

Dynamic Mechanism Design for Online Commerce

Jérémie Gallien

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, jgallien@mit.edu

This paper is a mechanism design study for a monopolist selling multiple identical items to potential buyers arriving over time. Participants in our model are 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. Invoking the revelation principle, we restrict our attention to direct dynamic mechanisms taking a sequence of valuations and arrival epochs as input. We define two properties (discreteness and stability), and prove under further distributional assumptions that we may at no cost of generality consider only mechanisms satisfying them. This effectively reduces the mechanism input to a sequence of valuations and leads to formulate the problem as a dynamic program (DP). As this DP is equivalent to a well-known infinite-horizon asset-selling problem, we finally characterize the optimal mechanism as a sequence of posted prices increasing with each sale. Remarkably, this result rationalizes somewhat the frequent restriction to dynamic pricing policies and impatient buyers assumption. Our numerical study indicates that, under various valuation distributions, 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.

Subject classifications: games/group decisions, bidding/auctions: mechanism design; dynamic programming/optimal control, applications: asset selling, optimal stopping; marketing, pricing: dynamic pricing, online auctions.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received November 2002; revision received January 2005; accepted January 2005.

1. Introduction

The problem of maximizing revenue when selling goods to potential buyers arriving over time has received considerable attention from researchers and practitioners alike (see Elmaghraby and Keskinocak 2003). However, while much work has been done on how posted (or reservation) prices should be set and updated dynamically in that environment (the active academic literature on dynamic pricing dates back to the early 1960s), there is surprisingly little justification available for why price-based policies should be used over other types of selling strategies. In particular, the recent spectacular growth of the Internet website eBay suggests that online auctions (see Lucking-Reiley 2000 for a description of their many possible formats) may sometimes constitute a viable, and perhaps even better, alternative. Our objective in the present paper is to identify, among these two strategies and as many others as possible, the profit-maximizing mechanism for selling identical goods to a stream of self-interested and time-sensitive buyers.

We begin with a (necessarily incomplete) review of the large literature available on the sale of identical items by a monopolist over time. When the seller announces prices, this challenge is usually referred to as the “posted price” or “dynamic pricing” problem (see Bitran and Caldentey 2003 and Elmaghraby and Keskinocak 2003 for recent surveys). When the seller instead either accepts or rejects the

buyers’ bids, it is known as the “house-selling,” “asset-selling,” or “reservation price” problem, which is one of the primary motivations for the development of the theory of optimal stopping (Chow et al. 1971) and the theory of search (Mortensen 1986). While that distinction is irrelevant to our study (where the selling mechanism is endogenous), we refer the reader to Arnold and Lippman (2001) for a discussion of the differences between price-taking and price-setting models. Another differentiating feature among these studies is the structure of the buyers’ arrival process: “discrete-time” models with constant interarrival times (e.g., Das Varma and Vettas 2001), “continuous-time” models where the arrivals follow a Poisson process (e.g., Kincaid and Darling 1963, Gallego and van Ryzin 1994, Arnold and Lippman 2001), or more generally a renewal process (e.g., Karlin 1962, Elfving 1967). We also assume a renewal process, but under the restrictions imposed (essentially increasing failure rate (IFR) interarrival time distribution) we show that the corresponding price determination problem is ultimately equivalent to that studied in Das Varma and Vettas (2001) and Arnold and Lippman (2001). In fact, our Proposition 3 is equivalent to their Proposition 1 and Theorem 11 (respectively), and our Proposition 4 is equivalent to Theorem 12 from the latter paper (Arnold and Lippman 2001). Most of these studies either assume “no recall,” “myopic,” “nonstrategic,” or “impulse” buyers who will never wait (any transaction must

occur immediately upon their arrival) or, on the contrary, “full recall” buyers who are completely time insensitive. In contrast, we assume (as in Besanko and Winston 1990) that buyers are “rational” or “time strategic,” i.e., that they may a priori choose to trade off waiting time against selling price strategically, according to a specified utility function with nontrivial time discounting. Finally, while a finite time horizon is more realistic in some settings (see Kincaid and Darling 1963, Gallego and van Ryzin 1994 and most other references cited in Bitran and Caldentey 2003), our study considers an infinite time horizon (as in Elfving 1967, Das Varma and Vettas 2001, and Arnold and Lippman 2001).

Note that all the papers cited in the previous paragraph restrict a priori the range of mechanisms considered to dynamic posted (or reservation) prices, and focus then on the value of the price parameters characterizing the optimal choice within that family. In contrast, the problem of mechanism design consists of identifying the allocation and payment rules (taking strategies of self-interested participants as an input) maximizing the expected revenue predicted by a specified game solution concept. The approach for solving such problems has largely been developed by Myerson (1981), who considers a market environment with one item for sale to a fixed number of risk-neutral bidders with private valuations drawn from a commonly known distribution. 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 environments, notably by Maskin and Riley (1989, 1984) to markets with multiple identical units and risk-averse bidders—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 a dynamic market environment and a mechanism design problem, the present paper is an attempt to bridge the branches of literature described in the previous two paragraphs. A remarkable similar attempt is in 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 related paper is McAfee and McMillan (1988), which also considers a discrete-time model with a fixed bidder acquisition cost and time-insensitive bidders, but a finite buyer population and a predetermined selling strategy; through an extension of the revelation principle to a dynamic market environment, these authors also identify the corresponding optimal mechanism. A significant part of our analysis consists of transforming the continuous-time problem considered into an equivalent discrete-time one. Consequently, it is unsurprising that the mechanism we eventually obtain bears some resemblance to the ones derived in those last

two papers. However, our model uniquely captures both time sensitivity of participants and probabilistic features of dynamic bidder arrivals, which we believe are key drivers of trading interactions occurring on online auction sites such as eBay. The present 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 strategic 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 investigated, and an exact formulation of the corresponding mechanism design problem in §2.2. Next, §2.3 contains the definition of two properties for dynamic mechanisms, along with the theoretical result that if an optimal mechanism exists, there must also exist an optimal mechanism satisfying them. This result is key to the solution of the mechanism design problem considered, which we develop in §2.4—see Theorem 3 for a summary. A numerical study is presented in §3. Finally, §4 contains concluding remarks, and the proofs of all stated results (except Theorem 3), as well as a summary of mathematical notations used in this paper, can be found in the appendix.

