)(y_1^o + \epsilon) + q_2(y_2^o + R(1,t))) \\ & - (1 - F(y_1^o + \epsilon))((q_1 - q_2) y_1^o + q_2(y_2^o + R(1,t))) \\ = \; & (q_1 - q_2)(1 - F(y_1^o + \epsilon))\epsilon \end{aligned}$$

For an infinitesimal change, this effect becomes $(q_1 - q_2)(1 - F(y_1^o))$*. To sum up, the total effect of the shift on expected revenue is*

$$(q_1 - q_2)\left[(1 - F(y_1^o)) - f(y_1^o)(y_1^o - R(1,t))\right].$$

The expression is linear in the difference $(q_1 - q_2)$ *and the optimal* y_1^o *- where the total effect of the shift should be equal to zero - does not depend on the characteristics of the available objects.*

Remark 2 *The equations for the revenue maximizing cutoff curves have an intuitive interpretation. Assume first that only one object with $q = 1$ is available. The allocation policy is described by the equation*

$$y_1(t) - \frac{1 - F(y_1(t))}{f(y_1(t))} = \lambda \int_t^T \frac{[1 - F(y_1(s))]^2}{f(y_1(s))} ds = R(1, t).$$

On the left hand side, we have the virtual valuation of an agent with type $y_1(t)$. As Claim 1 showed, the right hand side represents the expected revenue from time t on if the object is not sold at t, given that an optimal allocation policy is followed from time t on. Since the seller is able to extract as revenue only the virtual valuation of an arriving buyer, the equation shows that the optimal cut-off curve satisfies an indifference condition between immediate selling and a continuation that uses the optimal policy.

In the general case, if there are $k_t = |\Pi_t|$ available objects, then, no matter what their types are, the i'th cut-off curve, $1 \leq i \leq k_t$, in the dynamic revenue-maximizing policy is given by

$$y_i(t) - \frac{1 - F(y_i(t))}{f(y_i(t))} + \lambda \int_t^T \frac{[1 - F(y_{i-1}(s))]^2}{f(y_{i-1}(s))} ds = \lambda \int_t^T \frac{[1 - F(y_i(s))]^2}{f(y_i(s))} ds \qquad (2)$$

or, equivalently, by

$$y_i(t) - \frac{1 - F(y_i(t))}{f(y_i(t))} + R(\mathbf{1}_{i-1}, t) = R(\mathbf{1}_i, t) \qquad (3)$$

where $\mathbf{1}_i$ is the set of $1'$s of cardinality i and

$$R(\mathbf{1}_j, t) = \lambda \int_t^T \frac{[1 - F(y_j(s))]^2}{f(y_j(s))} ds \qquad (4)$$

is the expected revenue at time t from the optimal cut-off policy if j identical objects with $q = 1$ are still available. Since there will be sales only to the agents with positive virtual valuations, equation (3) implies that $R(\mathbf{1}_i, t) > R(\mathbf{1}_j, t)$ for any $i > j \geq 0$ and $t < T$.

While equation (3) has been obtained for the case of identical objects in the revenue-management literature (see for example Gallego and van Ryzin (1994), and Bitran and Mondschein (1997) for a discrete time model), the explicit expression in (4) is new, a by-product of our analysis that focused on the allocation policy rather than on prices.

For the general case with several distinct objects, note also that, if an object is sold at time t, then the lowest among the current optimal cut-off curves becomes irrelevant regardless of the characteristic of the sold object, while all the other $k_t - 1$ cutoff curves *do not change and remain relevant for the future allocation decisions. That is, the optimal cutoff curves depend only on the cardinality of* Π_t, k_t. *For any two sets of available objects* Π_t^1 *and* Π_t^2 *with* $k_t^1 = |\Pi_t^1|$ *and* $k_t^2 = |\Pi_t^2|$, *and for any* $1 \le i \le \min\{k_t^1, k_t^2\}$ *it holds that*

$$y_{i,\Pi_t^1}(t) = y_{i,\Pi_t^2}(t).$$

If Π_t *is a set of identical objects, then only the lowest cut-off curve* $y_{k_t}(t)$ *where* $k_t = |\Pi_t|$ *is relevant for the allocation decision.*

Note that in the optimal mechanism prices for the remaining objects increase after each sale. This follows because: 1) The invariance property implies that the allocation policy after a sale of any of k available objects will be based on the $k-1$ pre-sale highest cutoffs. This more conservative allocation is implemented via higher prices for all objects. 2) Each sale increases the difference between the remaining qualities, which, in turn, leads to an increase in the prices of all objects with qualities higher than the one just sold. As a consequence, arriving buyers do not necessarily have an incentive to delay their purchase because the average price may well go up for a period of time, before inevitably going down if the good is not sold and the deadline approaches.

A. Infinite Horizon with Discounting

In this subsection we very briefly present the characterization of the dynamic revenue maximizing allocation scheme in an infinite horizon setting. We assume here that the utility of the designer from a payment at time t is discounted at rate $e^{-\alpha t}$, where $\alpha > 0$ is a discount factor. For the case of identical objects, Gallien (2006) has shown that the optimal cutoff curves are stationary (i.e., time independent).

We show below that the revenue maximizing allocation policy for several distinct objects does not depend on the characteristics of the available objects. The analysis of the discounted, infinite-horizon case is similar and easier than the one we performed for the deadline case, and we omit the proof of the next Proposition.

Proposition 3 *The dynamic revenue- maximizing policy consists of* $n+1$ *constants* $y_n \le y_{n-1} ... \le y_1 \le y_0 \equiv \infty$ *that do not depend on the q's such that:*

1. *If an agent with type* x *arrives at a time* t, *it is optimal to assign to that agent the* j*'th highest element of* Π_t *if* $x \in [y_j,\ y_{j-1})$, *and not to assign any object if* $x < y_{k_t}$, *where* $k_t = |\Pi_t|$.

2. The constants y_j *satisfy:*

$$y_j - \frac{1-F(y_j)}{f(y_j)} + \frac{\lambda}{\alpha}\frac{[1-F(y_{j-1})]^2}{f(y_{j-1})} = \frac{\lambda}{\alpha}\frac{[1-F(y_j)]^2}{f(y_j)}$$

where

$$R(\mathbf{1}_j) = \frac{\lambda}{\alpha}\frac{[1-F(y_j)]^2}{f(y_j)}$$

is the expected revenue at time t *if* j *identical objects with* $q_1 = q_2 = ... = q_j = 1$ *are available at time* t*, given that the designer uses the optimal allocation policy from time* t *on.*

Since in the infinite horizon model prices only go up while the inventory is gradually depleted, buyers do not have here an incentive to delay their purchase after the arrival time.

