

A research and education initiative at the MIT
Sloan School of Management


Dynamic Mechanism Design for Online Commerce

Paper 165

Jérémie Gallien

November 2002

For more information,
please visit our website at <http://ebusiness.mit.edu>
or contact the Center directly at ebusiness@mit.edu or 617-253-7054



Dynamic Mechanism Design for Online Commerce¹

J r mie Gallien

Sloan School of Management

Massachusetts Institute of Technology

Cambridge, MA 02142

e-mail: jgallien@mit.edu

November 11, 2002

Abstract

Motivated by electronic commerce, this paper is a mechanism design study for sellers of multiple identical items. In the market environment we consider, participants are risk neutral and time-sensitive, with the same discount factor; potential buyers have unit demand and arrive sequentially according to a renewal process; and valuations are drawn independently from the same regular distribution. From the Revelation Principle, we can restrict our attention to direct dynamic mechanisms taking a sequence of valuations and arrival epochs as a strategic input. We define two properties (discreteness and stability), and prove that under a regularity assumption on the inter-arrival time distribution, we may at no cost of generality consider only mechanisms satisfying them. This effectively reduces the mechanism input to a sequence of valuations, allowing us to formulate the problem as a dynamic program (DP). Because this DP is equivalent to a well-known infinite horizon asset-selling problem, we can finally characterize the optimal mechanism as a sequence of posted prices increasing with each sale. Our numerical study indicates that, with uniform valuations, the benefit of dynamic pricing over a fixed posted price may be small. Besides, posted prices are preferable to online auctions for a large number of items or high interest rate, but in other cases auctions are close to optimal and significantly more robust.

1. Introduction

Recent years have seen a spectacular increase in the volume of commerce transactions realized through Internet websites. However, the trading mechanisms leading to the conclusion of these transactions are quite varied – they include for example electronic catalogues (as on the site Amazon) with either fixed or dynamic posted prices, but also online auctions (as on the site eBay), which may themselves take many forms (see Lucking-Reiley 1999). In this setting, a natural question facing any individual or firm considering the sale of goods on the Internet is: Which mechanism maximizes the revenue to be expected from this sale? Note that the question itself is by no means only relevant to transactions on the world wide web – as discussed below, some satisfying theoretical answers to this problem, although posed

¹ Copyright   2002 J r mie Gallien. Please do not distribute without permission from the author. Feedback on this paper is welcome and appreciated.

in a different market environment, have been provided as early as 1981. Still, electronic commerce situations seem to be characterized by some relatively unique features, namely low transaction costs along with dynamic arrivals and extreme time-sensitivity of market participants. In this paper we build a theoretical model capturing those market features, and analyze it to try and specifically shed some new light on the question stated above in an electronic commerce setting. We begin with a discussion of the relevant literature.

Deriving the acceptance policy maximizing expected discounted revenue when faced with a sequence of purchase offers for a given asset over time has been one of the very early and successful applications of dynamic programming (e.g. Karlin 1962); this situation, often referred to as the "house selling" or "asset selling" problem, is also one of the primary motivations for the development of the related theory of Optimal Stopping (see Chow et al. 1971) and the theory of Search (Kohn and Shavell 1974). In particular, Elfving (1967) and Siegmund (1967) study the optimal selling problem arising for a given continuous time-discounting function when successive offers arrive according to a non-homogeneous Poisson process, a set of assumptions relatively close to ours. An important characteristic of this stream of literature however is that all the game-theoretic aspects of the interaction are ignored, as the process through which potential buyers decide how much to offer is defined exogenously. This approach thus amounts to restricting a priori the range of mechanisms considered to dynamic posted prices, and searching the value of the price parameters characterizing the optimal choice within this family.

In contrast, for a given market environment the problem of mechanism (or "optimal auction") design consists of identifying the mechanism (broadly defined as a pair of allocation and payment functions taking any strategy profile as an input) maximizing the expected revenue predicted by a specified game solution concept (typically, Bayesian Nash equilibrium). The approach for solving such problems has been largely developed by Myerson (1981), who considers a market environment with one item for sale and n risk-neutral bidders with private valuations drawn from a commonly known distribution (possibly under a special correlation structure). Applying the so-called Revelation Principle (see §2.2), he proves that the optimal mechanism is any of the standard types of auctions with a suitably chosen reserve price. His method has since been extended to various other market environments, including notably to multiple identical units and to risk-averse bidders by Maskin and Riley (1989, 1984) – see

Klemperer (1999) for a survey. However, studies in this stream of literature are essentially static: they ignore both the process through which bidders arrive to the market in the first place, and the impact of the timing of transactions on the participants' utility functions.

In that it jointly considers dynamic and game-theoretic features, the present paper is an attempt to bridge the branches of literature described in the previous two paragraphs. A remarkable similar attempt is Riley and Zeckhauser (1983), who derive the optimal selling strategy for a discrete-time model with fixed, additive buyer acquisition costs and time-insensitive bidders, where the seller is free to update his selling strategy immediately before every new bidder arrival. Another important paper relevant to ours is McAfee and McMillan (1988), who also consider a discrete-time model with a fixed bidder acquisition cost and time-insensitive bidders, but a finite buyer population and a pre-determined selling strategy; after extending Myerson's proof of the Revelation Principle to a dynamic setting, they also find the optimal mechanism in this environment. In contrast with those last two papers, we strive to identify the continuous-time mechanism maximizing expected discounted revenue when selling multiple identical items to a stochastic stream of self-interested, time-sensitive bidders. However, because a significant part of our analysis consists of transforming our continuous-time problem into an equivalent discrete-time one, it is unsurprising that the mechanism we obtain eventually is close to theirs. Our paper also shows how Myerson's solution methodology can be adapted to a dynamic programming (DP) framework. Likewise, Vulcano et al. (2002) propose a DP formulation for a dynamic mechanism design problem, but in a discrete-time market environment where a random number of bidders arrive in each period, and there is no interaction between bidders across periods.

The remainder of this paper is structured as follows: The next section §2 includes in §2.1 a description of the market environment we assume and an exact formulation of the corresponding mechanism design problem in §2.2. Next, subsection §2.3 contains the statement of two mechanism properties, along with the result that if an optimal mechanism exists, there must also exist an optimal mechanism satisfying them. This result is key to the solution of our mechanism design problem, which we develop in §2.4 – see Theorem 3 for a summary. Finally, a numerical study is presented in §3, followed with concluding remarks in §4.

2. Model and Analysis

2.1. Market Environment.

Consider a risk-neutral seller with K identical items for sale. Starting at time 0 (when this sale opportunity is advertised), he faces an arrival stream of self-interested potential buyers, each with unit demand. We assume that the buyers' arrival epochs $t_1 \leq t_2 \leq \dots$ follow a renewal process characterized by its transform $\mathcal{G}(z) = E[z^x]$ with $x \sim t_{n+1} - t_n$ for $n \geq 0$ ². Each buyer n is risk-neutral and has a linear utility characterized by a privately known random valuation function $v_n(t)$ (maximum willingness to pay), where t denote the time when the transaction is realized (if applicable). Most of this paper focuses on the case $v_n(t) = \begin{cases} v_n & \text{if } t \geq t_i \\ 0 & \text{otherwise} \end{cases}$ for all potential buyers n , where v_1, v_2, \dots are i.i.d. random variables following a distribution known to all participants and characterized by a density $f(\cdot)$, cdf $F(\cdot)$ and compact support $V = [\underline{v}, \bar{v}]$ with $0 \leq \underline{v} < \bar{v} < +\infty$. We refer to this valuation structure as the *patient* bidders case, sometimes also referred to in the literature as the case "with recall". However, our analysis will also extend to the *impatient* bidders case (or "without recall") $v_n(t) = \begin{cases} v_n & \text{if } t = t_i \\ 0 & \text{otherwise} \end{cases}$. Finally, the seller and the potential buyers have a time-discount factor $\alpha \in (0, 1)$. While this concept will be defined more precisely later along with the participants' utility functions, its loose meaning is that the net value to the seller (resp. bidder n) of any profit y occurring at time t is $\alpha^t y$ (resp. $\alpha^{t-t_n} y$).

In addition to risk neutrality and the stationarity of demand implied by the assumptions of i.i.d valuations and renewal arrival process, four other salient assumptions about the market environment we have either made in this section or will make later in the paper are:

- (A1): $t \mapsto E[\alpha^{x-t} | x > t]$ is a non-decreasing function of t ;
- (A2): $j(v) \equiv v - \frac{1-F(v)}{f(v)}$ is a non-decreasing function on V ;
- (A3): The seller and the potential buyers have the same discount factor $\alpha \in (0, 1)$;
- (A4): Potential buyers have unit demand.

Assumption (A1) is satisfied by all inter-arrival time distributions with increasing failure rate (IFR). It will be used in our analysis to restrict our search to mechanisms where allocation and payment decisions occur upon bidder arrivals, and typically rules out multi-modal distributions, where the expected time until the next arrival could drastically increase when

² $t_0 = 0$ by convention.

conditioning on the time elapsed since the last arrival. (A2) is a very similar assumption that applies to the valuation distribution, and is typical in mechanism design studies; it is also satisfied by all IFR distributions. Intuitively, this assumption rules out mechanisms where potential buyers are discriminated based on which part of the valuation distribution support they are inferred to belong to. Although Myerson (1981) shows how it can be relaxed using convex analysis, we leave this possible extension of our model to future research in order to facilitate our focus on the dynamic aspect of mechanism design. Assumption (A3) is perhaps the strongest in our model (note however that the time to which this discount factor applies is different for each market participant). Interestingly, the buyers' discount factor seems immaterial in the optimal solution we eventually obtain, which could suggest that relaxing (A3) may not change the nature of the optimal mechanism – we do not currently have a proof of this statement though. Finally, while (A4) typically restricts the applicability of this model to situations where potential buyers are end-consumers, Maskin and Riley (1989) show how it can be relaxed in a static environment, so that even if we do not undertake this effort in the present paper, one may be able in the future to apply their approach to our dynamic setting.

2.2. Problem Formulation. In the game-theoretic framework we use in order to capture the self-interested behavior of potential buyers, the question we address in this paper is: Given the market environment just described, which selling mechanism supports a dominant equilibrium achieving the highest expected discounted revenue?