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 $t_1 \leq t_2 \leq \dots$ are exogenous and follow a renewal process characterized by its transform $\mathcal{G}(z) = E[z^x]$, with $x \sim t_{n+1} - t_n$ for $n \geq 0$ ($t_0 = 0$ by convention). Each buyer n has a linear utility characterized by a privately known random valuation function $v_n(t)$ (maximum willingness to pay), where t denotes the time when the transaction is realized (if applicable). Most of this paper explicitly focuses on the case $v_n(t) = v_n 1_{[t_n, +\infty)}(t)$ for all potential buyers n , where v_1, v_2, \dots are independent and identically distributed (i.i.d.) random variables following a distribution known to the seller and characterized by a density $f(\cdot)$, cumulative distribution function (c.d.f.) $F(\cdot)$, and compact support $V = [\underline{v}, \bar{v}]$ with $0 \leq \underline{v} < \bar{v} < +\infty$. The process $\varphi = \{(v_1, t_1), (v_2, t_2), \dots\}$ is thus a marked point process where arrival times and marks are independent. We refer to this valuation structure as the *patient* bidders case, sometimes also referred to in the literature as the case “with recall.” However, the analysis ultimately extends to the *impatient* bidders case (or “without recall”) $v_n(t) = v_n 1_{\{t_n\}}(t)$. As is discussed later in §2.2.2, the solution concept we use (dominant equilibrium) makes our analysis robust with respect to the information structure assumed for the buyers, so we voluntarily leave its specification

incomplete here. Finally, the seller and the potential buyers have a time-discount factor $\alpha \in (0, 1)$. While this concept will be defined more precisely in §2.2.4 along with the participants' utility functions, its loose meaning is that the net value to the seller (respectively, bidder n) of any profit y occurring at time t is $\alpha^t y$ (respectively, $\alpha^{t-t_n} y$).

In addition to the seller's risk neutrality and the stationarity of demand, four other salient assumptions about the market environment are made in this paper:

ASSUMPTION 1. $h(t) \equiv E[\alpha^{x-t} \mid x > t]$ is a nondecreasing function of t (where $x \sim t_{n+1} - t_n$).

ASSUMPTION 2. $j(v) \equiv v - (1 - F(v))/f(v)$ is a nondecreasing function of v .

ASSUMPTION 3. The seller and the potential buyers have the same discount factor $\alpha \in (0, 1)$.

ASSUMPTION 4. Potential buyers have unit demand.

Assumption 1 is satisfied by all interarrival time distributions with IFR. It allows restriction of the search to mechanisms where allocation and payment decisions occur upon bidder arrivals, and typically rules out multimodal distributions, where the expected time until the next arrival may drastically increase when conditioned on the time elapsed since the last arrival. Assumption 2 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, this extension is left to future research to facilitate our focus on the dynamic aspect of mechanism design. Assumption 3 is arguably more restrictive (note, however, that the time to which the discount factor applies is different for each market participant): In some situations, the seller may act for a company while buyers may be end consumers, in which case their respective time sensitivities would likely follow different structures. Interestingly, however, the buyers' discount factor seems immaterial in the optimal solution eventually obtained (see §2.4), which suggests that relaxing Assumption 3 may not change the nature of the optimal mechanism—this conjecture is not resolved here, however. Finally, while Assumption 4 typically restricts the applicability of this model to situations where potential buyers are indeed end consumers, Maskin and Riley (1989) show in a static environment how this assumption can be relaxed. Even though we do not undertake this task here, in the future one may thus be able to successfully apply their method in a dynamic setting.

2.2. Problem Formulation

This subsection begins in §2.2.1 with a reminder of the main concepts and basic terminology of mechanism design.

Readers seeking a more exhaustive or technical presentation may refer to Klemperer (1999) or Fudenberg and Tirole (1991), and those already familiar with this topic may directly proceed to §2.2.2, which contains a discussion of the solution concept we use (dominant equilibrium) and the information structure we assume. We introduce our formal definition of a dynamic mechanism in §2.2.3, and that of the participants' utility functions in §2.2.4, which also contains a mathematical statement of the optimization problem we study.

2.2.1. Mechanism Design Terminology. In game-theoretic terms, any set of rules describing in a given market which outcome will result from a trading interaction is known as a *mechanism*. It is usually defined as follows: Its *allocation* rule determines whether each given market participant gets an item and, if appropriate, which one and when; its *payment* rule determines what each participant must pay as a result of the interaction (and if appropriate, to whom and when); the *information structure* describes the information available to each participant when trading decisions are made; and the *strategy space* describes how those trading decisions are expressed, that is, the exact format in which participants provide competitive information. For example, in a simple fixed-price mechanism with no supply constraint, the relevant information structure (for buyers) is just the listed price for the item, a buyer's strategy consists of expressing the number of items requested at that price (possibly zero), the allocation rule assigns to each buyer the number of items he requested, and the payment rule specifies that each participant pays the listed price times the number of items requested. For a second example, consider the sale of a single item through a first-price sealed-bid auction with no reserve price: The strategy space consists of payment offers (bids) made privately by each bidder, the information structure describes what each participant knows about the number of competitors and their willingness to pay, and the allocation rule assigns the item to the highest bidder, who according to the payment rule must pay the amount he bid (while other bidders pay nothing).

In static mechanism design studies, the incentives of each potential buyer are typically modelled with a utility function entirely characterized by a privately known valuation (maximum willingness to pay). In our dynamic environment, however, the primitives of each bidder n include not only a valuation v_n , but also an exogenous arrival time t_n . These two quantities together define the bidder's *type* $\varphi_n = (v_n, t_n)$, which for our purposes entirely characterizes him.

Once equipped with a description of the participants' utility functions (we defer a rigorous definition of these functions until §2.2.4), any given mechanism fulfills the definition of a game of incomplete information (see Part III in Fudenberg and Tirole 1991). As such, any solution concept or equilibrium (e.g., Bayesian Nash equilibrium, dominant equilibrium) adapted to this type of game can be applied to try and predict the set of strategies that will be

played by all players, or *strategy profile*. Formally, a strategy is defined as a mapping from each player's type and information set onto the action space. In turn, for a given strategy profile, every realization of the players' types corresponds to a value for the seller's (discounted) revenue, and one can thus compute the expectation of the seller's discounted revenue over all players' types.

For a given solution concept, the problem of mechanism design consists of finding the mechanism for which the seller's expected revenue associated with a strategy profile predicted by that concept is the highest possible.

2.2.2. Game Solution Concept and Information Structure. As in several other mechanism design studies (e.g., Laffont and Maskin 1982, Chung and Ely 2002), the solution concept we use in this paper is dominant equilibrium (DE); this corresponds to the requirement that every participant's strategy is optimal for him for every possible combination of the other players' strategies (see §2.2.4 for a mathematical formulation). In contrast, the other game solution concept widely used in the literature, Bayesian Nash equilibrium (BNE), only assumes that each player's strategy maximizes his expected utility conditional on his beliefs of the other players' strategies. We observe that our DE formulation seems consistent with our assumption that buyers have unit demand, as the end consumers it typically applies to are arguably more risk averse in practice than the seller. The primary reason for that choice here, however, is mathematical tractability, as we are only able to prove the feasibility of various mechanisms constructed during the analysis (see in particular the proof of Theorem 1) under a dominant equilibrium formulation.