IV. Applications

A. Dynamic Pricing and Optimal Inventory Choice

The entire above analysis was based on the assumption that the inventory's size and composition was exogenously given. In this Subsection we show how our previous insights can be immediately used as a building block for a more general analysis where the seller can choose both the number of objects and the spectrum of offered qualities before sales start. This timing assumption is appropriate in cases where the necessary production lead-times are significant in comparison to the retail period. Our present treatment extends the seminal Mussa-Rosen (1978) analysis to the framework with random arrivals. Gallego and van Ryzin (1994) characterized the optimal inventory size in a model with deterministic demand, and showed that the solution coincides with the optimal size of the inventory in the original problem if the horizon goes to infinity. In contrast, we are able to characterize the optimal inventory size (and its composition) in the original problem for any finite horizon.

Since our analysis identified an invariant of the optimal selling scheme - the allocation policy - we can easily solve the larger problem backwards and characterize the optimal quality choice. Formally, the decision of the monopolist is to choose the optimal package of qualities Π_0 :

$$\max_{\Pi_0} R(\Pi_0, 0) - C(\Pi_0)$$

where $C(\Pi)$ denotes cost of producing the package of qualities Π. It is plausible to assume that the cost function is convex and symmetric (symmetry means here that the cost of producing any permutation of a given package of qualities is the same as the cost of producing the original package). Note that any convex and symmetric function is *Schur-convex.*[15] We assume below a very simple, separable form of Schur-convexity for the cost function, but the reader will have no difficulty adding the technical details needed for the more general result that does not assume separability.

Proposition 4 *Let $y = \{y_i(t)\}_{i=1}^n$ denote the allocation underlying the revenue maximizing policy with n objects, and assume that, for all n and for all vectors of qualities $(q_1, q_2, ..q_n)$, $C(q_1, q_2, ..q_n) = \sum_{i=1}^n \phi(q_i)$,where the function $\phi : R \to R$ is strictly increasing, convex and satisfies $\phi(0) = 0$. Then :*

1. *The optimal number of objects n^* is characterized by*[16]

$$\phi'(0) \in (y_{n^*+1}(0) - \frac{1 - F(y_{n^*+1}(0))}{f(y_{n^*+1}(0))},\ y_{n^*}(0) - \frac{1 - F(y_{n^*}(0))}{f(y_{n^*}(0))}] \tag{5}$$

2. *The optimal qualities q_i^* are given by:*

$$\phi'(q_i^*) = y_i(0) - \frac{1 - F(y_i(0))}{f(y_i(0))},\ i = 1, ..., n^* \tag{6}$$

Proof. See Appendix. ■

The above proposition and its analog for welfare maximization can be used to compare the quality/quantity optimal choices under the two regimes. For example, it is intuitive that a revenue maximizing monopolist will produce lower qualities than a welfare maximizing one, as in the Mussa-Rosen framework. Note though that some insights will differ from Mussa and Rosen's due to a difference in timing: while we assume that qualities are produced in advance - before buyers' characteristics get revealed - Mussa and Rosen implicitly assume that quality can be produced contingent on these. For instance, in our setup there is a distortion even in the quality provided to the highest type.

[15] Schur-convexity guarantees here that a higher marginal revenue is associated with a higher chosen quality for the respective object.

[16] It should be obvious how to adjust the formula in order to deal with the extreme cases where $n^* = 0, \infty$

B. The Pattern of Prices in Clearance Sales

Our characterization generates a wealth of empirically testable implications about the pattern of observed prices associated with the revenue maximizing policy. These implications can be compared to available data for relevant situations. For example, an important finding in the well-known empirical study of clearance sales for apparel conducted by Pashigian and Bowen (1991) p.1018 is that :

"More expensive apparel items within each product line are frequently sold at a higher average percentage markdown".

Pashigian and Bowen attempt to explain this phenomenon by using Lazear's (1986) theory of retail pricing and clearance sales. Lazear's theory only deals with the sale of a homogenous products and it cannot incorporate the parallel sale of several substitutes, as is common practice in most stores. Thus, the offered theoretical explanation for the empirical finding is not entirely convincing.[17] The next result shows how the empirical observation agrees with a fairly general prediction of our model that holds for any two different qualities, for any positive levels of inventories before and after the clearance sale, and for any distribution of values.

Proposition 5 *Assume that the seller uses the revenue maximizing policy in a situation where at time $t = 0$ there are $n_1 > 0$ items of quality q and $n_2 > 0$ items of quality s, $s < q$. Assume also that at time $t = T$ there are $0 < l_1 \leq n_1$ items of quality q , and $0 < l_2 \leq n_2$ items of quality s left unsold. Then the percentage markdown - defined as the difference between the prices of the same product at $t = 0$ and $t = T$ divided by the price at $t = 0$ - is always higher for the higher quality.*

Proof. See Appendix. ■

We conclude this subsection with another related observation that can also be taken to the data:

Proposition 6 *Consider the revenue maximizing policy. Then, at any point in time t where two different qualities are available, the price of the higher quality decreases quicker than the price of the lower quality.*

[17] For some special forms of the distribution of values, Lazear's theory predicts that increasing the dispersion in values for the product will lead to an increased markdown. Pashigian and Bowen identify more expensive items (which contain more elements of fashion and style) with an increased dispersion.

Proof. See Appendix. ■

V. Conclusion

We have studied the dynamic, revenue maximizing allocation of finite inventories consisting of substitute qualities to a stream of privately informed agents that arrive randomly before a deadline. The paper's main contributions were:

1. The integrated model allowed us to add quantity/quality decisions to the traditional pricing considerations for identical objects.

2. We have introduced a new method - based on a combination of the payoff equivalence principle and a variational argument - to the analysis of inventory pricing under random, sequential demand.

3. The analysis yields clear, testable implications about the pattern of optimal evolution of prices over time. These predictions can be taken to the data.

An important assumption has been that agents' purchase decisions are not strategic in their time dimension.[18] Of course, if all agents can delay their purchase till the deadline without cost, the optimal mechanism will be an auction that gathers all potential buyers - this is the opposite extreme assumption to the one made in this paper where buyers are infinitely impatient. It is of major theoretical and applied interest to study revenue maximizing schemes that take into account the agents' incentives to possibly delay their transactions in less extreme cases. Since, as sales occur, prices may go up and inventories get depleted, the decision whether to delay purchase to a later time is complex. The same applies to the seller's reaction to such buyer behavior which may need to include mechanisms that are more complex than menus of posted prices or standard auctions.