A useful first step is to formally define a mechanism as a pair of allocation and payment functions, both taking as input variable a *strategy profile*, that is the set of all the strategies or signals used individually by the market participants interacting through this game. A major apparent difficulty however is that there is no structural restriction inherent to this definition on what the strategy space of participants should be, so that the realm of mechanisms to search may seem beyond the reach of analysis. As in all optimal mechanism design studies we are aware of, we resolve this difficulty by invoking the so-called Revelation Principle. This principle allows us to restrict our search with no loss of generality to *direct* mechanisms, where the strategy space is nothing but the type space: in our setting, each potential buyer n only provides information to the mechanism through the required input of the type $\varphi_n = (v_n, t_n)$

that entirely characterizes him. More precisely, this principle states that given a mechanism with an arbitrary strategy space where a particular outcome is supported by an equilibrium, one can construct an associated direct mechanism where the same outcome will also be supported by an equilibrium. While we do not discuss here the rigorous game-theoretic justifications of this principle, we refer the reader to Myerson (1981) for a more formal statement of it in a classical mechanism design study, and to McAfee and McMillan (1988) for a justification of applying this principle to a dynamic market environment (see §1 for a detailed discussion of this last paper).

We formally define a *direct dynamic mechanism* ψ to be a sequence $(\psi^n)_{n \geq 1}$ where $\psi^n \equiv (q^n, y^n)$ is a pair of allocation and payment functions, respectively. They are applications $q^n : \Phi^n \mapsto \{0, 1\}^n$ and $y^n : \Phi^n \mapsto (\mathbb{R}^+)^n$ where

$$\Phi^n \equiv \{\varphi^n = ((v_1, t_1), \dots, (v_n, t_n)) \in (V \times \mathbb{R}^+)^n : t_i \leq t_j \text{ for } i \leq j \leq n\}$$

and $\Phi^n \equiv \{(\varphi^n, t), \varphi^n \in \Phi^n, t \in \mathbb{R}^+ : t_n \leq t\}$ ³. We also define the space of all demand streams as $\Phi = \{\varphi = (\varphi_n)_{n \geq 1} : \text{for all } n \geq 1, \varphi^n = (\varphi_k)_{1 \leq k \leq n} \in \Phi^n\}$, and we will refer to the probability measure relative to φ as P .

For $\psi = (q^n, y^n)_{n \geq 1}$ to be an *admissible* mechanism, we first impose that the allocation component functions satisfy the *availability constraint* (AC1) : $\sum_{i=1}^n q_i^n(\varphi^n, t) \leq K$ for all $n \geq 1$, $\varphi^n \in \Phi^n$ and $t \geq t_n$ (the total number of items allocated after the n^{th} bidder has arrived can never exceed K). More saliently, we also impose that the component functions $q_i^n(\cdot)$ and $y_i^n(\cdot)$ for $n \geq i \geq 1$ have *simultaneous step* realizations. That is, for every demand stream $\varphi \in \Phi$ and buyer index $i \geq 1$, either $q_i^n(\varphi^n, t) = y_i^n(\varphi^n, t) = 0$ for all $n \geq i$ and $t \geq t_n$ (bidder i never gets an item), or there exist an integer $s_i(\varphi) \geq i$ and real numbers $\tau_i(\varphi) \geq t_{s_i(\varphi)}$ and $y_i(\varphi) \geq 0$ such that $\tau_i(\varphi) < t_{s_i(\varphi)+1}$, $q_i^n(\varphi^n, t) = \begin{cases} 1 & \text{if } n \geq s_i(\varphi) \text{ and } t \geq \tau_i(\varphi); \\ 0 & \text{otherwise} \end{cases}$ and $y_i^n(\varphi^n, t) = \begin{cases} y_i(\varphi) & \text{if } n \geq s_i(\varphi) \text{ and } t \geq \tau_i(\varphi); \\ 0 & \text{otherwise} \end{cases}$ (bidder i gets an item at time $\tau_i(\varphi)$ and simultaneously pays $y_i(\varphi)$ for it⁴). We define by extension the r.v. $y_i(\varphi)$, $s_i(\varphi)$ and $\tau_i(\varphi)$ to be respectively equal to 0, $+\infty$ and $+\infty$ in the former case, and we define the r.v. $q_i(\varphi)$ to be equal to 0 in the former case and 1 in the latter. Also, we will omit from now on the dependence of all r.v. on φ and ψ when it is clear from context.

³ Note that random allocation rules are not investigated here.

⁴ $y_i^n(\cdot)$ thus corresponds to a cumulative transfer payment function.

While we will use the component functions of (q^n, y^n) in our study of the optimal mechanism properties of section 2.3 below, we will only need in subsequent sections the reduced sets of variables $(q_n, y_n, \tau_n)_{n \geq 1}$ or $(q_n, y_n, s_n)_{n \geq 1}$. Define the σ -algebras $\mathcal{F} = \sigma(\varphi)$, $\mathcal{F}_{n,t} = \sigma(\{\varphi^m$ and $t_{m+1} > t : m \geq n$ and $t_m \leq t\})$ and $\mathcal{F}_{\tau_n} = \{A \in \mathcal{F} : A \cap \{\tau_n \leq t\} \in \mathcal{F}_{n,t} \text{ for all } t \geq 0\}$. In terms of $(q_n, y_n, \tau_n)_{n \geq 1}$, the condition that the mechanism $(\psi^n)_{n \geq 1}$ is admissible (more feasibility conditions will be imposed shortly) can be summarized with $(AC) : \sum_{i=1}^n q_i \leq K$ for all $n \geq 1$ and $\varphi \in \Phi$; and

$$\begin{array}{llll}
\{\tau_n \leq t\} \in \mathcal{F}_{n,t} & \text{for all } t \geq 0 \text{ and } n \geq 1 & (cSS0) & \\
(q_n, y_n) \text{ is } \mathcal{F}_{\tau_n}\text{-measurable} & \text{for all } n \geq 1 & (cSS1) & \\
q_n = 1 \Leftrightarrow \tau_n < +\infty & \text{for all } n \geq 1 \text{ and } \varphi \in \Phi; & (cSS2) & (cSS) \\
y_n > 0 \Rightarrow \tau_n < +\infty & \text{for all } n \geq 1 \text{ and } \varphi \in \Phi; & (cSS3) & \\
q_n \in \{0, 1\}, y_n \geq 0, \tau_n \geq t_n & \text{for all } n \geq 1 \text{ and } \varphi \in \Phi. & (cSS4) &
\end{array}$$

In words, we restrict our attention to the class of mechanisms where full transfer payments from buyers occur instantaneously upon the allocation of an item to them. Note that delayed and continuous payment schemes in particular are thereby excluded.

We define the expected utility function of the seller as $U_o(\psi) \equiv E_\varphi [\sum_{n=1}^{+\infty} \alpha^{\tau_n} y_n]$, and the (random) utility function of bidder $n \geq 1$ as $u_n(v_n) \equiv \alpha^{\tau_n - t_n} (v_n q_n - y_n)$. Note that the definition of the n -th bidder's utility function u_n is relative to the exogenous time t_n at which he arrives. Introducing the notations $\varphi[v', n] \equiv \{\varphi^{n-1}, (v', t_n), \varphi_{n+1}, \dots\}$ and $\xi_n[v'] \equiv \xi_n(\varphi[v', n])$ for any $v' \in V$ and random variable ξ_n defined on Φ and relative to index $n \geq 1$, we can generalize this last definition to $u_n(v', v_n) \equiv \alpha^{\tau_n[v'] - t_n} (v_n q_n[v'] - y_n[v'])$. This expresses the utility of the n -th bidder with true valuation v_n when reporting a false valuation v' into the mechanism (but a correct arrival epoch t_n). We are now ready to formulate our optimization problem in the patient bidders case as:

$$\begin{array}{ll}
\text{Max}_{\psi} U_o(\psi) & \\
\text{Subject to } (AC), (cSS) \text{ and} & \\
u_n(v_n) \geq 0 & \text{for all } \varphi \in \Phi \text{ and } n \geq 1; \quad (IR) \\
u_n(v_n) \geq u_n(v', v_n) & \text{for all } \varphi \in \Phi \text{ and } n \geq 1, v' \in V. \quad (IC)
\end{array} \tag{1}$$

The last two constraints in (1) capture the self-interested behavior of participants and reflect the solution concept used to predict the outcome of the game: (IR) ensures *individual rationality*, namely that a potential buyer will only participate if his utility from doing so is non-negative, and (IC) guarantees *incentive compatibility*, namely that a bidder cannot

benefit from misrepresenting his valuation when interacting with the mechanism. In this form (*IC*) does not guarantee that a bidder could not benefit from misrepresenting his arrival epoch (which constitutes a part of his type), i.e. waiting for some time after his arrival for strategic purposes before communicating his valuation to the mechanism – in §2.4 when stating our main theoretical result Theorem 3, we will refer to such bidders as being *time-strategic*. Define $u_n(v', t, v_n) \equiv u_{m[t]}(\varphi[v', n, t])$, where $m[t] = \inf\{m : t_m > t\}$ and

$$\varphi[v', n, t] \equiv \begin{cases} \{\varphi^{n-1}, \varphi_{n+1}, \dots, \varphi_{m[t]-1}, (v', t), \varphi_{m[t]}, \dots\} & \text{if } m[t] > n + 1 \\ \{\varphi^{n-1}, (v', t), \varphi_{n+1}, \dots\} & \text{if } m[t] = n + 1 \end{cases}$$

for all $\varphi \in \Phi$, $n \geq 1$, $v' \in V$ and $t \geq 0$ such that $t \geq t_n$: an appropriate way to enforce full incentive compatibility would be to use instead the constraint (*IC'*) : $u_n(v_n) \geq u_n(v', v_n, t)$. However, it will turn out that the mechanisms obtained by solving (1) under assumptions (A1) – (A4) (see section 2.1) will also satisfy the more stringent condition (*IC'*).