There are, otherwise, respective advantages and disadvantages to both solution concepts (see Chapter 23 in Mas-Colell et al. 1995, and Chung and Ely 2003, for more complete discussions). One drawback of BNE is that it relies on fairly strong assumptions about the participants' beliefs and their cognitive or computational abilities, e.g., that all players share the exact same beliefs about the other players' types, that all players have the ability to compute a utility expectation conditional on those beliefs, and that all players believe that all the other players will perform that computation as well. As a result, the predictions only supported by a BNE are likely to be far less robust to misspecifications of the information structure or incorrect assumptions about the participants' behavior than predictions supported by a DE, and in that sense, are more risky. On the other hand, dominant equilibria do not always exist (this is not an issue in the present study), and the much stronger restrictions they impose (a DE is also a BNE, but the converse is not necessarily true) may result in optimal mechanisms yielding lower (predicted) revenue; note that the choice between BNE and DE is thus somewhat reminiscent of the classical trade-off in finance between expected return and volatility.

To the best of our knowledge, Mookherjee and Reichelstein (1992) describe the most generic conditions to date

under which any outcome supported by a BNE can also be supported with a DE. Unfortunately, the market environment they consider assumes single-dimensional types, and it is not clear how their results could be extended to our setting with a two-dimensional type space. As a result, our restriction to DE here may very well entail a loss of (predicted) revenue for the seller, and we are unable to determine how much that potential loss may amount to.

Finally, an important consequence of our DE formulation is that, in contrast with a BNE formulation, our revenue predictions are independent of the information structure we assume for the buyers (besides the private knowledge of their own individual type). That is, the conditions imposed by a DE insure that buyers' behavior will not be affected by any beliefs and/or knowledge about their competitors' valuations and arrival times. Consequently, we voluntarily left the buyers' information structure incompletely specified in our market environment description in §2.1. For a specific treatment of the issue of robustness with respect to the information structure, see Bergemann and Morris (2001) and references therein. We do assume, however, that the seller knows both the buyers' interarrival time and valuation distributions, and the fact that they are all independent.

2.2.3. Formal Mechanism Definition. In full generality, a mechanism can be defined as a pair of allocation and payment functions, both taking as input variable a strategy profile (see §2.2.1 for definitions). A major apparent difficulty, however, is that there is no 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 other mechanism design studies we are aware of, this difficulty is resolved by invoking the so-called revelation principle. This principle allows us to restrict the search with no loss of generality to *direct* mechanisms, where the strategy space is nothing but the type space: In the present 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 where participants truthfully reveal their type. We do not discuss here the rigorous game-theoretic justifications of this principle, and refer the reader instead to Myerson (1981) for a more formal statement 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 are now ready to define a *direct dynamic mechanism* ψ as 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 mappings $q^n: \Phi^n \mapsto \{0, 1\}^n$ and $y^n: \Phi^n \mapsto (\mathbb{R}^+)^n$, where

$$\Phi^n \equiv \{\varphi^n = ((v_1, t_1), \dots, (v_n, t_n)) \in (V \times \mathbb{R}^+)^n: t_i \leq t_j \text{ for } i \leq j \leq n\} \quad (1)$$

and $\Phi^n \equiv \{(\varphi^n, t), \varphi^n \in \Phi^n, t \in \mathbb{R}^+: t_n \leq t\}$. In that definition, $q_i^n(\varphi^n, t)$ and $y_i^n(\varphi^n, t)$, respectively, represent the cumulative allocation to, and transfer payment from, bidder i at time t , after the arrival of n bidders described by the stream φ^n —note that random allocation rules are not investigated here. We also define the space of all demand streams as $\Phi = \{\varphi = (\varphi_n)_{n \geq 1}: \text{for all } n \geq 1, \varphi^n \in \Phi^n\}$, and refer to the probability measure relative to φ as P .

A first natural restriction on the mechanisms $\psi = (q^n, y^n)_{n \geq 1}$ we consider is an *availability* constraint: The total number of items allocated at any point in time can never exceed the total number of items available, K . This can be stated mathematically as (AC): $\sum_{i=1}^n q_i^n(\varphi^n, t) \leq K$ for all $n \geq 1, \varphi^n \in \Phi^n$, and $t \geq t_n$.

A second, more salient, restriction we impose for the sake of tractability is that each bidder n may either: (i) never get any item, and not ever pay anything; or (ii) be allocated an item at some time $\tau_n(\varphi) \geq t_n$ after he arrives, and simultaneously pay a positive amount $y_n(\varphi)$ for it, exclusive of any other transfer payment to the seller. In particular, we exclude delayed and continuous payment schemes. Accordingly, the mechanism can alternatively be characterized through the following random variables, defined for each bidder i : $q_i(\varphi) \in \{0, 1\}$ (final allocation), $y_i(\varphi) \in \mathbb{R}^+$ (final payment), $s_i(\varphi) \in \{i, i+1, \dots\}$ (allocation time index), and $\tau_i(\varphi) \in [t_i, +\infty]$ (allocation time). The relationship with the original definition of a dynamic mechanism $\psi = (q^n, y^n)_{n \geq 1}$ is

$$(q_i^n(\varphi^n, t), y_i^n(\varphi^n, t)) = \begin{cases} (q_i(\varphi), y_i(\varphi)) & \text{if } n \geq s_i(\varphi) \text{ and } t \geq \tau_i(\varphi), \\ (0, 0) & \text{otherwise,} \end{cases} \quad (2)$$

and the availability constraint can be expressed as (AC): $\sum_{i=1}^{+\infty} q_i(\varphi) \leq K$ for all $\varphi \in \Phi$. In the case (i) described above where bidder i is never allocated an item and $(q_i(\varphi), y_i(\varphi)) = (0, 0)$, the r.v. $s_i(\varphi)$ and $\tau_i(\varphi)$ are defined by extension to be both equal to $+\infty$. From now on, the dependence of all r.v. on φ and ψ will be omitted when it is clear from context.

While we will use the component functions of (q^n, y^n) in §2.3 to study some properties of the optimal mechanism, in subsequent sections we will only need the sets of limit variables $(q_n, y_n, \tau_n)_{n \geq 1}$ or $(q_n, y_n, s_n)_{n \geq 1}$. When using those limit variables in our analysis, we will enforce the condition that they should exclusively be a function of the information available to date (or adapted to the stochastic process of bidder arrivals). That is, the mechanism may generate allocation and payment decisions at a given time based only upon the information available at that point. In (loose) measure-theoretic terms, the variables τ_i should define stopping times relative to an appropriately defined filtration associated with $(\varphi_n)_{n \geq 1}$, and the variables (q_i, y_i) should be \mathcal{F} -measurable, where \mathcal{F} is an appropriately defined σ -algebra associated with τ_i . Because no ambiguity about this issue will arise, however, we will not use these technical concepts further in the exposition.