Other major - but also difficult - generalizations would be the introduction of heterogeneous objects that are not necessarily similarly ranked by all buyers, or the introduction of multi-unit demand. The latter would allow the study of dynamic bundling policies. Such models need a multi-dimensional specification of private information.

[18] Recall that the model with infinite horizon is immune to this restriction.

Finally, our analysis has considered a monopolist seller, but it would be very interesting to characterize dynamic pricing equilibria among oligopolists that face randomly arriving buyers who can decide where to buy.

Appendix

Proof of Proposition 1. :

$\Longrightarrow$ If two reports of the agent that arrives at t lead to the same physical allocation, then, in any incentive compatible mechanism, the associated payments should be the same as well. Denote by P_j the payment that will be charged for the object with quality q_j. A direct mechanism is equivalent to a mechanism where the agent arriving at time t chooses an object and a payment from a menu $(q_j, P_j)_{j=1}^{k_t}$. If some type x prefers the pair (q_k, P_k) over any other pair (q_l, P_l) with $q_k > q_l$, then any type $\widetilde{x} > x$ also prefers (q_k, P_k) over (q_l, P_l). This implies that $Q_t(\widetilde{x}, \Pi_t) \geq Q_t(x, \Pi_t)$ for any t and Π_t. Finally, noting that $Q_t(x, \Pi_t) = \emptyset$ is equivalent to allocating an object with quality equal to zero, implies that an agent who arrives at time t gets object $q_{(k)}$ if he reports a type contained in the interval $(y_{k,\Pi_t}(t), y_{k-1,\Pi_t}(t))$. A similar argument shows that $Q_t(y_{i,\Pi_t}(t), \Pi_t) \in \left\{q_{(i+1:\Pi_t)}, q_{(i:\Pi_t)}\right\}$ for $i \in \{1, 2, ..., k_t\}$.

$\Longleftarrow$ The proof is constructive: given a partition-based policy, we design a payment scheme $P_t(x, \Pi_t)$ that, for any $j \in \{1, ..., k_t\}$, induces type $x \in [y_{j,\Pi_t}(t), y_{j-1,\Pi_t}(t))$ to choose the object with type $q_{(j:\Pi_t)}$. Without loss of generality, we assume that an agent whose type is on the boundary between two intervals in the partition chooses the item with higher type. Consider then the following payment scheme

$$P_t(x, \Pi_t) = \sum_{i=j}^{k_t}(q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)})y_{i,\Pi_t}(t) + S(t) \text{ if } x \in [y_{j,\Pi_t}(t), y_{j-1,\Pi_t}(t)) \tag{A1}$$

where $S(t)$ is is some allocation- and type-independent function. Note that type $x = y_{j,\Pi_t}(t)$ is indifferent between $\left(q_{(j:\Pi_t)}, P_j\right)$ and $\left(q_{(j+1:\Pi_t)}, P_{j+1}\right)$. Moreover, any type above $y_{j,\Pi_t}(t)$ prefers $\left(q_{(j:\Pi_t)}, P_j\right)$ over $\left(q_{(j+1:\Pi_t)}, P_{j+1}\right)$, while any type below prefers $\left(q_{(j+1:\Pi_t)}, P_{j+1}\right)$ over $\left(q_{(j:\Pi_t)}, P_j\right)$. Therefore, any type $x \in [y_{j,\Pi_t}(t), y_{j-1,\Pi_t}(t))$ prefers $\left(q_{(j:\Pi_t)}, P_j\right)$ over any other pairs in the menu. ■

Proof of Proposition 2: The proof proceeds by a sequence of three arguments: first, we derive the expected revenue for any Markovian, non-randomized allocation policy (Proposition 7); second, we derive the revenue-maximizing cutoff curves when only one object remains (Claim 1); finally, we derive the revenue maximizing allocation policy if two objects are left (Claim 2).

Proposition 7 *Assume that*

1. *If $k_t = 2$, the designer uses the dynamic allocation cutoff-curves $y_{2,2}(t) \leq y_{1,2}(t)$, i.e., the agent that arrives at time t gets: the object with quality q_1 if his type is $x_i \geq y_{1,2}(t)$; the object with quality q_2, if his type is $x_i \in [y_{2,2}(t), y_{1,2}(t))$; no object if $x_i < y_{2,2}(t)$.*

2. *If $k_t = 1$, the designer uses the dynamic cutoff–curves $y_{1,q_j}(t)$, i.e., the agent that arrives at time t gets the remaining object with characteristic p_j if $x_i \geq y_{1,q_j}(t)$, and no object otherwise.*

   Then, the expected revenue from this policy is given by

$$\begin{aligned}&\int_0^T \left(q_2 y_{2,2}(t) + R(q_1,t)\right)\lambda\left(1 - F(y_{2,2}(t))\right) e^{-\int_0^t \lambda(1-F(y_{2,2}(s)))ds} dt \ + \\ &\int_0^T \left((q_1 - q_2)\, y_{1,2}(t) + R(q_2,t) - R(q_1,t)\right) \cdot \\ &\lambda\left(1 - F(y_{1,2}(t))\right) e^{-\int_0^t \lambda(1-F(y_{2,2}(s)))ds} dt\end{aligned}$$

   where

$$R(q_j,t) = q_j \int_t^T y_{1,q_j}(s)\, \lambda(1 - F(y_{1,q_j}(s)) e^{-\int_t^s \lambda[1-F(y_{1,q_j}(z))]dz} ds \qquad \text{(A2)}$$

   is the expected revenue at time t if only one object with quality q_j remains, given that the dynamic allocation function y_{1,q_j} is used from t on.

Proof of Proposition 7. If only one object with characteristic q_i is available at time t, then the expected revenue is given by

$$q_i \int_t^T y_{1,q_i}(s) h_{1,q_i}(s) ds$$

where $h_{1,q_i}(s)$ represents the density of the waiting time till the first arrival of an agent with a value that is at least $y_{1,q_i}(s)$. Note that this density is equal to the density of the first arrival in a non-homogenous Poisson process with rate $\lambda(s)(1 - F(y_{1,q_i}(s))$. The density of the time of the $n-$th arrival in a non-homogenous Poisson process with rate $\delta(s)$ is given by (see Ross (1983))

$$g_n(s) = \delta(s) e^{-m(s)} \frac{m(s)^{n-1}}{(n-1)!}\ , \text{ where } m(s) = \int_t^s \delta(z) dz \qquad \text{(A3)}$$

Thus, in our case, we obtain

$$h_{1,q_i}(s) = \lambda(s)(1 - F(y_{1,q_i}(s)) e^{-\int_t^s \lambda(z)[1-F(y(z))]dz} \text{ for } t \leq s \leq T$$

and (A2) follows.