An important remark is that (*IR*) and (*IC*) must hold for every realization of the demand stream φ , because we require the outcome prediction associated with all mechanisms considered to be supported by a dominant equilibrium. Note that under the alternative, less demanding requirement that the seller's revenue be only supported by a Bayesian Nash equilibrium, these constraints would only have to hold in expectation, conditional on the bidders' information set. There are two reasons why we use dominant equilibrium rather than Bayesian equilibrium as our solution concept in this paper: A first generic reason is that when they exist, dominant equilibria predict much more reliably the outcomes of games with incomplete information than Bayesian equilibria. A second more technical and idiosyncratic reason is that it turns out to be considerably easier to prove the feasibility of various mechanisms we construct during our analysis when the constraints of (1) are required to hold for every sample path rather than in expectation.

Finally, we observe that a formulation of the mechanism design problem corresponding to the impatient bidders case is obtained by adding to (1) the constraint (*IB*) : $q_n = 1 \Rightarrow \tau_n = t_n$ for all n , that is by merely considering a subset of the feasible solution space for (1).

2.3. Optimal Mechanism Properties. In this section we introduce two properties of dynamic mechanisms holding under assumption (A1), and show that it is costless to restrict our search to mechanisms satisfying them; This later enables us to significantly simplify the

formulation of our problem.

We define the r.v. s'_k and τ'_k for $1 \leq k \leq K$ as $(s'_k, \tau'_k) \equiv \inf \{(n, t) : \sum_{i=1}^n q_i^n(\varphi^n, t) \geq k\}$ (number of arrived bidders and time epoch at the k^{th} item sale, respectively), and let n_k be the buyer to whom the k^{th} item is sold, so that $(s_{n_k}, \tau_{n_k}) = (s'_k, \tau'_k)$ and $q_{n_k}^n(\varphi^n, t) = 1 \Leftrightarrow \begin{cases} n \geq s'_k \\ t \geq \tau'_k \end{cases}$. We will say that a mechanism ψ is *discrete* if $\tau'_k = t_{s'_k}$ for all $1 \leq k \leq K$, that is if items are only sold upon a customer arrival (although not necessarily to the arriving customer). Because all mechanisms of interest for impatient bidders are discrete, the following first proposition is non-trivial only with patient bidders:

Proposition 1 *If $E[\alpha^{x-t}|x > t]$ is a nondecreasing function of t , then any optimal mechanism is discrete.*

Proof. Let $\psi = (q^n, y^n)_{n \geq 1}$ be an optimal mechanism, and take $0 \leq k \leq K - 1$. Consider now any $n \geq 1$ and φ^n for which $k(t_n) \equiv \lim_{t \rightarrow t_n^-} \sum_{i=1}^{n-1} q_i^{n-1}(\varphi^{n-1}, t) = k$ (the number of items sold before time t_n is equal to k), and define the set $I(t_n) = \{n_j : 1 \leq j \leq k(t_n)\}$ (bidders to whom these k items have been sold). From the principle of dynamic optimality, $(q^m, y^m)_{m \geq n}$ maximizes $E_{\bar{\varphi}^n} \left[\sum_{j=k(t_n)+1}^K \alpha^{\tau'_j} y_{n_j} | (\varphi_i)_{\{i:1 \leq i \leq n\} \setminus I(t_n)} \right]$ among all sequences $(\hat{q}^m, \hat{y}^m)_{m \geq n}$ such that $(q^m, y^m)_{1 \leq m < n} \times (\hat{q}^m, \hat{y}^m)_{m \geq n}$ is feasible. Denote now by $J(\varphi^{n-1}, \varphi_n, I(t_n))$ the optimal value of this expected partial objective value. Because φ is a renewal process, for all $(t', t) \in (0, +\infty)^2$ the distribution of $\bar{\varphi}^n$ given $t_n = t'$ is equal to the distribution of $\bar{\varphi}^n$ given $t_n = t$ translated by $t' - t$. Therefore $J(\varphi^{n-1}, (v_n, t'), I) = \alpha^{t'-t} J(\varphi^{n-1}, (v_n, t), I)$ for all $(t', t) \in (t_{n-1}, +\infty)^2$ and any set $I \subset \{i : 1 \leq i \leq n-1\}$ such that $|I| \leq K$, so there exists a function $V(\varphi^{n-1}, v_n, I)$ independent of t such that $J(\varphi^{n-1}, (v_n, t), I) = \alpha^t V(\varphi^{n-1}, v_n, I)$. Applying now the principle of dynamic optimality to the case $k = K - 1$, it is optimal to sell the last item for a payment of $y_{n_K} = y$ at time $\tau'_{n_K} = t$ between the n^{th} and $(n+1)^{\text{th}}$ bidder arrivals and after the $K - 1$ first items have been sold to the bidders in set I_{K-1} if and only if this maximizes discounted revenue among all feasible sales alternatives available at this point and

$$\begin{aligned} \alpha^t y &\geq E_{\varphi_{n+1}} [J(\varphi^n, \varphi_{n+1}, I_{K-1}) | t_{n+1} > t] \\ &= E_{(v_{n+1}, t_{n+1})} [\alpha^{t_{n+1}} V(\varphi^n, v_{n+1}, I_{K-1}) | t_{n+1} > t] \quad (\text{from the result just proven}) \\ &= E_{t_{n+1}} [\alpha^{t_{n+1}} | t_{n+1} > t] E_{v_{n+1}} [V(\varphi^n, v_{n+1}, I_{K-1})] \quad (\text{from the independence of } v_{n+1} \text{ and } t_{n+1}) \end{aligned}$$

Because the resulting inequality holds for $t \geq \max(t_n, \tau'_{K-1})$, invoking now the hypothesis that $E_{t_{n+1}}[\alpha^{t_{n+1}-t} | t_{n+1} > t]$ is nondecreasing in t and the "if" part of the equivalence stated above, we have proven so far that either $\tau'_K = t_{s'_K}$ or $\tau'_K = \tau'_{K-1}$. But note that in the latter case we can apply the same reasoning to $y = y_{n_K} + y_{n_{K-1}}$, $t = \tau'_K = \tau'_{K-1}$ and I_{K-2} (and so on), it follows therefore that there exists k_1 , $1 \leq k_1 \leq K$ such that $\tau'_k = t_{s'_k}$ if and only if $k_1 \leq k \leq K$. Assuming $k_1 > 1$ (otherwise the result is proven), applying again the same reasoning to $y_{n_{k_1-1}} = y$, $\tau'_{k_1-1} = t$ and I_{k_1-2} shows that there exists k_2 , $1 \leq k_2 < k_1$ such that $\tau'_k = t_{s'_{k_1-1}}$ if and only if $k_2 \leq k < k_1$. Because the sequence k_1, k_2, \dots, k_ℓ is strictly decreasing until the point where $k_\ell = 1$, repeating this procedure iteratively shows that $\tau'_k = t_{s'_k}$ for all $k : 1 \leq k \leq K$, or equivalently that ψ is discrete. ■

We will assume in the rest of the paper that the hypothesis (A1) required by Proposition 1 holds; it is satisfied in particular by all IFR distributions. This justifies that we restrict ourselves from now on to discrete mechanisms, and use instead the simplified definitions $q^n : \Phi^n \mapsto \{0, 1\}^n$ and $y^n : \Phi^n \mapsto (\mathbb{R}^+)^n$, where

$$\Phi^n \equiv \{\varphi^n = ((v_1, t_1), \dots, (v_n, t_n)) \in (V \times \mathbb{R}^+)^n : t_i \leq t_j \text{ for } i \leq j \leq n\}.$$

In this setting, the two variables τ_n and s_n become redundant, since $\tau_n = t_{s_n}$ for all φ . The set of conditions (cSS0) – (cSS4) is then equivalent to its discrete-time analogue (dSS0) – (dSS4), with (dSS0) : $\{s_n \leq m\} \in \mathcal{F}_m$ for all $m \geq n \geq 1$, where $\mathcal{F}_m = \sigma(\varphi^m)$ and (dSS1) – (dSS4) is obtained by replacing τ_n with s_n in (cSS1) – (cSS4) and t_n with n in (cSS4), where $\mathcal{F}_{s_n} = \{A \in \mathcal{F} : A \cap \{s_n \leq m\} \in \mathcal{F}_m \text{ for all } m \geq 1\}$.

We introduce next the following definition: A discrete mechanism $(\psi^n)_{n \geq 1}$ is said to be *stable* if for all $n \geq 1$, $v^n = \hat{v}^n$ implies⁵ $\psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n)$. In words, a mechanism is stable if the allocations and payments it generates are independent of the bidder arrival epochs. By extension, we will say that a mechanism is *N-stable* if its allocations and payments are only independent of the arrival epochs of the N first bidders for $N \geq 1$, that is for all $n \geq 1$ and $(\varphi^n, \hat{\varphi}^n) \in \Phi^n$ such that $v^n = \hat{v}^n$,

$$\begin{cases} \psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n) & \text{if } n \leq N \\ \psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n) & \text{if } n > N \text{ and } t_\ell = \hat{t}_\ell \text{ for } N < \ell \leq n \end{cases} ;$$

note that the notions of stability and N -stability for all $N \geq 1$ coincide. Our interest in

⁵ In this definition, $\varphi = \{(v_1, t_1), (v_2, t_2), \dots\}$, $\hat{\varphi} = \{(\hat{v}_1, \hat{t}_1), (\hat{v}_2, \hat{t}_2), \dots\}$ and $v^n = (v_i)_{1 \leq i \leq n}$.

these properties is justified by the main theorem of this section, stated immediately after the following technical lemma.

Lemma 1 *Let $(\psi^n)_{n \geq 1}$ be a N -stable mechanism feasible for (1), and let $(g^n)_{n \geq 1}$ be a family of applications onto $(\Phi^n)_{n \geq 1}$ such that:*

(i) $(g^n)_{n \geq 1}$ is valuation-preserving, that is $\Delta^n(g^n(\varphi^n)) = \Delta^n(\varphi^n)$ for all $n \geq 1$ and $\varphi^n \in \Phi^n$, where $\Delta^n(\varphi^n) \equiv v^n$;

(ii) $(g^n)_{n \geq 1}$ is N -consistent, that is $\Lambda^n(g^m(\varphi^m)) = \Lambda^n(g^n(\varphi^n))$ for all $m \geq n \geq N + 1$, where $\Lambda^n(\varphi^m) \equiv t^n$;

(iii) $(g^n)_{n \geq 1}$ is ψ -commutative, that is $\psi^m(g^m(\varphi^m)[v'_n, n]) = \psi^m(g^m(\varphi^m[v'_n, n]))$ for each $m \geq n \geq 1$, $\varphi^m \in \Phi^n$ and $v'_n \in V$.