2.2.4. Utility Functions and Optimization Problem

Statement. We define the expected utility function of the seller as $U_o(\psi) \equiv E_\varphi[\sum_{n=1}^{+\infty} \alpha^{t_n} y_n]$, and the (random) utility function of bidder $n \geq 1$ as $u_n(v_n) \equiv \alpha^{t_n - t_n}(v_n q_n - y_n)$. We thus consider an infinite-horizon setting so that our model does not capture the market environments with an intrinsic time limit (e.g., perishable goods) that are sometimes encountered in practice.

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. Also, while that definition describes the utility of bidder n truthfully reporting his valuation v_n (and arrival time t_n) to the mechanism, we can more generally define $u_n(v, v_n)$ as the utility of the n th bidder with true valuation v_n when reporting any other valuation v —but still his correct arrival epoch t_n —to the mechanism (clearly, $u_n(v_n, v_n) = u_n(v_n)$). Mathematically, $u_n(v, v_n) \equiv \alpha^{t_n[v] - t_n}(v_n q_n[v] - y_n[v])$, where $\varphi[v, n] \equiv \{\varphi^{n-1}, (v, t_n), \varphi_{n+1}, \dots\}$ and $\xi_n[v] \equiv \xi_n(\varphi[v, n])$ for any $v \in V$ and random variable ξ_n (e.g., q_n, y_n , and τ_n) defined on Φ and relative to index $n \geq 1$. The optimization problem can now be stated in the patient bidders case as

$$\begin{aligned} & \text{Max}_{\psi} U_o(\psi) \\ & \text{subject to (AC) and} \\ & u_n(v_n) \geq 0 \quad \text{for all } \varphi \in \Phi \text{ and } n \geq 1, \quad (IR) \\ & u_n(v_n) \geq u_n(v, v_n) \\ & \quad \text{for all } \varphi \in \Phi \text{ and } n \geq 1, v \in V. \quad (IC) \end{aligned} \quad (3)$$

The last two constraints in (3) 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 nonnegative, and (IC) guarantees *incentive compatibility*, namely, that a bidder cannot benefit from misrepresenting his valuation when interacting with the mechanism. In that 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. Defining $u_n(v, t, v_n, t_n)$ analogously as the utility obtained by the n th bidder with type (v_n, t_n) when pretending instead that his valuation is v and his arrival epoch $t \geq t_n$, an appropriate way to enforce full incentive compatibility would be to use instead the constraint (IC'): $u_n(v_n) \geq u_n(v, t, v_n, t_n)$. However, it turns out that the mechanisms obtained by solving (3) under Assumptions 1–4 (see §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 DE. Note that under the alternative requirement that the seller's revenue

prediction be supported by a BNE, these constraints would only have to hold in expectation, conditional on the bidders' information set. Note that in our formulation (IR) implies that no feasible mechanism will involve participation fees (i.e., $q_n = 0$ implies $y_n = 0$). These fees could conceivably arise and result in higher predicted revenues under a BNE formulation—see §2.2.2 for a related discussion. We observe, however, that the absence of participation fees seems preferable in a market where most potential buyers are end consumers (there are no buyer participation fees on the popular auction site eBay, for example), which is consistent with our unit-demand Assumption 4.

Finally, observe that a formulation of the mechanism design problem corresponding to the impatient bidders case is obtained by adding to (3) 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 (3).

2.3. Optimal Mechanism Properties

In this section, we introduce two properties of dynamic mechanisms holding under Assumption 1, and show that it is costless to restrict the search to mechanisms satisfying them; this later allows us to significantly simplify the problem formulation. Note that, as in the rest of this paper, the notion of mechanism optimality involved in the results of this section relates exclusively to our DE formulation (3).

DEFINITION 1. A mechanism ψ is said to be *discrete* if the sales it generates occur exclusively upon bidder arrivals. Equivalently, $\tau_n = t_{s_n}$ for all φ and $n \geq 1$.

Note that the definition of a discrete mechanism does not preclude the possibility of recall: It imposes a restriction on when an allocation may be generated, not on to whom it should be made. Because all mechanisms of interest for impatient bidders are discrete, the following proposition is nontrivial only with patient bidders:

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

The proof of Proposition 1 relies on a comparison between the seller's optimal expected discounted revenue to go at an arbitrary time and the expected future value of that quantity at the next bidder arrival given the same information; in that comparison, $h(t)$ represents the conditional additional discounting associated with delaying the future expected revenue stream until the next arrival. We will assume throughout that $h(t)$ is indeed a nondecreasing function of t , as required by Proposition 1. As stated in the discussion of that assumption above, it is satisfied in particular when the interarrival time x follows an IFR distribution. This justifies the restriction from now on to discrete mechanisms, and the simplified definitions $q^n: \Phi^n \mapsto \{0, 1\}^n$ and $y^n: \Phi^n \mapsto (\mathbb{R}^+)^n$ (see (1)). In this setting, the two variables τ_n and s_n become redundant, and we will exclusively use the latter from now on.

DEFINITION 2. A discrete mechanism $(\psi^n)_{n \geq 1}$ is said to be *stable* if the allocations and payments it generates for two arrival streams with identical valuation sequences but different arrival times are identical. Mathematically, if $\varphi = \{(v_1, t_1), (v_2, t_2), \dots\}$ and $\hat{\varphi} = \{(v_1, \hat{t}_1), (v_2, \hat{t}_2), \dots\}$, then $\psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n)$ for all n .

In other words, the allocations and payments generated by a stable mechanism do not depend on the bidders' arrival times, just their valuations and order of arrival—note that this definition is only sensible when applied to discrete mechanisms. Our interest in this second property is justified by the following theorem, which constitutes the main result of this section:

THEOREM 1. *If there exists a discrete optimal mechanism, there also exists an optimal mechanism that is stable.*

Found in the appendix, the proof of Theorem 1 consists of modifying an arbitrary optimal discrete mechanism to construct a mechanism that is shown by induction to still be optimal, but also stable. The combination of Proposition 1 and Theorem 1 justifies the restriction to discrete and stable mechanisms, and that we will work from now on with the even 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:

$$\begin{aligned}
 E[\alpha^{\tau_n} y_n] &= E[\alpha^{t_{s_n}} y_n] \\
 &\quad \text{because the mechanism is discrete,} \\
 &= E[E[\alpha^{t_{s_n}} y_n | s_n]] \\
 &\quad \text{by the law of total probability,} \\
 &= E[E[\alpha^{t_{s_n}} | s_n] E[y_n | s_n]] \\
 &\quad \text{because the mechanism is stable,} \\
 &= E[\mathcal{G}(\alpha)^{s_n} E[y_n | s_n]] \\
 &\quad \text{by the definition of } \mathcal{G}(\alpha), \\
 &= E[\mathcal{G}(\alpha)^{s_n} y_n] \\
 &\quad \text{by the law of total probability.}
 \end{aligned} \tag{4}$$