If two objects are still available, the expected revenue is given by

$$\int_0^T \left[P_t^{(2)}\left(\{q_1,q_2\}\right)+R\left(q_1,t\right)\right] h_{2,2}(t)dt+\int_0^T \left[P_t^{(1)}\left(\{q_1,q_2\}\right)+R\left(q_2,t\right)\right] h_{1,2}(t)dt \quad \text{(A4)}$$

Here $h_{1,2}(t)$ represents the density of the waiting time till the first arrival of an agent with a value that is at least $y_{1,2}(t)$ if no arrival of an agent with value in the interval $[y_{2,2}(t), y_{1,2}(t))$ has occurred. Similarly, $h_{2,2}(t)$ represents the density of the waiting time till the first arrival of an agent with a value in the interval $[y_{2,2}(t), y_{1,2}(t))$ if no arrival of an agent with value in the interval $[y_{1,2}(t), \infty)$ has occurred. Since the arrival processes of agents with types in the intervals $[y_{2,2}(t), y_{1,2}(t))$ and $[y_{1,2}(t), \infty)$, respectively, are independent non-homogenous Poisson processes (see Proposition 2.3.2 in Ross (1983)), using (A3) we obtain

$$\begin{aligned} h_{1,2}(t) &= \lambda\left(1-F(y_{1,2}(t))\right) e^{-\int_0^t \lambda[1-F(y_{1,2}(s))]ds} e^{-\int_0^t \lambda[F(y_{1,2}(s))-F(y_{2,2}(s))]ds} \\ &= \lambda\left(1-F(y_{1,2}(t))\right) e^{-\int_0^t \lambda[1-F(y_{2,2}(s))]ds} \end{aligned}$$

and

$$\begin{aligned} h_{2,2}(t) &= \lambda\left(F(y_{1,2}(t))-F(y_{2,2}(t))\right) e^{-\int_0^t \lambda[F(y_{1,2}(s))-F(y_{2,2}(s))+1-F(y_{1,2}(s))]ds} \\ &= \lambda\left(F(y_{1,2}(t))-F(y_{2,2}(t))\right) e^{-\int_0^t \lambda[1-F(y_{2,2}(s))]ds} \end{aligned}$$

Finally, recall that incentive compatibility implies that

$$P_t^{(2)}\left(\{q_1,q_2\}\right)=q_2 y_{2,2}(t) \text{ and } P_t^{(1)}\left(\{q_1,q_2\}\right)=q_2 y_{2,2}(t)+\left(q_1-q_2\right) y_{1,2}(t),$$

Plugging the expressions for $P_t^{(2)}\left(\{q_1,q_2\}\right), P_t^{(1)}\left(\{q_1,q_2\}\right), h_{1,2}(t)$ and $h_{2,2}(t)$ into the expression for expected revenue (A4) yields the required formula. ■

Claim 1 *If only one object remains, the dynamic revenue maximizing allocation curve $y_1(t)$ solves*

$$y_1(t)=\frac{1-F\left(y_1(t)\right)}{f\left(y_1(t)\right)}+\lambda\int_t^T \frac{\left[1-F\left(y_1(s)\right)\right]^2}{f\left(y_1(s)\right)}ds. \quad \text{(A5)}$$

The expected revenue at time t where $\Pi_t=q_j$ is given by $R\left(q_j,t\right)=q_j R\left(1,t\right)$ where

$$R\left(1,t\right)=\lambda\int_t^T \frac{\left[1-F\left(y_1(s)\right)\right]^2}{f\left(y_1(s)\right)}ds \quad \text{(A6)}$$

Proof of Claim 1. If only the object with characteristic q_j is available, it follows from Proposition 7 that the expected revenue at time t is given by

$$R(q_j,t) = q_j \int_t^T y_{1,q_j}(s)\,\lambda(1-F(y_{1,q_j}(s))e^{-\int_t^s \lambda[1-F(y_{1,q_j}(z))]dz}ds.$$

Let $H(s) = \int_t^s \lambda[1-F(y_{1,q_j}(z))]dz$. Then, we obtain

$$R(q_j,t) = q_j \int_t^T F^{-1}\left[1-\frac{H'(s)}{\lambda}\right] H'(s)e^{-H(s)}ds.$$

This expression for revenue is appropriate for using a variational argument with respect to the function H. The corresponding necessary condition for the variational problem (i.e., the Euler-Lagrange equation) is

$$-(H'(s))^2 + 2H''(s) + \frac{H'(s)H''(s)f'\left(F^{-1}\left(1-\frac{H'(s)}{\lambda}\right)\right)}{\left(f\left(1-\frac{H'(s)}{\lambda}\right)\right)^2} = 0$$

Plugging back the expression for $H(s)$ gives

$$-\lambda[1-F(y_{1,q_j}(s))]^2 - 2f(y_{1,q_j}(s))y'_{1,q_j}(s) - \frac{[1-F(y_{1,q_j}(s))]f'(y_{1,q_j}(s))y'_{1,q_j}(s)}{f(y_{1,q_j}(s))} = 0$$

This implies that for any $s \in [0,T]$, the solution $y_{1,q_j}(s)$ should satisfy

$$-y'_{1,q_j}(s) - y'_{1,q_j}(s)\left(1+\frac{(1-F(y_{1,q_j}(s)))\,f'(y_{1,q_j}(s))}{(f(y_{1,q_j}(s)))^2}\right) = \lambda\frac{(1-F(y_{1,q_j}(s)))^2}{f(y_{1,q_j}(s))} \qquad \text{(A7)}$$

Since for any t , and for any differentiable $y(t)$ it holds that

$$-y'(t)\left(1+\frac{(1-F(y(t)))\,f'(y(t))}{(f(y(t)))^2}\right) = \frac{d}{dt}\left(\frac{1-F(y(t))}{f(y(t))}\right),$$

we can rewrite the necessary condition as

$$y'_{1,q_j}(s) + \lambda\frac{(1-F(y_{1,q_j}(s)))^2}{f(y_{1,q_j}(s))} = \frac{d}{ds}\left(\frac{1-F(y_{1,q_j}(s))}{f(y_{1,q_j}(s))}\right)$$