The mechanism $(\dot{\psi}^n)_{n \geq 1} \equiv (\psi^n \circ g^n)_{n \geq 1}$ is then also N -stable and feasible for (1).

Proof. It is immediate to check that $\dot{\psi}$ is N -stable and that it satisfies (dSS) and (AC) when (i) and (ii) hold. Denoting by $g(\varphi)$ the event $(g^n(\varphi^n))_{n \geq 1}$, observe that

$$(q_n(g(\varphi)), y_n(g(\varphi)), s_n(g(\varphi))) = (\dot{q}_n(\varphi), \dot{y}_n(\varphi), \dot{s}_n(\varphi)).$$

Because ψ satisfies in particular (IR) on the events $g(\varphi)$ for $\varphi \in \Phi$, $\dot{\psi}$ satisfies (IR). For all $n \geq 1$, $\varphi \in \Phi$ and $v'_n \in V$ we have

$$\begin{aligned} \alpha^{t_{s_n(g(\varphi))}}(v_n q_n(g(\varphi)) - y_n(g(\varphi))) &\geq \alpha^{t_{s_n(g(\varphi)[v'_n, n])}}(v_n q_n(g(\varphi)[v'_n, n]) - y_n(g(\varphi)[v'_n, n])) \\ \alpha^{t_{s_n(\varphi)}}(v_n \dot{q}_n(\varphi) - \dot{y}_n(\varphi)) &\geq \alpha^{t_{s_n(g(\varphi)[v'_n, n])}}(v_n q_n(g(\varphi)[v'_n, n]) - y_n(g(\varphi)[v'_n, n])) \\ &= \alpha^{t_{s_n(\varphi[v'_n, n])}}(v_n \dot{q}_n(\varphi[v'_n, n]) - \dot{y}_n(\varphi[v'_n, n])) \end{aligned}$$

where the first inequality holds because ψ satisfies in particular (IC) on the events $g(\varphi)$ for $\varphi \in \Phi$, and the second inequality follows from assumption (iii). This is just the statement that $\dot{\psi}$ satisfies (IC), which concludes the proof. ■

Theorem 2 *If there exists a discrete optimal mechanism, there also exists an optimal mechanism which is stable.*

Proof. Let $\psi = (\psi^n)_{n \geq 1}$ be a discrete solution to (1). In the following, we prove by induction that for all $N \geq 1$ there exists a mechanism $\tilde{\psi}$ which is optimal and N -stable.

Let $N \in \{0, 1, \dots\}$ and assume that ψ is optimal and N -stable (by extension, we define all feasible mechanisms to be 0-stable – the base of the induction is thus included in what follows). Let $k(\varphi^n) \equiv \sum_{i=1}^n q_i^n(\varphi^n)$ and $\varphi^N \in \Phi^N$ such that $k(\varphi^N) < K$ – if such a φ^N does not exist, then ψ is $(N + 1)$ -stable. From the principle of dynamic optimality, $(\psi^j)_{j>N}$ maximizes $E \left[\sum_{\ell=k(\varphi^N)+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \varphi^{N+1} \right]$ among all sequences $(\dot{\psi}^j)_{j>N}$ such that $(\psi^j)_{1 \leq j \leq N} \times (\dot{\psi}^j)_{j>N}$ is feasible for (1). We now prove that for any $(\hat{\varphi}^{N+1}, \tilde{\varphi}^{N+1}) \in \Phi^{N+1}$ such that $\hat{v}^{N+1} = \tilde{v}^{N+1}$ and $k(\hat{\varphi}^N) = k(\tilde{\varphi}^N) = k < K$,

$$\alpha^{-\hat{t}_{N+1}} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \hat{\varphi}^{N+1} \right] = \alpha^{-\tilde{t}_{N+1}} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \tilde{\varphi}^{N+1} \right]. \quad (2)$$

By contradiction, assume instead

$$\alpha^{-\hat{t}_{N+1}} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \hat{\varphi}^{N+1} \right] < \alpha^{-\tilde{t}_{N+1}} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \tilde{\varphi}^{N+1} \right].$$

Consider the family of applications $(g^n)_{n \geq 1}$ defined by $g^n(\varphi^n) = (v^n, \pi^n(t^n))$ and

$$\begin{cases} \pi^n(t^n) = t^n \text{ for } n \leq N; \\ \pi^n(t^n) = (\tilde{t}^{N+1}, t_{N+2} - t_{N+1} + \tilde{t}_{N+1}, \dots, t_n - t_{N+1} + \tilde{t}_{N+1}) \text{ for } n \geq N + 1. \end{cases}$$

From Lemma 1 the mechanism $\dot{\psi} = \psi \circ g$ is feasible, yet

$$\begin{aligned} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} \dot{y}_{n_\ell} \mid \hat{\varphi}^{N+1} \right] &= \alpha^{\hat{t}_{N+1} - \tilde{t}_{N+1}} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \tilde{\varphi}^{N+1} \right] \\ &> E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \hat{\varphi}^{N+1} \right], \end{aligned}$$

where the equality above follows from the definition of $\dot{\psi}$ and the fact that $(\varphi^n)_{n \geq 1}$ is a renewal process, and the inequality from the contradiction hypothesis. But this inequality is strict and $(\dot{\psi}^j)_{j \leq N} = (\psi^j)_{j \leq N}$, which contradicts the optimality of $(\psi^j)_{j>N}$ for $\text{Max}_{(\dot{\psi}^j)_{j>N}} E \left[\sum_{\ell=k(\hat{\varphi}^N)+1}^K \alpha^{t_{s'_\ell}} y_{n_\ell} \mid \hat{\varphi}^{N+1} \right]$ stated above, and proves (2).

Define now the family of applications $(g^n)_{n \geq 1}$ on $(\Phi^n)_{n \geq 1}$ by $g^n(\varphi^n) = (v^n, \pi^n(t^n))$ and

$$\begin{cases} \pi^n(t^n) = t^n \text{ for } n \leq N; \\ \pi^{N+1}(t^{N+1}) = (t^N, \hat{t}_{N+1}(t_N)) \text{ where } \hat{t}_{N+1}(t_N) = t_N + E[x]; \\ \pi^n(t^n) = (t^N, \hat{t}_{N+1}(t_N), t_{N+2} - t_{N+1} + \hat{t}_{N+1}(t_N), \dots, t_n - t_{N+1} + \hat{t}_{N+1}(t_N)) \text{ for } n \geq N + 2. \end{cases}$$

Note that the mechanism $\tilde{\psi}$ defined by $\tilde{\psi}^n = \psi^n \circ g^n$ is $(N + 1)$ -stable, and it is easy to

check that it satisfies the conditions of Lemma 1, so that it is feasible for (1). Besides, for any $\varphi^{N+1} \in \Phi^{N+1}$ such that $k(\varphi^N) = k < K$ we have $\tilde{k}(\varphi^N) = k$ because ψ is N -stable from the induction hypothesis and

$$\begin{aligned}
E \left[\sum_{\ell=1}^K \alpha^{t_{s'\ell}} \tilde{y}_{\tilde{n}_\ell} \middle| \varphi^{N+1} \right] &= E \left[\sum_{\ell=1}^k \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] + E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'\ell}} \tilde{y}_{\tilde{n}_\ell} \middle| \varphi^{N+1} \right] \\
&= E \left[\sum_{\ell=1}^k \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] + \alpha^{t_{N+1} - \hat{t}_{N+1}(t_N)} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'\ell}} \tilde{y}_{\tilde{n}_\ell} \middle| g^{N+1}(\varphi^{N+1}) \right] \\
&= E \left[\sum_{\ell=1}^k \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] + \alpha^{t_{N+1} - \hat{t}_{N+1}(t_N)} E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'\ell}} y_{n_\ell} \middle| g^{N+1}(\varphi^{N+1}) \right] \\
&= E \left[\sum_{\ell=1}^k \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] + E \left[\sum_{\ell=k+1}^K \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] \\
&= E \left[\sum_{\ell=1}^K \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right], \tag{3}
\end{aligned}$$

where the first equality follows from the fact that $\alpha^{t_{s'\ell}} y_{n_\ell}$ is \mathcal{F}_N -measurable for $\ell \leq k$ and $\tilde{\psi}^n = \psi^n$ for $n \leq N$; the second because $g^n(g^n(\varphi^n)) = g^n(\varphi^n)$ and $(\varphi^n)_{n \geq 1}$ is a renewal process; the third from the construction of $\tilde{\psi}$; and the fourth from (2). Equation (3) and the fact that $\tilde{\psi}$ is N -stable imply finally

$$\begin{aligned}
U_0(\tilde{\psi}) &= E \left[E \left[\sum_{\ell=1}^K \alpha^{t_{s'\ell}} \tilde{y}_{\tilde{n}_\ell} \middle| \varphi^{N+1} \right] \right] \\
&= E \left[E \left[\sum_{\ell=1}^K \alpha^{t_{s'\ell}} y_{n_\ell} \middle| \varphi^{N+1} \right] \right] \\
&= U_0(\psi).
\end{aligned}$$

This last equality proves that $\tilde{\psi}$ is optimal, which concludes the proof. ■

Because we assume (A1) to hold throughout this paper, the combination of Proposition 1 and Theorem 2 justifies that we restrict our attention to discrete and stable mechanisms, and work from now on with the yet simpler definitions $q^n : V^n \mapsto \{0, 1\}^n$ and $y^n : V^n \mapsto (\mathbb{R}^+)^n$.