Note that the same reasoning applies when y_n is replaced with q_n in the equalities above, and it can also be used to prove that $E[\alpha^{\tau_n - t_n} \xi_n] = E[\mathcal{G}(\alpha)^{s_n - t_n} \xi_n]$, where ξ_n is either q_n or y_n . Now applying the monotone convergence theorem to the series with the general term given by (4) yields

$$U_o(\psi) = \sum_{n=1}^{+\infty} E[\mathcal{G}(\alpha)^{s_n} y_n]. \tag{5}$$

We next adapt and apply to this problem to the methodology developed in §4 of Myerson (1981). In particular,

the following lemma and proposition, justifying alternative expressions for the feasible space and objective function of (3) correspond, respectively, to his Lemmas 2 and 3. Their respective proofs (see the appendix) are mere adaptations to our setting of the ones found in his paper, with a couple of modifications because we use a DE (as opposed to a BNE) formulation, and consider an infinite number of bidders with time-discounted payments (as opposed to a finite number of bidders with no discounting).

LEMMA 2. *The conditions (IR) and (IC) in (3) are equivalent to the following set of conditions:*

$$u_n(\underline{v}) \geq 0 \quad \text{for all } n \geq 1 \text{ and } \varphi \in \Phi, \quad (IR1)$$

$$\alpha^{t_{s_n} v' 1^{-t_n}} q_n[v'] \leq \alpha^{t_{s_n} v 1^{-t_n}} q_n[v] \quad \text{for all } n \geq 1, \varphi \in \Phi, \text{ and } v \geq v' \in V^2, \quad (IP1) \quad (6)$$

$$u_n(v_n) = u_n(\underline{v}) + \int_{\underline{v}}^{v_n} \alpha^{t_{s_n} v 1^{-t_n}} q_n[v] dv \quad \text{for all } n \geq 1, \varphi \in \Phi, \text{ and } v_n \in V. \quad (IC1)$$

PROPOSITION 2. *For any mechanism ψ satisfying the feasibility conditions (6),*

$$U_o(\psi) = E \left[\sum_{n=1}^{+\infty} \left(v_n - \frac{1 - F(v_n)}{f(v_n)} \right) \mathcal{G}(\alpha)^{s_n} q_n \right] - \sum_{n=1}^{+\infty} \mathcal{G}(\alpha)^n E[u_n(\underline{v})]. \quad (7)$$

The second main step is based on the (classical) observation that only the variables $(q_n, s_n)_{n \geq 1}$ relative to the allocation decisions (i.e., not the payment variables y_n) appear in the first term of the right-hand side of (7). Besides, under (IR1) the maximum possible value of the second term is zero. Therefore, if $(q_n, s_n)_{n \geq 1}$ solves

$$\text{Max}_{(q_n, s_n)_{n \geq 1}} E \left[\sum_{n=1}^{+\infty} \left(v_n - \frac{1 - F(v_n)}{f(v_n)} \right) \mathcal{G}(\alpha)^{s_n} q_n \right] \quad (8)$$

subject to (AC) and (IP1),

and $(y_n)_{n \geq 1}$ can then be found such that (IC1) and (IR1) are satisfied as well as $E[u_n(\underline{v})] = 0$ for all n , then $(q_n, y_n, s_n)_{n \geq 1}$ is an optimal solution to the original problem (3).

When a single item is for sale ($K = 1$), the relaxation of (8) 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 $w \sim v - (1 - F(v))/f(v)$ (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 $\bar{w}/\mathcal{G}(\alpha) = \bar{w}P(w \leq \bar{w}) + \int_{\bar{w}}^{\infty} w dH(w)$, and reject all other offers (see Chapters 4 and 7 in Bertsekas 1995).

In the common case where $j(v) \equiv v - (1 - F(v))/f(v)$ is nondecreasing (see Assumption 2 discussed in §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

$$j(p_1) = p_1 \frac{\mathcal{G}(\alpha)(1 - F(p_1))}{1 - \mathcal{G}(\alpha)F(p_1)}. \quad (9)$$

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 (8) 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, defining by recurrence an increasing finite sequence $(n_l)_{1 \leq l \leq K}$ by $n_1 = \min\{n: v_n \geq p_K\}$, $n_l = \min\{n > n_{l-1}: v_n \geq p_{K-l+1}\}$ for $l \geq 2$ (arrival index of first and l th buyers, respectively), and

$$l(n) = \begin{cases} \max\{l: n_l < n\} & \text{if } n_1 < n, \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

(number of items sold by the n th arrival), the optimal allocation policy can be written as $\hat{q}_n = 1_{[p_{K-l(n)}, \bar{v}]}(v_n)$ and $\hat{s}_n = n/\hat{q}_n$. Note that $\mathcal{G}(\alpha)^{\hat{s}_n - n} \hat{q}_n = 1_{[p_{K-l(n)}, \bar{v}]}(v_n)$ and $p_{K-l(n)}$ 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 (8).

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 (7) can be expressed for all $v' \in V$ as

$$\mathcal{G}(\alpha)^n E[u_n(\underline{v})] = E[\mathcal{G}(\alpha)^{s_n} (v_n q_n - y_n) \mid v_n = v'] - \int_{\underline{v}}^{v'} E[\mathcal{G}(\alpha)^{s_n} q_n \mid v_n = v] dv. \quad (11)$$

Therefore, setting $\hat{y}_n = p_{K-l(n)} \hat{q}_n$ (the first bidder n arriving after $l(n)$ items have been sold and having a valuation greater or equal to $p_{K-l(n)}$ gets the next item, and his payment equals $p_{K-l(n)}$), $E[u_n(\underline{v})] = 0$ for all n as

$$\begin{aligned} & \int_{\underline{v}}^{v'} E[\mathcal{G}(\alpha)^{\hat{s}_n} \hat{q}_n \mid v_n = v] dv \\ &= \mathcal{G}(\alpha)^n \int_{\underline{v}}^{v'} 1_{[p_{K-l(n)}, \bar{v}]}(v) dv \\ &= \mathcal{G}(\alpha)^n (v' - p_{K-l(n)}) 1_{[p_{K-l(n)}, \bar{v}]}(v') \\ &= E[\mathcal{G}(\alpha)^{\hat{s}_n} (v_n \hat{q}_n - \hat{y}_n) \mid v_n = v']. \end{aligned} \quad (12)$$

In addition, it is straightforward to verify that \hat{y}_n satisfies constraints (IC1) and (IR1); we have therefore proven that the mechanism $(\hat{q}_n, \hat{y}_n, \hat{s}_n)_{n \geq 1}$ is an optimal solution

to the mechanism design problem (3). In the remainder of this paper, we refer to this mechanism as DP^* , standing for optimal dynamic pricing. An important remark is that because DP^* satisfies the additional constraint (IB) (see §2.2.4), 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 this environment—this result would likely break down if instead the number of potential bidders was assumed 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

$$R_k = \mathcal{G}(\alpha) \max_{p \in V} [(1 - F(p))(p + R_{k-1}) + F(p)R_k], \quad (13)$$

where R_k is the optimal expected discounted revenue-to-go when k items are available for sale at time 0. Thus, $R_k = E[\mathcal{G}(\alpha)^{N(p_k)}](p_k + R_{k-1})$, 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 (13) yields the recursive system appearing in the following theorem, which summarizes the analysis.