Taking now the integral between t and T yields

$$\begin{aligned} &\int_t^T y'_{1,q_j}(s)\,ds + \lambda\int_t^T \frac{(1-F(y_{1,q_j}(s)))^2}{f(y_{1,q_j}(s))}ds \\ = &\int_t^T \frac{d}{ds}\left(\frac{1-F(y_{1,q_j}(s))}{f(y_{1,q_j}(s))}\right)ds \end{aligned}$$

This is equivalent to:

$$
\begin{aligned}
&y_{1,q_j}(T) - y_{1,q_j}(t) + \lambda \int_t^T \frac{\left(1 - F(y_{1,q_j}(s))\right)^2}{f(y_{1,q_j}(s))} ds \\
= \; & \frac{1 - F(y_{1,q_j}(T))}{f(y_{1,q_j}(T))} - \frac{1 - F(y_{1,q_j}(t))}{f(y_{1,q_j}(t))}
\end{aligned}
$$

Together with the boundary condition

$$
y_{1,q_j}(T) - \frac{1 - F(y_{1,q_j}(T))}{f(y_{1,q_j}(T))} = 0
$$

we get (A5). The assumptions of increasing virtual type and finite density ensure that a solution to (A5) exists for any t.

To complete the proof and obtain the expression for revenue (A6), note that the expected revenue is given by $R(q_j, t) = q_j R(1, t)$ where

$$
R(1,t) = \int_t^T y_1(s) \lambda (1 - F(y_1(s)) e^{-\int_t^s \lambda [1 - F(y_1(z))] dz} ds
$$

Differentiating the above with respect to t gives

$$
R'(1,t) = \lambda (1 - F(y_1(t)) \left(R(1,t) - y_1(t)\right)
$$

It is then straightforward to verify that the function $\int_t^T \frac{[1-F(y_1(s))]^2}{f(y_1(s))} ds$ satisfies the above differential equation with the boundary condition $R(1,T) = 0$.

We proceed now to characterize the revenue-maximizing allocation policy if there are two objects left. ■

Claim 2 *If two objects remain, the dynamic revenue maximizing policy is characterized by two cutoff curves, $y_1(t)$ and $y_2(t)$, where $y_1(t)$ satisfies equation (A5) and where $y_2(t)$ satisfies:*

$$
y_2(t) = \frac{1 - F(y_2(t))}{f(y_2(t))} + \lambda \int_t^T \frac{[1 - F(y_2(s))]^2}{f(y_2(s))} ds - R(1,t) \tag{A8}
$$

Moreover, the expected revenue at time t for the case $\Pi_t = \{1, 1\}$ is given by

$$
R(\{1,1\}, t) = \lambda \int_t^T \frac{[1 - F(y_2(s))]^2}{f(y_2(s))} ds. \tag{A9}
$$

Proof of Claim 2. We split the proof into two cases **I.** We consider first the case where $q_1 > q_2$. That is, the seller needs to specify two different prices, and hence two different cutoff curves, $y_{1,2}(t)$ and $y_{2,2}(t)$. We can re-write the expected revenue given by Proposition 7 as

$$\begin{aligned}&\int_0^T \left(q_1 F^{-1}\left(1-\frac{H'(t)}{\lambda}\right)+q_2R(1,t)\right)H'(t)e^{-H(t)}dt\\&+(q_2-q_1)\int_0^T\left[F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)-R(1,t)\right]G'(t)e^{-H(t)}dt\end{aligned}$$

where

$$\begin{aligned}\int_0^t \lambda\left[1-F(y_{2,2}(s))\right]ds &:= H(t)\\ \int_0^t \lambda\left[1-F(y_{1,2}(s))\right]ds &:= G(t).\end{aligned}$$

The necessary conditions for the variational problem (i.e., the Euler-Lagrange equation) with respect to the functions $H(t)$ and $G(t)$, respectively, are:

$$\begin{aligned}&-(q_2-q_1)\,G'(t)\left[F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)-R(1,t)\right]-q_1\frac{\frac{1}{\lambda}\left(H'(t)\right)^2}{f\left(F^{-1}\left(1-\frac{H'(t)}{\lambda}\right)\right)}\\&-2q_1\frac{\frac{1}{\lambda}H''(t)}{f\left(F^{-1}\left(1-\frac{H'(t)}{\lambda}\right)\right)}+q_2R'(1,t)-q_1\frac{\frac{1}{\lambda^2}H'(t)H''(t)f'\left(F^{-1}\left(1-\frac{H'(t)}{\lambda}\right)\right)}{\left[f\left(F^{-1}\left(1-\frac{H'(t)}{\lambda}\right)\right)\right]^3}=0\end{aligned}$$

and

$$\begin{aligned}&-\frac{2\frac{1}{\lambda}G''(t)}{f\left(F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)\right)}-R'(1,t)-\frac{\frac{1}{\lambda^2}G'(t)G''(t)f'\left(F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)\right)}{\left[f\left(F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)\right)\right]^3}\\&-H'(t)\left[-\frac{\frac{1}{\lambda}G'(t)}{f\left(F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)\right)}+F^{-1}\left(1-\frac{G'(t)}{\lambda}\right)-R(1,t)\right]=0.\end{aligned}$$

Plugging the expressions for $H(t)$ and $G(t)$ allows us to write the necessary conditions in the following way:

$$\begin{aligned}&-(q_2-q_1)\,\lambda\left[1-F(y_{1,2}(t))\right](y_{1,2}(t)-R(1,t))-q_1\frac{\lambda\left[1-F(y_{2,2}(t))\right]^2}{f(y_{2,2}(t))}\\&-2q_1y'_{2,2}(t)-q_2R'(1,t)-q_1\frac{y'_{2,2}(t)\left[1-F(y_{2,2}(t))\right]f'(y_{2,2}(t))}{\left[f(y_{2,2}(t))\right]^2}=0\end{aligned}\tag{A10}$$

and

$$\begin{aligned}&\left[1-F(y_{2,2}(t))\right]\left[\frac{1-F(y_{1,2}(t))}{f(y_{1,2}(t))}-y_{1,2}(t)+R(1,t)\right]-2y'_{1,2}(t)\\&-R'(1,t)+\frac{y'_{1,2}(t)\left[1-F(y_{1,2}(t))\right]f'(y_{1,2}(t))}{\left[f(y_{1,2}(t))\right]^2}=0\end{aligned}\tag{A11}$$

Next, we show that a solution to the system of differential equations A10 and A11 is given by $y_{1,2}(t) = y_1(t)$ and $y_{2,2}(t) = y_2(t)$ where $y_1(t)$ and $y_2(t)$ solve the system of equations:

$$y_1(t) = \frac{1-F(y_1(t))}{f(y_1(t))} + \lambda \int_t^T \frac{[1-F(y_1(s))]^2}{f(y_1(s))} ds \tag{A12}$$

$$y_2(t) = \frac{1-F(y_2(t))}{f(y_2(t))} + \lambda \int_t^T \frac{[1-F(y_2(s))]^2}{f(y_2(s))} ds - R(t). \tag{A13}$$

Again, the assumptions of increasing virtual types and finite density guarantee the existence of solutions for (A12) and (A13). Note also that, since $R(t) \geq 0$, we must have $y_2(t) \leq y_1(t)$, $\forall t$.