2.4. Optimization Problem Solution. The first part of this section is a proof culminating with the statement of Theorem 3, the main theoretical result of this paper. Its first

step is an alternative expression for the expected utility function of the seller:

$$\begin{aligned}
E[\alpha^{\tau_n} y_n] &= E[\alpha^{t_{s_n}} y_n] \text{ because the mechanism is discrete;} \\
&= E[E[\alpha^{t_{s_n}} y_n | s_n]] \text{ by the law of total probability;} \\
&= E[E[\alpha^{t_{s_n}} | s_n] E[y_n | s_n]] \text{ because the mechanism is stable;} \\
&= E[\mathcal{G}(\alpha)^{s_n} E[y_n | s_n]] \text{ by the definition of } \mathcal{G}(\alpha); \\
&= E[\mathcal{G}(\alpha)^{s_n} y_n] \text{ by the law of total probability.}
\end{aligned}$$

Applying the monotone convergence theorem to this last equation yields

$$U_o(\psi) = \sum_{n=1}^{+\infty} E[\mathcal{G}(\alpha)^{s_n} y_n]. \quad (4)$$

We next apply to our problem the methodology developed in section 4 of Myerson (1981), so we only provide an outline of this part of the derivation and refer the interested reader to that paper for the justifications we omitted. First, the feasibility conditions (*AC*), (*dSS*), (*IR*) and (*IC*) for (1) are equivalent to the set of conditions composed of (*AC*), (*dSS*) and

$$\begin{aligned}
u_n(\underline{v}) &\geq 0 && \text{for all } n \geq 1 \text{ and } \varphi \in \Phi; && (IR1) \\
\alpha^{t_{s_n[v']}-t_n} q_n[v'] &\leq \alpha^{t_{s_n[v]}-t_n} q_n[v] && \text{for all } n \geq 1, \varphi \in \Phi \text{ and } v \geq v' \in V^2; && (IP1) \\
u_n(v_n) &= u_n(\underline{v}) + \int_{\underline{v}}^{v_n} \alpha^{t_{s_n[v]}-t_n} q_n[v] dv && \text{for all } n \geq 1, \varphi \in \Phi \text{ and } v' \in V. && (IC1)
\end{aligned}$$

Second, the objective function $U_o(\psi)$ can be further transformed. Specifically, (*IC1*) implies that

$$\begin{aligned}
E[\mathcal{G}(\alpha)^{s_n} y_n] &= E[\mathcal{G}(\alpha)^{s_n} v_n q_n] - E\left[\int_{\underline{v}}^{v_n} E[\mathcal{G}(\alpha)^{s_n} q_n | v_n = v] dv\right] - \mathcal{G}(\alpha)^n E[u_n(\underline{v})]; \\
&= E\left[\left(v_n - \frac{1 - F(v_n)}{f(v_n)}\right) \mathcal{G}(\alpha)^{s_n} q_n\right] - \mathcal{G}(\alpha)^n E[u_n(\underline{v})],
\end{aligned}$$

where the second equality is obtained after applying Fubini's theorem and integrating by part the second term of the right-hand side of the first equality and re-arranging terms. Combining this last expression with (4) and applying the monotone convergence theorem yields the alternative equivalent formulation:

$$U_o(\psi) = E\left[\sum_{n=1}^{+\infty} \left(v_n - \frac{1 - F(v_n)}{f(v_n)}\right) \mathcal{G}(\alpha)^{s_n} q_n\right] - \sum_{n=1}^{+\infty} \mathcal{G}(\alpha)^n E[u_n(\underline{v})]. \quad (5)$$

Note that only the variables $(q_n, s_n)_{n \geq 1}$ relative to the allocation decisions appear in the

first term of the right-hand side of (5). We can thus exploit the following consequence of this last observation combined with (IR1): if $(q_n, s_n)_{n \geq 1}$ solves

$$\begin{aligned} & \text{Max}_{(q_n, s_n)_{n \geq 1}} E \left[\sum_{n=1}^{+\infty} \left(v_n - \frac{1-F(v_n)}{f(v_n)} \right) \mathcal{G}(\alpha)^{s_n} q_n \right] \\ & \text{subject to } (AC), (IP1), (dSS0), (dSS2) \text{ and:} \\ & (q_n)_{n \geq 1} \text{ is adapted to } (\mathcal{F}_{s_n})_{n \geq 1}; \quad (dSS1') \\ & q_n \in \{0, 1\} \text{ and } s_n \geq n, \quad (dSS4') \end{aligned} \quad (6)$$

and $(y_n)_{n \geq 1}$ can be found such that (IC1), (dSS1), (dSS3) and (dSS4) are satisfied and $E[u_n(\underline{v})] = 0$ for all n , then $(q_n, y_n, s_n)_{n \geq 1}$ is an optimal solution to the original problem (1).

When a single item is for sale ($K = 1$), the relaxation of (6) where constraint (IP1) has been removed is equivalent to the infinite horizon discounted asset selling problem with recall, where the discrete discount factor is $\mathcal{G}(\alpha)$ and the distribution for the sequential purchasing offers is $w \sim v - \frac{1-F(v)}{f(v)}$ (we denote the corresponding probability law by H). The solution to this classical problem is to immediately accept any offer w_n such that $w_n \geq \bar{w}$, where \bar{w} satisfies $\frac{\bar{w}}{\mathcal{G}(\alpha)} = \bar{w}P(w \leq \bar{w}) + \int_{\bar{w}}^{\infty} w dH(w)$, and reject all other offers (see Chapters 4 and 7 in Bertsekas 1995). In the common case where $j(v) \equiv v - \frac{1-F(v)}{f(v)}$ is nondecreasing (assumption (A2) discussed in section 2.1), this is equivalent to immediately selling the item to the first bidder n such that $v_n \geq p_1$, where p_1 satisfies

$$j(p_1) = p_1 \frac{\mathcal{G}(\alpha)(1 - F(p_1))}{1 - \mathcal{G}(\alpha)F(p_1)}. \quad (7)$$

When multiple items are for sale ($K > 1$), the optimal policy is found by introducing a state variable $k \in \{1, \dots, K\}$ representing the number of items remaining to be sold, and applying the same result through backwards induction on k by adding to each item's selling price in the Bellman equation for (6) the optimal expected discounted revenue associated with the sale of all the other items still remaining at that point. Consequently, there exists a finite sequence $(p_k)_{1 \leq k \leq K}$ such that it is optimal to sell the $(K - k + 1)$ -th item to the first bidder n arriving after the $(K - k)$ -th sale such that $v_n \geq p_k$. Equivalently, if we define by recurrence an increasing finite sequence $(n_\ell)_{1 \leq \ell \leq K}$ by $n_1 = \min\{n : v_n \geq p_K\}$ and $n_\ell = \min\{n > n_{\ell-1} : v_n \geq p_{K-\ell+1}\}$ for $\ell \geq 2$ (arrival index of first and ℓ -th buyers, respectively) and $\ell(n) = \begin{cases} \max\{\ell : n_\ell < n\} & \text{if } n_1 < n \\ 0 & \text{otherwise} \end{cases}$ (number of items sold by the n -th arrival), then the optimal allocation policy can be written $\hat{q}_n = 1\{v_n \geq p_{K-\ell(n)}\}$ and $\hat{s}_n = n/\hat{q}_n$. Note that $\mathcal{G}(\alpha)^{\hat{s}_n - n} \hat{q}_n = 1\{v_n \geq p_{K-\ell(n)}\}$ and $p_{K-\ell(n)}$ is independent of v_n , so

$(\hat{q}_n, \hat{s}_n)_{n \geq 1}$ satisfies constraint (IP1) and is therefore an optimal solution to (6).

Turning now to the payment functions $(y_n)_{n \geq 1}$, observe that according to (IC1) the term $\mathcal{G}(\alpha)^n E[u_n(\underline{v})]$ appearing in the objective (5) can be expressed for all $v' \in V$ as

$$\mathcal{G}(\alpha)^n E[u_n(\underline{v})] = E[\mathcal{G}(\alpha)^{s_n} (v_n q_n - y_n) | v_n = v'] - \int_{\underline{v}}^{v'} E[\mathcal{G}(\alpha)^{s_n} q_n | v_n = v] dv$$

Therefore, if we set $\hat{y}_n = p_{K-\ell(n)} 1\{v_n \geq p_{K-\ell(n)}\}$, we have $E[u_n(\underline{v})] = 0$ for all n as:

$$\begin{aligned} \int_{\underline{v}}^{v'} E[\mathcal{G}(\alpha)^{\hat{s}_n} \hat{q}_n | v_n = v] dv &= \mathcal{G}(\alpha)^n \int_{\underline{v}}^{v'} 1\{v \geq p_{K-\ell(n)}\} dv \\ &= \mathcal{G}(\alpha)^n (v' - p_{K-\ell(n)}) 1\{v' \geq p_{K-\ell(n)}\} \\ &= E[\mathcal{G}(\alpha)^{\hat{s}_n} (v_n \hat{q}_n - \hat{y}_n) | v_n = v']. \end{aligned}$$

In addition, it is straightforward to verify that \hat{y}_n satisfies constraints (IC1), (dSS1), (dSS3) and (dSS4); we have therefore proven that the mechanism $(\hat{q}_n, \hat{y}_n, \hat{s}_n)_{n \geq 1}$ is an optimal solution to our mechanism design problem (1). In the remainder of this paper, we will refer to this mechanism as DP^* , standing for optimal dynamic pricing. An important remark is that because DP^* satisfies the additional constraint (IB) (see §2.2), it is also optimal for the impatient bidders case: in line with results described in the literature for discrete asset selling problems, the option to recall past offers is worthless in our environment -this result would likely break down if we were to assume instead the number of potential bidders to be finite or the valuation distribution to be unknown (see Riley and Zeckhauser 1983).

Finally, we provide a method to compute explicitly the sequence of prices $(p_k)_{1 \leq k \leq K}$ associated with DP^* : for $k \geq 2$, p_k maximizes the right-hand side of the Bellman equation

$$R_k = \mathcal{G}(\alpha) \max_{p \in V} [(1 - F(p))(p + R_{k-1}) + F(p)R_k], \quad (8)$$

where R_k is the optimal expected discounted revenue-to-go when k items are available for sale at time 0. Thus we also have $R_k = E[\mathcal{G}(\alpha)^{N(p_k)}] (p_k + R_{k-1})$, where $N(p_k)$ is a positive geometric random variable with parameter $F(p_k)$. Substituting the right-hand side of this last expression calculated explicitly into the first-order condition associated with (8) yields the recursive system appearing in the following theorem, which summarizes our analysis.