THEOREM 3. *Under Assumptions 1–4, the dynamic mechanism DP^* solves (3): 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*

$$\begin{cases} p_1 \text{ solves (9), } R_1 = j(p_1), \\ p_k \text{ solves } j(p_k) = p_k \frac{\mathcal{G}(\alpha)(1 - F(p_k))}{1 - \mathcal{G}(\alpha)F(p_k)} \\ \quad - R_{k-1} \frac{1 - \mathcal{G}(\alpha)}{1 - \mathcal{G}(\alpha)F(p_k)}, \\ R_k = R_{k-1} + j(p_k), \end{cases} \quad (14)$$

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 problems with impatient and time-strategic bidders obtained by adding, respectively, constraint (IB) and (IC') to (3).

To interpret system (14), we use the analogy between mechanism design and monopoly pricing introduced in Bulow and Roberts (1989). 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 (9). Besides, the right-hand side of (9) 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 $F(p_1)$). Therefore, Equation (9) 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 that it entails becomes less penalizing. The equation in the second line of system (14) is identical to (9), 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 that $R_k = \sum_{i=1}^k j(p_i)$ 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.

In the final part of this section, we prove and interpret two basic properties of the mechanism DP^* —the first one is also established and interpreted in a discrete-time setting equivalent to ours by Das Varma and Vettas (2001).

PROPOSITION 3. *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.*

The intuition for this 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 p_K 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 fewer 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 3 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 (3) by the more stringent constraint (IC') (see the discussion after (3) in §2.2). In other 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. Interestingly, the principle of selling a limited inventory of (near) identical items at posted prices increasing over time seems to be common in the sale of art photographs and real-estate units within a development project—although these situations clearly exhibit some unique specificities, we infer that some of the rationale we develop in this paper for using such a mechanism may also be relevant in those markets.

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

PROPOSITION 4. *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 $j(p^*) = 0$.*

To understand Proposition 4, it is useful to recall the interpretation for the second equation in (14): 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 this numerical study is to compare under different scenarios the performance of the optimal mechanism DP^* defined in the statement of Theorem 3 with that of two other dynamic mechanisms widely used in practice: fixed posted price and online auction. In the remainder of this paper, a mechanism ψ characterized by a vector of parameters ω is referred to as $\psi(\omega)$, the expected discounted revenue associated with this mechanism as $E[\psi(\omega)]$, and the mechanism obtained when ω is chosen to maximize $E[\psi(\omega)]$ as ψ^* . In addition, the sub-optimality of a mechanism ψ relative to DP^* is denoted by $S[\psi] \equiv (E[DP^*] - E[\psi])/E[DP^*]$. Finally, we define for convenience the *interest rate* β through the relation $\alpha = 1/(1 + \beta)$. We now turn 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. Its distribution satisfies $\tau'_k \sim \tau'_{k-1} + \sum_{i=n_{k-1}+1}^{n_k} x_i$, 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 $F(p)$. A straightforward calculation yields

$$E[FP(p)] = \frac{\mathcal{G}(\alpha)}{1 - \mathcal{G}(\alpha)} p(1 - F(p)) \left(1 - \left(\frac{\mathcal{G}(\alpha)(1 - F(p))}{1 - \mathcal{G}(\alpha)F(p)} \right)^K \right), \quad (15)$$

which we maximize over p 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 by then, 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 1989). 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 $OA(T) \equiv OA(T, r^*)$.

We assume in the next two §§3.1 and 3.2 that the bidder arrival process is Poisson with rate 1 and that valuations are uniformly distributed in $[0, 10]$. A lengthy but straightforward calculation shows that $E[OA(T)]$ specializes then to

$$10\alpha^T \left[\frac{\lambda T}{4} \left(1 - \Gamma\left(K, \frac{\lambda T}{2}\right) \right) + K\Gamma\left(K + 1, \frac{\lambda T}{2}\right) - \frac{K(K + 1)}{\lambda T} \Gamma\left(K + 2, \frac{\lambda T}{2}\right) \right], \quad (16)$$

where $\Gamma(a, z) \equiv \int_0^z t^{a-1} e^{-t} dt / \int_0^\infty t^{a-1} e^{-t} dt$ is the incomplete gamma function ratio; $E[OA^*]$ is then obtained using numerical methods. In the following §3.3, we assume instead that valuations follow more general $c \times \text{Beta}[a, b]$ distributions, and that the arrival process is either Poisson or deterministic with rate 1; $E[OA^*]$ is then computed using a commercial simulation-based optimization search software (OptQuest). Although this last method does not provide any guarantee of optimality, in all validation runs

we performed with uniform valuations and Poisson arrivals its output was within 0.01% of the value of $E[OA^*]$ computed through numerical optimization of (16). The next three subsections, respectively, address the following questions: (i) What is the suboptimality of the mechanisms FP^* and OA^* with respect to DP^* for various values of β and K ? (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? (iii) What is the impact of the distributional assumptions made to represent valuations and interarrival times?

3.1. Mechanism Suboptimality

Table 1 summarizes the results of this first set of experiments (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). Our main observations are the following:

- 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 .
- 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., $E[DP^*] \geq E[FP^*] \geq E[OA^*]$. While the first inequality and $E[DP^*] \geq E[OA^*]$ follow from the theory developed in §2, the second inequality is merely an empirical observation that may not always apply. In the environments we did study, 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 $S[FP^*]$ never exceeds a couple of percentage points over the range of scenarios considered,

one may question in practice the benefit of using DP^* over FP^* , as a dynamic pricing mechanism is arguably harder to implement and may not be as popular with buyers as a fixed price one.

- 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 β 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 (14) disappears), so that the mechanisms DP^* and FP^* become identical—note that this is also consistent with the intuitive interpretation provided for Proposition 3 above. The data in Table 1 suggest 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 the 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 $\lambda = 1$.

- 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 up for sale.

To summarize these 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

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

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

increases. 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 observed for $S[OA^*]$ when the interest rate goes to zero.