Differentiation of (A12) with respect to t gives

$$2y_1'(t) = -y_1'(t)\frac{[1-F(y_1(t))]\, f'(y_1(t))}{[f(y_1(t))]^2} - \lambda\frac{[1-F(y_1(t))]^2}{f(y_1(t))}.$$

Plugging the above expression into (A11), and using the fact that

$$R'(1,t) = -y_1(t)\lambda(1-F(y_1(t))) + \lambda(1-F(y_1(t)))\, R(1,t) \tag{A14}$$

yields

$$\left[\lambda \int_t^T \frac{[1-F(y_1(s))]^2}{f(y_1(s))} ds - R(1,t)\right] [\lambda(1-F(y_1(t))) - (1-F(y_{2,2}(t)))] = 0$$

where last equality follows from Claim 1. Thus, we have showed that $y_{1,2}(t) = y_1(t)$ solves (A11) for any $y_{2,2}(t)$. We still need to show that $y_{1,2}(t) = y_1(t)$ and $y_{2,2}(t) = y_2(t)$ solve equation A10. Differentiation of (A13) with respect to t gives

$$2y_2'(t) = -y_2'(t)\frac{[1-F(y_2(t))]\, f'(y_2(t))}{[f(y_2(t))]^2} - \lambda\frac{[1-F(y_2(t))]^2}{f(y_2(t))} - R'(1,t).$$

Plugging this equality into (A10), we have to show that

$$-(q_2-q_1)\lambda[1-F(y_{1,2}(t))](y_{1,2}(t) - R(1,t)) - (p_2-p_1)R'(1,t) = 0.$$

For $y_{1,2}(t) = y_1(t)$, this equality holds by (A14).

II. We now consider the case with $q_1 = q_2 = q$. Since $R(q,t) = qR(1,t)$, Proposition 7 implies that we can rewrite the expected revenue as

$$q\int_0^T (y_{2,2}(t) + R(1,t))\,\lambda(1-F(y_{2,2}(t)))\, e^{-\int_0^t \lambda(1-F(y_{2,2}(s)))ds} dt.$$

The proof that the revenue maximizing cutoff curves are given by $y_1(t)$ and $y_2(t)$ is analogous to the above case, and we omit it here.

Proposition 7 implies then that

$$R(\{1,1\},t) = \int_t^T (y_2(s) + R(1,s))\, \lambda (1 - F(y_2(s)))\, e^{-\int_t^s \lambda(1-F(y_2(z)))dz} ds.$$

Differentiation with respect to t yields

$$R'(\{1,1\},t) = \lambda (1 - F(y_2(t)))\, (R(\{1,1\},t) - y_2(t) - R(1,t)). \tag{A15}$$

Recall that $y_2(t)$ solves

$$y_2(t) + R(1,t) = \frac{1 - F(y_2(t))}{f(y_2(t))} + \lambda \int_t^T \frac{[1 - F(y_2(s))]^2}{f(y_2(s))} ds \tag{A16}$$

Using equation (A16), it is easy to verify that $R(\{1,1\},t)$ given by equation (A9) satisfies differential equation (A15) with the boundary condition $R(\{1,1\},T) = 0$.

To complete the proof of the Proposition we have to show that the resulting cutoffs are implementable, i.e. $y_2(t) \leq y_1(t)$ for any $t \leq T$. Note that (A12) and (A13) imply that $y_1(t)$ is the solution to $y(t) = H(y(t))$ while $y_2(t)$ is the solution to $y(t) = G(y(t))$. Since $G(y) \leq H(y)$ holds for any y, the result follows. ■

Proof of Proposition 4. We start by showing that the expected revenue from the optimal mechanism is linear in the q's. That is, if at time $t \leq T$ the set Π_t of the object is still available, then the expected revenue is given by

$$R(\Pi_t,t) = \sum_{i=1}^{k_t} q_{(i:\Pi_t)} \left(y_i(t) - \frac{1 - F(y_i(t))}{f(y_i(t))} \right).$$

Our first step toward the above mentioned result consists of showing, by induction on the number of objects, that there exist factors that depend on time but not on the qualities, $a_1(t), ..., a_{k_t}(t)$, such that

$$R(\Pi_t,t) = \sum_{i=1}^{k_t} q_{(i:\Pi_t)} a_i(t). \tag{A17}$$

When only one object of quality q is available, the price for the object is given by $qy_1(t)$ and from Claim 1 it follows that the expected revenue is given by

$$R(q,t) = q\lambda \int_t^T \frac{[1 - F(y_1(s))]^2}{f(y_1(s))} ds.$$

Assume then that if at time t there are k objects with qualities $\Pi_t = (q_1, q_2, ...q_k)$, the expected revenue is given by

$$R(\Pi_t,t) = \sum_{i=1}^{k_t} q_{(i:\Pi_t)} a_i.$$

and consider now a situation with $k+1$ objects at time t. Recall that

$$R(\Pi_t,t)=\sum_{i=1}^{k_t}\int_t^T P_s^i(\Pi_t)\,h_{i,\Pi_t}(s)\,ds+\sum_{i=1}^{k_t}\int_t^T R\left(\Pi_t\backslash\{q_{(i:\Pi_t)}\},s\right)h_{i,\Pi_t}(s)\,ds.$$

The second element of the sum is linear in qualities by the induction argument, while the first element is linear by the definition of $P_s^i(\Pi_t)$ (see equation (1)).