Theorem 3 *Under assumptions (A1) – (A4), the dynamic mechanism DP^* solves (1): it supports a dominant equilibrium maximizing the seller's expected discounted revenue among*

all mechanisms where full transfer payments occur instantaneously upon item allocations. It is characterized by a sequence of prices $(p_k)_{1 \leq k \leq K}$, and consists of using a fixed price of p_k for the $(K - k + 1)$ -th sale. This sequence can be computed using the recursion

$$\begin{cases} p_1 \text{ solves (7); } R_1 = j(p_1) \\ p_k \text{ solves } j(p_k) = p_k \frac{\mathcal{G}(\alpha)(1-F(p_k))}{1-\mathcal{G}(\alpha)F(p_k)} - R_{k-1} \frac{1-\mathcal{G}(\alpha)}{1-\mathcal{G}(\alpha)F(p_k)} \\ R_k = R_{k-1} + j(p_k) \end{cases} \quad (9)$$

where R_k is the optimal expected discounted revenue-to-go when k items are available for sale at time 0. DP^* is also optimal for the impatient bidders problem and the time-strategic bidders problem obtained by respectively adding constraint (IB) and (IC') to (1).

In order to interpret system (9), we use the analogy between mechanism design and monopoly pricing introduced in Bulow and Roberts (1989). Let us first consider the single-item case $K = 1$: Replacing in the usual Cournot framework the "quantity" q with the probability of sale $1 - F(p_1)$, the marginal revenue associated with a bidder with valuation p_1 is precisely given by $\frac{d}{dq} [p_1(1 - F(p_1))] = p_1 + \frac{dp_1}{dq} q = j(p_1)$, the left-hand side of (7). Besides, the right-hand side of (7) is equal to the expected discounted revenue $R_1 = p_1 E [\mathcal{G}(\alpha)^{N(p_1)}]$ obtained when starting at $t = 0$ to sell an item for a fixed price of p_1 (as before $N(p_1)$ denotes a positive geometric random variable with parameter $F(p_1)$). So equation (7) is just the statement that the winning bidder should be determined so that marginal revenue from the current sale equals marginal cost or salvage value, that is the (discounted) revenue to be expected if the current sale attempt were to fail. Incidentally, note that as intuition suggests p_1 increases as α increases: when the seller's time value decreases, he can charge a higher price as the longer time to sell it entails becomes less penalizing. The equation in the second line of (9) is identical to (7), except that its right-hand side has a second term capturing the negative impact of the postponement of the $k-1$ remaining sales on the marginal cost/salvage value of the $(K - k + 1)$ -th item; The equations in its first and third lines imply $R_k = \sum_{i=1}^k j(p_i)$ for $k \geq 1$, which expresses that the optimal expected discounted revenue when k items are available is the sum of the individual marginal revenues associated with the sales of those items -the analogue of this result in the classical static optimal mechanism design framework is that the expected revenue from an auction is equal to the expected marginal revenue of the winning bidder (see Bulow and Roberts 1989).

In the final part of this section, we prove and interpret two basic properties of the mechanism DP^* - the first one is also established in a discrete-time setting equivalent to ours by

Das Varma and Vettas (2001).

Proposition 2 *The sequence of prices $(p_k)_{1 \leq k \leq K}$ characterizing the mechanism DP^* is decreasing with k , so that unit prices increase as sales occur over time.*

Proof. Using the previous remark that $R_k = E [\mathcal{G}(\alpha)^{N(p_k)}] (p_k + R_{k-1})$, we can write $p_k = \underset{p \in V}{\operatorname{argmax}} \mathcal{T}(p, R_{k-1})$, where $\mathcal{T} : V \times [0, +\infty) \rightarrow R$ is defined by $\mathcal{T}(p, y) = \frac{1-F(p)}{1-\mathcal{G}(\alpha)^{F(p)}}(p+y)$. Since \mathcal{T} is twice differentiable and $\frac{\partial^2 \mathcal{T}(p,y)}{\partial p \partial y} = \frac{f(p)}{[1-\mathcal{G}(\alpha)^{F(p)}]^2} (\mathcal{G}(\alpha) - 1) \leq 0$, this function has decreasing differences in (p, y) on $V \times [0, +\infty)$, thus Theorem 2.8.1 in Topkis (1998) applies, so that $\underset{p \in V}{\operatorname{argmax}} \mathcal{T}(p, y)$ is decreasing with y on $[0, +\infty)$. Because $(R_k)_{k \geq 1}$ is an increasing sequence (the discounted revenue of any mechanism when $k - 1$ items are available at time 0 can a fortiori be achieved when k items are available), it follows that $(p_k)_{1 \leq k \leq K}$ decreases with k . ■

The intuition for this last Proposition is that delaying the first sale when K items are available increases the discounting of the revenues associated with all K future sales, whereas a delay in (say) the last sale only impacts its own time-discount factor. As a result, the optimal mechanism DP^* initially uses a relatively low price for the first sale, which reduces on average not only the time until this first sale but also the epoch (and discounting) of all subsequent sales. As more sales occur and less items are left to be sold, DP^* progressively increases the selling price as short-term revenues become more attractive in the trade-off between selling price and sales epochs introduced by time discounting. Proposition 2 is important from a theoretical standpoint, because as mentioned in §2.2 it implies that DP^* is also optimal for the mechanism design problem obtained by replacing the incentive compatibility constraint (IC) in (1) by the more stringent constraint (IC') (see the discussion after (1) in §2.2). In words, because prices prescribed by DP^* increase over time, for any bidder and any demand realization the only potential consequence of intentionally delaying the type input into the mechanism is a reduction of discounted profit.

The following Proposition describes the limiting behavior of the optimal mechanism as the number of items initially available becomes large.

Proposition 3 *The sequence of prices $(p_k)_{k \geq 1}$ prescribed by DP^* has a finite limit $p^* = \lim_{k \rightarrow +\infty} p_k$ such that $j(p^*) = 0$.*

Proof. The existence of $\lim_{k \rightarrow +\infty} p_k$ is immediate as the sequence $(p_k)_{k \geq 1}$ is decreasing (from

Proposition 2) and it is bounded from below (by \underline{v}). It follows from the last equation in (9) that for $k \geq 1$, $R_k = \sum_{i=1}^k j(p_i)$. But because $(R_k)_{k \geq 1}$ is an increasing sequence, this implies $j(p_k) \geq 0$ for all $k \geq 1$. Besides, $R_k = E[\mathcal{G}(\alpha)^{N(p_k)}](p_k + R_{k-1})$ and $\begin{cases} N(p_k) \geq 1 \\ p_k \leq \bar{v} \end{cases}$ implies $R_k \leq \mathcal{G}(\alpha)(\bar{v} + R_{k-1})$; this implies in turn $R_k \leq \frac{\mathcal{G}(\alpha)\bar{v}}{1-\mathcal{G}(\alpha)}$ for all $k \geq 1$, thus $(R_k)_{k \geq 1}$ is both increasing and bounded from above and therefore has a finite limit. Because $(R_k)_{k \geq 1}$ is the serie with general term $j(p_k)$, this entails $\lim_{k \rightarrow +\infty} j(p_k) = 0$ and by continuity $j(p^*) = 0$. ■

To understand Proposition 3, it is useful to recall the interpretation for the second equation in (9): as the number of items becomes large, the marginal cost/salvage value of the very first item to be sold goes to zero, because the short-term profits to be derived from a postponed sale of this item tend to be overshadowed by the resulting heavier discounting of the (many) subsequent sales. Note that this result is also consistent with the sale of an item with no salvage value to a single potential buyer having a valuation distribution with c.d.f. $F(\cdot)$: the first order condition associated with $\max_p[p(1 - F(p))]$ is precisely equivalent to $j(p) = 0$.

3. Numerical Study

The purpose of our numerical study is to compare under different scenarios the performance of the optimal mechanism DP^* derived in the previous section with that of two other dynamic mechanisms widely used in practice: fixed posted price and online auction. In the remainder of this paper, we will refer to a mechanism ψ characterized by a vector of parameters ω as $\psi(\omega)$, to the expected discounted revenue associated with this mechanism as $E[\psi(\omega)]$, and to the mechanism obtained when ω is chosen to maximize $E[\psi(\omega)]$ as ψ^* . We turn now to the precise definition of the two mechanisms just mentioned:

Optimal Fixed Price (FP^*): A fixed posted price p is used for all K transactions, and this price is chosen optimally. That is, p maximizes over V the expected discounted revenue $E[FP(p)] = pE[\sum_{k=1}^K \alpha^{\tau'_k}]$ obtained with this mechanism, where τ'_k is the epoch at which the k^{th} transaction is concluded. We have $\tau'_k \sim \tau'_{k-1} + \sum_{i=n_{k-1}+1}^{n_k} x_i$, where $(x_i)_{i \geq 1}$ are i.i.d with $x_i \sim x$ and $n_k - n_{k-1} \geq 1$ is geometric with parameter $F(p)$. An easy calculation yields

$$E[FP(p)] = \frac{\mathcal{G}(\alpha)}{1 - \mathcal{G}(\alpha)} p (1 - F(p)) \left(1 - \left(\frac{\mathcal{G}(\alpha)(1 - F(p))}{1 - \mathcal{G}(\alpha)F(p)} \right)^K \right), \quad (10)$$

which we maximize using numerical methods to obtain $E[FP^*]$.