This also points to a limitation of the comparison 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 et al. 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 for introducing 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 $E[MA_1(T_1)] = E[OA(T_1)] / (1 - \alpha^{T_1} e^{-\lambda T_1/2})$ 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 $(T_l^*)_{1 \leq l \leq k}$ for the values of $(T_l)_{1 \leq l \leq k}$ maximizing $E[MA_k(T_1, \dots, T_k)]$, for $k > 1$,

$$\begin{aligned} E[MA_k((T_l^*)_{1 \leq l \leq k-1}, T_k)] \\ = E[OA_k(T_k)] + \alpha^{T_k} E[MA_k((T_l^*)_{1 \leq l \leq k-1}, T_k)] \\ \cdot P(N_{r^*}(T_k) = 0) + \alpha^{T_k} \sum_{l=1}^{k-1} P(N_{r^*}(T_k) = l) E[MA_{k-l}^*], \end{aligned}$$

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

$$\begin{aligned} E[MA_k^*] \\ = \max_{T_k} \left[\frac{E[OA_k(T_k)] + \alpha^{T_k} \left(\sum_{l=1}^{k-1} P(N_{r^*}(T_k) = l) E[MA_{k-l}^*] \right)}{1 - \alpha^{T_k} P(N_{r^*}(T_k) = 0)} \right]. \end{aligned} \quad (17)$$

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 increases, 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 other words, the mechanism MA_K^* then initially adopts a fixed-price strategy with posted price r^* such that $j(r^*) = 0$ (the right-hand side of the second equation in (14) 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 4). When fewer 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, the argument that the worse performance of auction mechanisms in this setting can be explained by their potential failure to sell all the items may therefore be countered on experimental grounds. Rather, a bidding mechanism inherently seems to not be as time-efficient a revenue generator then. However, while the switching behavior of MA_K^* may seem interesting, it unfortunately prevents us from drawing a similar conclusion for the cases where β and/or K are high.

3.2. Parameter Robustness

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 (15), (16), and

$$E[DP((p_k)_{1 \leq k \leq K})] = \sum_{k=1}^K p_k \prod_{l=1}^k \frac{\mathcal{G}(\alpha)(1 - F(p_l))}{1 - \mathcal{G}(\alpha)F(p_l)}, \quad (18)$$

respectively. 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

$$p_k(p) = \begin{cases} \frac{p}{p^*} p_k^* & \text{for } 0 \leq p \leq p^*, \\ p_k^* + (10 - p_k^*) \frac{p - p^*}{10 - p^*} & \text{for } p^* < p \leq 10. \end{cases} \quad (19)$$

Note that we thus consider the prices obtained by proportionately scaling the optimal prices, so that the resulting assessment of the robustness of DP^* may be somewhat optimistic. 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).

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. Observe first 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 underestimated 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 increases. 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, this numerical study suggests unambiguous guidelines for when to use a (possibly dynamic) posted-price mechanism versus an auction mechanism.

3.3. Distributional Robustness

3.3.1. Distribution Shape. In this third set of experiments, bidder valuations were assumed instead to follow $8 \times \text{Beta}[0.5, 0.3]$ and $20 \times \text{Beta}[2, 6]$ distributions. By design, these distributions have the same first two moments as the distribution $U[0, 10]$ used in the previous subsection, but a different support and shape ($[0, 8]$ with a right-skewed U shape and $[0, 20]$ with a left-skewed bell shape, respectively). Table 2 summarizes our results; it reports not only $E[DP^*]/K$, $S[FP^*]$, and $S[OA^*]$ as before, but also (in parentheses) the relative (respectively, absolute) variation of each expected discounted revenue (respectively, suboptimality) value relative to the corresponding one from Table 1 obtained with uniform valuations.

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

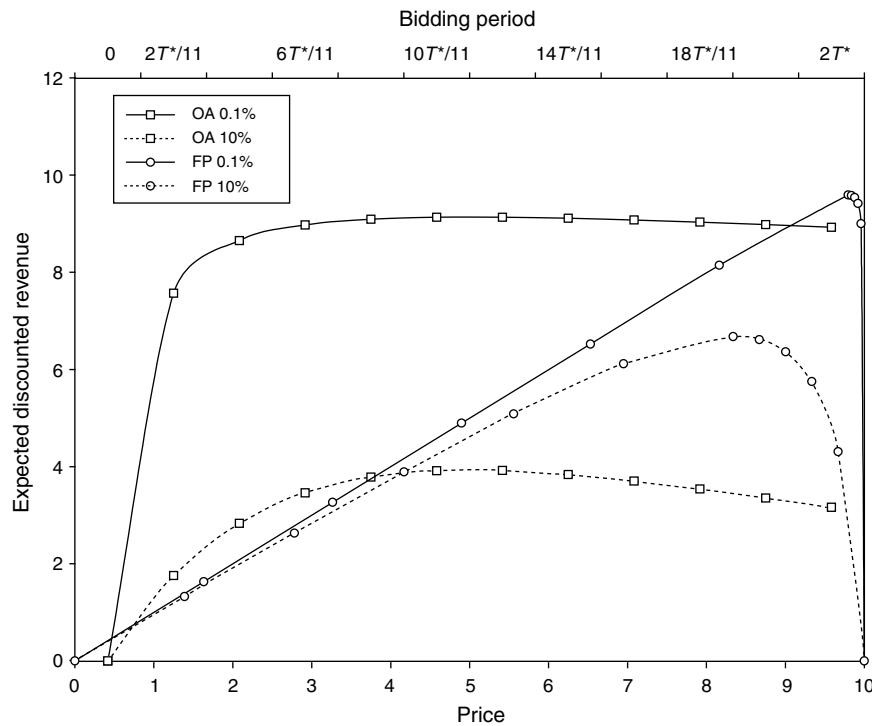



Figure 2. $E[DP((p_k(p))_{k \leq K})]$, $E[FP(p)]$, and $E[OA(T)]$ for $K = 50$ and $\beta \in \{0.1\%, 10\%\}$.



Our main observations are the following:

- The previous finding that the suboptimality of OA^* and FP^* increases with β and K and the inequalities $E[DP^*] \geq E[FP^*] \geq E[OA^*]$ remain valid here.

- Although limited as before to a couple of percentage points, $S[FP^*]$ was relatively larger under $20 \times \text{Beta}[2, 6]$ and smaller under $8 \times \text{Beta}[0.5, 0.3]$. Our explanation is that a larger support and longer right tail of the valuation distribution provide more leverage to the possibility of selling different items at different prices.