To see that $a_i(t)=y_i(t)-\frac{1-F(y_i(t))}{f(y_i(t))}$, note that, since $a_i(t)$ does not depend on Π_t, equations (A17) and (3) imply that

$$a_i(t)=R(\mathbf{1}_i,t)-R(\mathbf{1}_{i-1},t)=y_i(t)-\frac{1-F(y_i(t))}{f(y_i(t))}$$

The above argument shows that, at time $t=0$, the marginal returns to qualities are ordered, with the $i-$highest produced quality having also the $i-$highest marginal return $y_i(0)-\frac{1-F(y_i(0))}{f(y_i(0))}$. By the convexity of ϕ, the $i-$highest quality will have the $i-$highest marginal cost as well, and the characterization of the optimal produced qualities follows. As soon as the marginal return drops below $\phi'(0)$, the cost of producing another object with a positive quality cannot be recovered, and this determines the optimal number of produced objects. ∎

Proof of Proposition 5. Denote by $y=\{y_i\}_{i=1}^{n_1+n_2}$ the optimal policy given the inventory in period $t=0$. Our previous results imply that

$$\begin{aligned}
P_T^{(l_1+l_2)}(\Pi_T) &= sy_{l_1+l_2}(T);\quad P_0^{(n_1+n_2)}(\Pi_0)=sy_{n_1+n_2}(0)\\
P_T^{(l_1)}(\Pi_T) &= sy_{l_1+l_2}(T)+(q-s)y_{l_1}(T);\quad P_0^{(n_1)}(\Pi_0)=sy_{n_1+n_2}(0)+(q-s)y_{n_1}(0)
\end{aligned}$$

We obtain the following chain:

$$\begin{aligned}
\frac{P_0^{(n_1)}(\Pi_0)-P_T^{(l_1)}(\Pi_T)}{P_0^{(n_1)}(\Pi_0)} &\geq \frac{P_0^{(n_1+n_2)}(\Pi_0)-P_T^{(l_1)}(\Pi_T)}{P_0^{(n_1+n_2)}(\Pi_0)}\Leftrightarrow\\
\frac{P_T^{(l_1)}(\Pi_T)}{P_0^{(n_1)}(\Pi_0)} &\leq \frac{P_T^{(l_1)}(\Pi_T)}{P_0^{(n_1+n_2)}(\Pi_0)}\Leftrightarrow\\
\frac{sy_{l_1+l_2}(T)+(q-s)y_{l_1}(T)}{sy_{n_1+n_2}(0)+(q-s)y_{n_1}(0)} &\leq \frac{sy_{l_1+l_2}(T)}{sy_{n_1+n_2}(0)}\Leftrightarrow\\
\frac{qy_{l_1+l_2}(T)}{sy_{n_1+n_2}(0)+(q-s)y_{n_1}(0)} &\leq \frac{y_{l_1+l_2}(T)}{y_{n_1+n_2}(0)}\Leftrightarrow\\
\frac{q}{sy_{n_1+n_2}(0)+(q-s)y_{n_1}(0)} &\leq \frac{1}{y_{n_1+n_2}(0)}\Leftrightarrow\\
qy_{n_1+n_2}(0) &\leq sy_{n_1+n_2}(0)+(q-s)y_{n_1}(0)\Leftrightarrow\\
(q-s)y_{n_1+n_2}(0) &\leq (q-s)y_{n_1}(0)\Leftrightarrow y_{n_1+n_2}(0)\leq y_{n_1}(0)
\end{aligned}$$

The third line uses Proposition 1 and Remark 2 . The fourth line uses the fact that for all $1 \leq i, j \leq n_1 + n_2$, $y_i(T) = y_j(T)$ which is a consequence of equation 2. The last inequality follows from Proposition 1. ■

Proof of Proposition 6. Assume that the inventory at time t is Π_t, and consider two different qualities $q_{(j_1:\Pi_t)}$ and $q_{(j_2:\Pi_t)}$ such that $q_{(j_1:\Pi_t)} > q_{(j_2:\Pi_t)}$. Recall that the prices associated with each quality satisfy

$$P_t^{(j_l)}(\Pi_t) = \sum_{i=j_l}^{k_t} (q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)}) y_{i,\Pi_t}(t), \; l = 1, 2$$

where y denotes the (quality independent) revenue maximizing allocation policy. Thus,

$$P_t^{(j_1)}(\Pi_t) = P_t^{(j_2)}(\Pi_t) + \sum_{i=j_1}^{j_2-1} (q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)}) y_{i,\Pi_t}(t)$$

which implies

$$\frac{dP_t^{(j_1)}(\Pi_t)}{dt} = \frac{dP_t^{(j_2)}(\Pi_t)}{dt} + \frac{d[\sum_{i=j_1}^{j_2-1} (q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)}) y_{i,\Pi_t}(t)]}{dt}.$$

Lemma 1 (proven below) yields that

$$\frac{dy_{i,\Pi_t}(t)}{dt} \leq 0, \forall i$$

which implies

$$\frac{dP_t^{(j_1)}(\Pi_t)}{dt} \leq \frac{dP_t^{(j_2)}(\Pi_t)}{dt}$$

Since these last derivatives are negative (which is implied again by Lemma 1) the result follows.[19] ■

Lemma 1 *In the revenue maximizing mechanism, the cutoffs determining the allocation policy are decreasing with time. That is*

$$y_i'(t) \leq 0 \textit{ for } i \in \{1, ..., k_t\} \textit{ and } t \in [0, T).$$

Proof. We prove the result by induction on the number of the available objects k_t. If $k_t = 1$, differentiating (A5) w.r.t. t gives

$$y_1'(t) \frac{\partial}{\partial y_1(t)} \left(y_1(t) - \frac{1 - F(y_1(t))}{f(y_1(t))} \right) = -\lambda \frac{[1 - F(y_1(t))]^2}{f(y_1(t))}.$$

[19] The same argument also shows that the difference between the two prices is monotonically decreasing.

The result follows then since the r.h.s. is negative, and since the virtual valuation is assumed to be increasing. Assume then that the statement holds for $k_t = l$, and consider the case where $k_t = l + 1$. The properties of the revenue maximizing mechanism imply that l highest curves $y_i(t)$ coincide with those relevant for the case where $k_t = l$. Thus $y_i'(t) \leq 0$ for $i \in \{1, ..., l\}$. Differentiating w.r.t. t the expression (2) for $y_{l+1}(t)$ gives

$$y_{l+1}'(t)\frac{\partial}{\partial y_{l+1}(t)}\left(y_{l+1}(t) - \frac{1 - F(y_{l+1}(t))}{f(y_{l+1}(t))}\right) = \lambda\left[\frac{[1 - F(y_l(t))]^2}{f(y_l(t))} - \frac{[1 - F(y_{l+1}(t))]^2}{f(y_{l+1}(t))}\right]. \tag{A18}$$

Proposition 1 implies that $y_{l+1}(t) \leq y_l(t)$ for any $t \in [0, T)$. There are two cases:

1. There exists $t \in [0, T)$ such that $y_{l+1}(t) = y_l(t)$. Then, (A18) implies that $y_{l+1}'(t) = 0$.

2. For any $t \in [0, T)$ we have $y_{l+1}(t) < y_l(t)$. Notice that the function $\frac{[1-F(y)]^2}{f(y)}$ is decreasing if and only if the virtual valuation is increasing. Therefore, an increasing virtual valuation implies that the r.h.s. of (A18) is negative and that $\frac{\partial}{\partial y_{l+1}(t)}\left(y_{l+1}(t) - \frac{1-F(y_{l+1}(t))}{f(y_{l+1}(t))}\right) > 0$. This yields $y_{l+1}'(t) < 0$, and completes the proof.