Optimal Online Auction (OA^*): At a specified closing date T , a maximum of K items are awarded by decreasing order of the bids submitted so far, provided they are higher than the reserve price r , for a price equal to the maximum of the highest rejected bid and r . Known as *Dutch* or *Open* auction, this mechanism is the most common auction format for selling multiple identical items on the internet (it is used in particular on the site eBay). Interestingly, with an appropriate reserve price this mechanism is optimal when time is not discounted and the number of bidders is fixed (Maskin and Riley 1989). Also, this mechanism supports a dominant equilibrium whereby bidders truthfully bid their valuation; if $\mathcal{V}_r(T) = \{v'_1, v'_2, \dots, v'_{N_r(T)}\}$ is the set of valuations larger than or equal to r of the bidders who arrived before T , the corresponding expected discounted revenue is thus $E[OA(T, r)] = \alpha^T E[\max(r, v'_{(K+1)}) \times \min(N_r(T), K)]$, where $v'_{(K+1)}$ is the $(K + 1)$ -th highest valuation in $\mathcal{V}_r(T)$ and by convention $v'_{(K+1)} = 0$ if $N_r(T) \leq K$. Note that the optimal value r^* of r is independent of T and is obtained by solving $j(r) = 0$ (Proposition 4 in Maskin and Riley). Consequently, for simplicity we will assume from now on that the reserve price r is always set to r^* , and omit the dependence on r when writing $OA(T) \equiv OA(T, r^*)$. To finally obtain $E[OA^*]$, we find in our experiments the optimal value of T using numerical methods.

In the interest of tractability, we assume for the remainder of this section that the bidders arrival process is Poisson with rate 1, and that valuations are uniformly distributed in $[0, 10]$; we also define the interest rate β through the relation $\alpha = 1/(1 + \beta)$. Under these assumptions, an easy calculation shows in particular that $E[OA(T)]$ specializes to

$$10\alpha^T \left[\frac{\lambda T}{4} \left(1 - \Gamma^*\left(K, \frac{\lambda T}{2}\right) \right) + K\Gamma^*\left(K + 1, \frac{\lambda T}{2}\right) - \frac{K(K + 1)}{\lambda T} \Gamma^*\left(K + 2, \frac{\lambda T}{2}\right) \right], \quad (11)$$

where $\Gamma^*(a, z) \equiv \int_z^\infty t^{a-1} e^{-t} dt / \int_0^\infty t^{a-1} e^{-t} dt$ is the incomplete gamma function ratio. In the two next subsections we address the main following questions: (i) What is the suboptimality of the mechanisms FP^* and OA^* with respect to DP^* for various values of β and K ? and (ii) How robust are the mechanisms FP^* , OA^* and DP^* relative to the choice of the parameters p , T and $(p_k)_{1 \leq k \leq K}$, respectively?

3.1. Mechanism Suboptimality. Table 1, in which the suboptimality of a mechanism ψ is denoted $S[\psi] \equiv (E[DP^*] - E[\psi]) / E[DP^*]$, summarizes the results of our suboptimality

β	K=1				K=10				K=50			
	E[DP*]	S[FP*]	S[OA*]	S[MA ₁ *]	E[DP*]/K	S[FP*]	S[OA*]	S[MA ₁₀ *]	E[DP*]/K	S[FP*]	S[OA*]	S[MA ₅₀ *]
0.1%	9.39	0%	2.68%	2.68%	8.66	0.6%	6.89%	6.89%	7.31	1.4%	15.23%	15.23%
0.2%	9.15	0%	3.83%	3.83%	8.15	0.8%	9.80%	9.80%	6.39	1.8%	21.65%	18.96%
0.3%	8.96	0%	4.73%	4.73%	7.78	0.9%	12.06%	12.06%	5.75	2.0%	26.60%	18.34%
0.4%	8.81	0%	5.49%	5.49%	7.47	1.0%	13.97%	13.97%	5.25	2.1%	30.77%	16.94%
0.5%	8.68	0%	6.17%	6.17%	7.21	1.1%	15.67%	15.67%	4.85	2.2%	34.43%	15.44%
0.6%	8.57	0%	6.79%	6.78%	6.99	1.2%	17.21%	17.21%	4.50	2.2%	37.70%	14.02%
0.7%	8.46	0%	7.36%	7.35%	6.78	1.3%	18.63%	18.37%	4.21	2.2%	40.65%	12.71%
0.8%	8.37	0%	7.90%	7.88%	6.60	1.3%	19.96%	19.04%	3.95	2.2%	43.32%	11.54%
0.9%	8.28	0%	8.41%	8.38%	6.43	1.4%	21.21%	19.40%	3.73	2.2%	45.72%	10.48%
1.0%	8.19	0%	8.89%	8.85%	6.28	1.4%	22.38%	19.57%	3.52	2.1%	47.87%	9.54%

Table 1: $E[DP^*]$, $S[FP^*]$, $S[OA^*]$ and $S[MA_K^*]$ for $\beta \in [0.1\%, 1\%]$ and $K \in \{1, 10, 50\}$.

study (we defer the definition of the mechanism MA^* appearing in the 5th, 9th and last columns of this table to the end of the present subsection).

A first observation is that for a fixed value of K the expected discounted revenue associated with the three mechanisms considered decreases with the interest rate β – this is hardly surprising as an increase in the interest rate alone only amounts to a higher time penalty through a lower value of the discount factor α . Secondly, over the range of market environments considered, the ranking of FP^* and OA^* with respect to expected discounted revenue is always the same, i.e. $E[DP^*] \geq E[FP^*] \geq E[OA^*]$. However, the suboptimality of OA^* observed is much higher than that of FP^* : the loss in discounted revenue resulting from a restriction of the optimal dynamic pricing mechanism to a single price seems to be considerably smaller than the loss incurred when using instead an online auction, a more drastic departure from that mechanism structure. In fact, because the suboptimality $S[SP^*]$ never exceeds a couple of percentage points over the range of scenarios we consider, one may question in practice the benefit of using DP^* over SP^* , as a dynamic pricing mechanism is arguably harder to implement and may not be as popular with buyers than a fixed price one. Another point is that the performances of FP^* and OA^* relative to DP^* both deteriorate when time becomes more valuable, but this effect is much more sensitive for OA^* than it is for FP^* . In the limit where the value of time goes to zero (or equivalently $\alpha \rightarrow 1$), the dispersion of the prices $(p_k)_{1 \leq k \leq K}$ characterizing DP^* becomes negligible (as the second term in the right-hand side of the recurrence equation satisfied by $(p_k)_{1 \leq k \leq K}$ in (9) disappears), so

that the mechanisms DP^* and FP^* become identical – note that this is also consistent with the intuitive interpretation provided for Proposition 2 above. The data in Table 1 suggests that the suboptimality of OA^* also goes to zero with β ; although our model is nonsensical when time has no value at all because the potential number of bidders is infinite, its limit when the interest rate goes to zero thus seems to be consistent with the optimality of OA^* for a static market environment proven in Maskin and Riley (1989) (see also our discussion of Wang 1993 below). Note that this limit behavior is highly relevant given the range of likely values for the interest rate in most practical settings: for example, with an average arrival rate of one bidder every four hours, an interest rate of 30% per annum corresponds to $\beta = 0.012\%$ when the unit of time is set, as in our experiments, such that $\lambda = 1$. Finally, increasing the number of items for sale K amplifies the suboptimality of FP^* and OA^* relative to DP^* . However, here again this phenomenon is more sensitive for OA^* than it is for FP^* ; as a result, the performance of OA^* relative to FP^* also deteriorates as more and more items are put for sale.

In summary, in our experiments DP^* always performs the best (this is predicted by the theory developed in §2) and FP^* is very close to optimal, always outperforming OA^* . In addition, both the performance of FP^* and OA^* relative to DP^* and that of OA^* relative to FP^* deteriorate when the value of time and/or the number of items for sale increase. However, for low values of the interest rate and moderate numbers of items for sale, which should be the norm in many applications, the suboptimality of the optimal online auction mechanism for the criteria of expected discounted revenue only amounts to a few percentage points.

These results should be interpreted in light of Wang (1993), who compares the expected revenue derived from a sequence of optimal auctions with that of an optimal posted price mechanism in a model with a single item for sale, fixed display/storage cost rates, auction setup costs and Poisson bidder arrivals. Interestingly, when the cost rates are the same and there are no auction setup costs, the hypotheses enabling a meaningful comparison with our model, he finds that the two mechanisms generate the same expected revenue. Because there is no time discounting in Wang’s model, this result is consistent with the limiting behavior we observe for $S[OA^*]$ when the interest rate goes to zero.

This also points out to a limitation of the comparison we have performed so far between

posted prices and online auction mechanisms: in our model, a seller using an auction will only run one single auction, regardless of how many items are left unsold after the initial bidding period. In contrast, Wang considers a sequence of however many auctions are necessary to sell one item; this suggests more generally the multiple auctions mechanism $MA_K(T_1, \dots, T_K)$ consisting of a sequence of however many online auctions it takes to sell the K items, each with the same structure as $OA(T)$, but where the bidding period T_k of each auction in the sequence is chosen dynamically as a function of the number of items k still unsold at that point (Vulcano, van Ryzin and Maglaras 2002 treat a somewhat related problem). Note that the reserve price r_k of each auction starting with k items left is set as before such that $j(r) = 0$, independently of k , which (in contrast to the single auction case) is not optimal. In fact, setting $(r_k, T_k) = (p_k, 0)$ for each k , where $(p_k)_{1 \leq k \leq K}$ are the prices characterizing DP^* , precisely achieves the optimal mechanism DP^* – the only reason why we introduce the mechanism MA_K here is to explore the hypothesis that the better performance of DP^* and FP^* over OA^* may be explained by the "unfair" advantage of never leaving any items unsold.

With uniform valuations and a Poisson bidder arrival process with rate λ , the expression $E[MA_1(T_1)] = E[OA(T_1)] / \left(1 - \alpha^{T_1} e^{-\frac{\lambda T_1}{2}}\right)$ is readily derived, and we can maximize it over $T_1 > 0$ using numerical methods to obtain $E[MA_1^*]$. Introducing next the notations $OA_k(T)$ for the (single) online auction mechanism with bidding period T when k items are for sale and $(T_\ell^*)_{1 \leq \ell \leq k}$ for the values of $(T_\ell)_{1 \leq \ell \leq k}$ maximizing $E[MA_k(T_1, \dots, T_k)]$, we have for $k > 1$

$$E[MA_k((T_\ell^*)_{1 \leq \ell \leq k-1}, T_k)] = E[OA_k(T_k)] + \alpha^{T_k} E[MA_k((T_\ell^*)_{1 \leq \ell \leq k-1}, T_k)] P(N_{r^*}(T_k) = 0) + \alpha^{T_k} \sum_{\ell=1}^{k-1} P(N_{r^*}(T_k) = \ell) E[MA_{k-\ell}^*],$$

and we finally obtain $E[MA_k^*]$ for $k > 1$ through the recursion

$$E[MA_k^*] = \max_{T_k} \left[\frac{E[OA_k(T_k)] + \alpha^{T_k} \left(\sum_{\ell=1}^{k-1} P(N_{r^*}(T_k) = \ell) E[MA_{k-\ell}^*] \right)}{1 - \alpha^{T_k} P(N_{r^*}(T_k) = 0)} \right].$$

The columns of Table 1 reporting $S[MA_K^*]$ show that in situations with relatively low values of the interest rate β and the number of items for sale K , including some where the suboptimality of OA^* is larger than 17%, the performances of OA^* and MA_K^* are virtually

identical. In contrast, when β and/or K increase, the advantage of MA_K^* over OA^* can become significant. However, an inspection of the optimal bidding periods $(T_k^*)_{1 \leq k \leq K}$ of the auction sequence MA_K^* (not shown here) reveals that in the cases where $E[MA_K^*] \gg E[OA^*]$, there always exists a k_0 such that $T_k^* = 0$ for $k_0 \leq k \leq K$. In words, the mechanism MA_K^* initially adopts then a fixed price strategy with posted price r^* such that $j(r^*) = 0$ (the right-hand side of the second equation in (9) expressing the marginal cost of selling in DP^* with k items left is close to zero when β and/or k are large, see Proposition 3). When less than k_0 items are left, MA_K^* switches to an auction strategy (when the marginal cost of selling at r^* becomes too high). This switching behavior also explains the observation from the last column of Table 1 that the suboptimality $S[MA_K^*]$ is a monotonic function of neither K nor β : $S[MA_K^*]$ starts increasing as a function of K or β precisely when the impact of the initial fixed price behavior mode begins to shadow that of the final auction mode.

For a relatively low interest rate and number of items for sale, we may therefore counter on experimental grounds the argument that the worse performance of auction mechanisms in our setting can be explained by their potential failure to sell all the items. Rather, a bidding mechanism seems to inherently not be as time-efficient a revenue generator then. However, while the switching behavior of MA_K^* is interesting, it unfortunately prevents us from drawing a similar conclusion for the cases where β and/or K are high.

3.2. Mechanism Robustness. In order to evaluate the robustness of DP^* , FP^* and OA^* with respect to the choice of parameters, we plot $E[FP(p)]$, $E[OA(T)]$ and $E[DP((p_k)_{1 \leq k \leq K})]$ in the same market environments, using equations (10), (11) and

$$E[DP((p_k)_{1 \leq k \leq K})] = \sum_{k=1}^K p_k \prod_{\ell=1}^k \frac{\mathcal{G}(\alpha)(1 - F(p_\ell))}{1 - \mathcal{G}(\alpha)F(p_\ell)}, \quad (12)$$

respectively. In order to plot $E[DP((p_k)_{1 \leq k \leq K})]$ as a function of a single variable for comparison purposes, we introduce the price p^* maximizing $E[FP(p)]$, the prices $(p_k^*)_{1 \leq k \leq K}$ characterizing DP^* , and define

$$p_k(p) = \begin{cases} \frac{p}{p^*} p_k^* & \text{for } 0 \leq p \leq p^* \\ p_k^* + (10 - p_k^*) \frac{p - p^*}{10 - p^*} & \text{for } p^* < p \leq 10 \end{cases} .$$

Figures 1 and 2 contain graphs of $E[FP(p)]$, $E[DP((p_k(p))_{1 \leq k \leq K})]$ and $E[OA(T)]$ for the two cases $\beta = 0.1\%$ and $\beta = 10\%$, when $K = 1$ and $K = 50$, respectively. Functions

$E[FP(p)]$ and $E[DP((p_k(p))_{1 \leq k \leq K})]$ are plotted over the range $p \in [0, 10]$ (lower X -axis), while $E[OA(T)]$ is plotted over $T \in [0, 2T^*]$ (upper X -axis).




Figure 1: $E[FP(p)]$ and $E[OA(T)]$ for $K = 1$ and $\beta \in \{0.1\%, 10\%\}$.

In these graphs, the flatness of each curve around its maxima provides a measure for the robustness of the corresponding mechanism: if a discounted revenue curve is relatively flat around its maxima, a deviation from the optimal parameter value is not likely to be very penalizing, while a sharp curve indicates the contrary. We first observe that according to this measure, OA^* seems to be quite robust in all cases, even if it becomes slightly less robust when the number of items K is high and the interest rate β is low, or when K is low and β is high. In contrast, the mechanisms FP^* and DP^* , which seem roughly equivalent from the perspective of robustness, are significantly less robust than OA^* across the range of cases considered – this should not be under-estimated as a potential reason for the popularity of online auctions in practice. However, the robustness of FP^* and DP^* improves relatively significantly when the interest rate and/or the number of items for sale




Figure 2: $E[DP((p_k(p))_{k \leq K})]$, $E[FP(p)]$ and $E[OA(T)]$ for $K = 50$ and $\beta \in \{0.1\%, 10\%\}$.

increase. This last observation is quite remarkable: the environments where the dynamic and fixed pricing mechanisms DP^* and FP^* seem to most significantly outperform the auction mechanism OA^* in terms of expected discounted revenue (large K and/or high β , see §3.1) thus coincide with those where OA^* has only a reduced advantage over DP^* and FP^* in terms of robustness. In that sense, our numerical study suggests unambiguous guidelines for when to use a (possibly dynamic) posted price mechanism versus an auction mechanism.

4. Conclusion

In this paper, we formulate and solve a continuous-time dynamic mechanism design problem for the sale of multiple identical items when participants are time-sensitive. Although we do make some relatively restrictive assumptions about the market environment, we still find it noteworthy that the optimal mechanism, a dynamic pricing scheme where the posted price increases after each sale, can be fully characterized; this is achieved in particular by

equation (7) and system (9), which gather in concise forms all the basic problem data. In particular, our analysis relies on the study of two generic mechanism properties, discreteness and stability, which we hope to be of interest to others researching this topic.

While our numerical study makes the strong assumption that valuations are uniformly distributed, so that its conclusions should be treated with caution, it may still suggest some useful guidelines for practitioners. More specifically, we first find that the relative benefit of dynamic pricing over a well-chosen fixed posted price mechanism is relatively small in this environment. Secondly, according to the criteria of both expected discounted revenue and robustness, posted prices should be used over online auctions when the number of items for sale is large and/or the market is particularly time-sensitive. With a small number of items for sale and/or low time-discounting, the significantly better robustness of online auctions should more than make up for their only slightly lower performance in terms of expected discounted revenue.

Natural extensions to the present study include relaxing the unit demand assumption (perhaps using some of the techniques developed for a static environment by Maskin and Riley), and the assumption that all market participants share the same time-discount factor, arguably the two strongest in our model. Another particularly interesting extension would be to consider in our framework an adaptive mechanism design problem, whereby the valuation distribution is not known initially but rather progressively inferred by observing participants' actions. Finally, we hope to further validate the experimental results presented here, either by studying actual transaction data or by conducting other experiments with more diverse distributional assumptions.

Acknowledgment. I would like to thank Rakesh Vohra, Martin Lariviere, Paul Milgrom and Donald Rosenfield for suggesting relevant references. Lawrence Wein, Steve Graves, Jay Sethuraman, Leonid Kogan, Costis Maglaras and the participants of the Industrial Engineering seminar at Lehigh University, the Operations Management and Applied Probability seminars at MIT and the 2002 Revenue Management Conference at Columbia University offered many insightful remarks. Finally, I am particularly grateful to Andreea Cepoiu for conducting some of the numerical experiments, and to the MIT Undergraduate Research Opportunities Program for funding her.

References

- Bertsekas, D. P. (1995). *Dynamic Programming and Optimal Control*, Volume 1. Belmont, MA: Athena Scientific.
- Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *Journal of Political Economy* 97(5), 1060–1090.
- Chow, Y. S., H. Robbins, and D. Siegmund (1971). *Great Expectations: The Theory of Optimal Stopping*. Boston: Houghton Mifflin Company.
- Elfving, G. (1967). A persistency problem connected with a point process. *Journal of Applied Probability* 4, 77–89.
- Karlin, S. (1962). Stochastic models and optimal policy for selling an asset. In K. J. Arrow, S. Karlin, and H. Scarf (Eds.), *Studies in Applied Probability and Management Science*, pp. 148–158. Stanford University Press.
- Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys* 3(13), 227–286.
- Kohn, M. G. and S. Shavell (1974). The theory of search. *Journal of Economic Theory* 9, 93–123.
- Lucking-Reiley, D. (2000). Auctions on the internet: What’s being auctioned, and how? *Journal of Industrial Economics* 48(3), 227–252.
- Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica* 52(6), 1473–1518.
- Maskin, E. and J. Riley (1989). Optimal multi-unit auctions. In F. Hahn (Ed.), *The Economics of Missing Markets, Information and Games*, pp. 312–335. Oxford University Press.
- McAfee, R. P. and J. McMillan (1988). Search mechanisms. *Journal of Economic Theory* 44, 99–123.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6, 58–73.
- Riley, J. and R. Zeckhauser (1983). Optimal selling strategies: When to haggle, when to hold firm. *Quarterly Journal of Economics* 98(2), 267–289.
- Siegmund, D. O. (1967). Some problems in the theory of optimal stopping rules. *Annals of*

Mathematical Statistics 38(6), 1627–1640.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton, New Jersey: Princeton University Press.

Varma, G. D. and N. Vettas (2001). Optimal dynamic pricing with inventories. *Economics Letters* 72(3), 335–340.

Vulcano, G., G. V. Ryzin, and C. Maglaras (2002). Optimal dynamic auctions for revenue management. *Management Science* 48(11).

Wang, R. (1993). Auctions versus posted-price selling. *The American Economic Review* 83(4), 838–851.