■

References

Albright, Christian S. 1974. "Optimal Sequential Assignments with Random Arrival Times." *Management Science,* 21(1): 60-67.

Arnold, Michael A. and Steven A. Lippman. 2001. "The Analytics of Search with Posted Prices." *Economic Theory,* 17**:** 444-466.

Athey, Susan, and Ilja Segal. 2007. "An Efficient Dynamic Mechanism." discussion paper Stanford University.

Aviv, Yossi, and Amit Pazgal. 2008. "Optimal Pricing of Seasonal Products in the Presence of Forward-Looking Consumers." *Manufacture & Service Operations Management*, 10(3): 339-359.

Bergemann, Dirk, and Juuso Välimäki. 2006. "Efficient Dynamic Auctions." Cowles Foundation Discussion Paper No. 1584.

Bitran, Gabriel R., and Susana V. Mondschein. 1997. "Periodic Pricing of Seasonal Products in Retailing." *Management Science,* 43(1): 64-79.

Bitran, Gabriel R., and René Caldentey. 2003. "An Overview of Pricing Models for Revenue Management." *Manufacture & Services Operations Management,* 5(3): 203-229.

Bulow, Jeremy, and John Roberts. 1989. "The Simple Economics of Optimal Auction." *Journal of Political Economy,* 97(5): 1060-1090.

Derman, Cyrus, Gerald J. Lieberman, and Sheldon M. Ross. 1972. "A Sequential Stochastic Assignment Process." *Management Science,* 18(7): 349-355.

Das Varma, Gopal, and Nikolaos Vettas. 2001. "Optimal Dynamic Pricing with Inventories." *Economic Letters,* 72(3): 335-340.

Elmaghraby, Wedad, and Pinar Keskinocak. 2003. "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions." *Management Science,* 49(10): 1287-1309.

Gallego, Guillermo, and Garrett van Ryzin. 1994. "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons." *Management Science,* 40(8): 999-1020.

Gallego, Guillermo, and Garrett van Ryzin. 1997. "A Multiproduct Dynamic Pricing Problem and its Applications to Network Yield Management." *Operations Research,* 45(1): 24-41.

Gallien, Jérémie. 2006. "Dynamic Mechanism Design for Online Commerce." *Operations Research,* 54(2): 291-310.

Harris, Milton, and Artur Raviv. 1981. "A Theory of Monopoly Pricing Schemes with Deman Uncertainty." *American Economic Review,* 71(3): 347-365.

Kincaid, Wilfred M. and Donald A. Darling. 1963. "An Inventory Pricing Problem." *Journal of Mathematical Analysis and Applications,* 7: 183-208.

Lazear, Edward P. 1986. "Retail Pricing and Clearance Sale." *The American Economic Review,* 76(1): 14-32.

Lippman, Steven, and John J. McCall. 1981. *"The Economics of Uncertainty: Selected Topics and Probabilistic Methods."* in Handbook of Mathematical Economics, Ed. K.J. Arrow and M.D. Intriligator, Volume I, Elsevier, 211-284.

McAfee, Preston, and Vera te Velde. 2007. "*Dynamic Pricing in the Airline Industry.*" in Handbook on Economics and Information Systems, Ed: T.J. Hendershott, Elsevier Handbooks in Information Systems, Volume I.

McAfee, Preston, and Vera te Velde. 2008. "Dynamic Pricing with Constant Demand Elasticity." *Production and Operations Management,* 17(4): 432-438.

McAfee, Preston, and Thomas Wiseman. 2008. "Capacity Choice Counters the Coase Conjecture." *Review of Economic Studies,* 75(1): 317-332.

Mortensen, Dale T. 1986. "*Job Search and Labor Market.*" in Handbook of Labor Economics, Ed: O. Ashenfelter and R. Layard, Volume II, Elsevier, 849-919.

Mussa, Michael, and Sherwin Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory,* 18: 301-317.

Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operation Research,* 6: 58-73.

Nocke, Volker, and Martin Peitz. 2007. "A Theory of Clearance Sales." *Economic Journal,* 117: 964-990.

Parkes, David C., and Satinder Singh. 2003. "An MDP-Based Approach to Online Mechanism Design." *Proceedings of 17th Annual Conference on Neural Information Processing Systems* (NIPS 03).

Pashigian, B.Peter. 1988. "Demand Uncertainty and Sales: A Study of Fashion and Markdown Prices." *The American Economic Review,* 78(5): 936-953.

Pashigian, B.Peter, and Brian Bowen. 1991. "Why Are Products Sold on Sale?: Explanations of Pricing Regularities." *The Quarterly Journal of Economics,* 106(4): 1015-1038.

Pashigian, B.Peter, Brian Bowen, and Eric Gould. 1995. "Fashion, Styling, and the Within-Season Decline in Automobile Prices." *Journal of Law and Economics,* 38(2): 281-309.

Puterman, Martin L. 2005. *Markov Decision Processes.* Wiley: Hoboken, NJ

Riley, John. and Richard Zeckhauser. 1983. "Optimal Selling Strategies: When to Haggle, When to Hold Firm." *The Quarterly Journal of Economics,* 98(2): 267-289.

Ross, Sheldon M. 1983. *Stochastic Processes*, Wiley: New York.

Su, Xuanming. 2007. "Inter-temporal Pricing with Strategic Customer Behavior." *Management Science,* 53(5): 726-741.

Talluri, Kalyan T., and Garrett J. van Ryzin. 2004. *The Theory and Practice of Revenue Management.* Springer: New York.

Wang, Ruqu. 1993. "Auctions versus Posted-Price Selling." *The American Economic Review,* 83(4): 838-851.

Wilson, Charles A. 1988. "On the Optimal Pricing Policy of a Monopolist." *Journal of Political Economy,* 96(1): 164-176.