- For low values of β and K , the optimal discounted revenue $E[DP^*]$ is highest under $20 \times \text{Beta}[2, 6]$, while for high values of β and K it is highest under $8 \times \text{Beta}[0.5, 0.3]$. Our interpretation is that setting a value for the interest rate dictates how the trade-off between high price and time to sell should be addressed; this in turn effectively amounts to selecting which section of the valuation distribution support is most attractive. Also, increasing the number of items for sale K effectively increases the time sensitivity for the sale of the initial items (see the discussion after Proposition 3); thus, K also participates in selecting which part of the valuation distribution matters most. As a result, sale opportunities at prices larger than 10 provided by $20 \times \text{Beta}[2, 6]$ can be leveraged even though they have a small probability when time sensitivity and/or the number of items are low (cases $K = 1$ and $K = 10$, $\beta = 0.1\%$), while the much shorter time to sell at prices between 5 and 8 under $8 \times \text{Beta}[0.5, 0.3]$ becomes more attractive when time sensitivity and the number of items are high (case $\beta = 1\%$, $K = 10$). Finally, the case $\beta = 0.5\%$,

$K = 10$ illustrates the intermediary situation where the most relevant part of the support is (8, 10), providing an advantage to $U[0, 10]$ —note that this interpretation is supported by the relative values of the density functions of $8 \times \text{Beta}[0.5, 0.3]$, $U[0, 10]$, and $20 \times \text{Beta}[2, 6]$ on $[0, 20]$.

- The performance of OA^* relative to DP^* improved under $8 \times \text{Beta}[0.5, 0.3]$ and deteriorated under $20 \times \text{Beta}[2, 6]$. Our interpretation for this is rooted in the observation that $E[OA^*]$ primarily depends on the distribution of highest valuations, whereas $E[DP^*]$ depends on the probability that valuations are higher than optimally chosen prices. As a result, the relative concentration of likely values for the high-order statistics of $8 \times \text{Beta}[0.5, 0.3]$ resulting from the “┘” shape of its right tail effectively reduces the discrepancy between $E[DP^*]$ and $E[OA^*]$. On the contrary, this discrepancy increases when the distribution of the high-order statistics is more variable, as is the case under $20 \times \text{Beta}[2, 6]$.

3.3.2. Distribution Variability. Our last set of experiments, with results reported in Table 3, was designed to assess the impact of arrival and valuation variability on the relative performance of DP^* , FP^* , and OA^* . While the structure of Table 3 is the same as that of Table 2, the valuation distributions considered were instead $10 \times \text{Beta}[0.54, 0.54]$ and $10 \times \text{Beta}[1.84, 1.84]$, which have the same mean and support as the distribution $U[0, 10]$ used earlier, but with standard deviations 20% larger and 20% smaller, respectively. In addition, the last three rows of the table report the results of experiments conducted

Table 2. $E[DP^*]$, $S[FP^*]$, and $S[OA^*]$ with Poisson(1) arrivals and $20 \times \text{Beta}[2, 6]$ or $8 \times \text{Beta}[0.5, 0.3]$ distributions.

β (%)	$K = 1$			$K = 10$					
	$20 \times \text{Beta}[2, 6]$	$8 \times \text{Beta}[0.5, 0.3]$	$20 \times \text{Beta}[2, 6]$	$8 \times \text{Beta}[0.5, 0.3]$	$20 \times \text{Beta}[2, 6]$	$8 \times \text{Beta}[0.5, 0.3]$			
	$E[DP^*]$	$S[OA^*]$	$E[DP^*]/K$	$S[OA^*]$	$E[DP^*]/K$	$S[OA^*]$			
0.1	11.7 (+25%)	5.2% (+2.5%)	7.9 (-16%)	1.4% (-1.3%)	9.7 (+11%)	12.3% (+5.4%)	7.7 (-11%)	0.1% (-0.5%)	3.3% (-3.6%)
0.5	9.7 (+11%)	7.4% (+1.2%)	7.7 (-11%)	4.5% (-1.7%)	7.2 (-0.7%)	2.4% (+1.2%)	19.3% (+3.7%)	7.0 (-3%)	0.3% (-0.8%)
1.0	8.7 (+6%)	10.1% (+1.2%)	7.5 (-8%)	7.1% (-1.8%)	6.0 (-5%)	2.6% (+1.1%)	24.8% (+2.4%)	6.4 (+2%)	0.5% (-0.9%)

Table 3. $E[DP^*]$, $S[FP^*]$, and $S[OA^*]$ with Poisson(1) or deterministic arrivals and $10 \times \text{Beta}[0.54, 0.54]$ or $10 \times \text{Beta}[1.84, 1.84]$ distributions.

β (%)	$K = 1$			$K = 10$							
	$10 \times \text{Beta}[0.54, 0.54]$	$10 \times \text{Beta}[1.84, 1.84]$	$10 \times \text{Beta}[0.54, 0.54]$	$10 \times \text{Beta}[1.84, 1.84]$	$10 \times \text{Beta}[0.54, 0.54]$	$10 \times \text{Beta}[1.84, 1.84]$					
	$E[DP^*]$	$S[OA^*]$	$E[DP^*]$	$S[OA^*]$	$E[DP^*]$	$S[OA^*]$					
Poisson arrivals	0.1	9.72 (+3.6%)	2.0% (-0.7%)	8.81 (-6.1%)	3.0% (+0.3%)	9.21 (+6.4%)	0.3% (-0.3%)	5.3% (-1.6%)	7.95 (-8.2%)	0.9% (+0.3%)	7.6% (+0.7%)
	0.5	9.24 (+6.4%)	5.4% (-0.8%)	7.96 (-8.3%)	5.1% (-1.0%)	7.93 (+9.9%)	0.7% (-0.4%)	14.3% (-1.4%)	6.52 (-9.6%)	1.4% (+0.3%)	14.9% (-0.8%)
	1.0	8.84 (+7.9%)	8.1% (-0.7%)	7.45 (-9.1%)	7.7% (-1.2%)	6.98 (+11.1%)	1.0% (-0.4%)	22.1% (-0.3%)	5.68 (-9.5%)	1.7% (+0.3%)	20.5% (-1.9%)
Deterministic arrivals	0.1	9.72 (+3.6%)	1.9% (-0.8%)	8.81 (-6.1%)	3.0% (+0.3%)	9.21 (+6.4%)	0.3% (-0.3%)	5.2% (-1.6%)	7.95 (-8.2%)	0.9% (+0.3%)	7.6% (+0.7%)
	0.5	9.24 (+6.4%)	5.3% (-0.8%)	7.96 (-8.3%)	5.8% (-0.3%)	7.93 (+9.9%)	0.7% (-0.4%)	13.9% (-1.8%)	6.52 (-9.7%)	1.4% (+0.3%)	14.8% (-0.9%)
	1.0	8.84 (+7.8%)	7.7% (-1.1%)	7.44 (-9.2%)	7.1% (-1.8%)	6.97 (+11.0%)	1.0% (-0.4%)	21.3% (-1.1%)	5.67 (-9.6%)	1.7% (+0.3%)	19.9% (-2.5%)

instead with deterministic arrivals. Our observations are the following:

- Again, the previous finding that the suboptimality of OA^* and FP^* increases with β and K , the inequalities $E[DP^*] \geq E[FP^*] \geq E[OA^*]$, and the small suboptimality of FP^* all remain valid here.

- The marginal impact of switching to deterministic interarrival times was barely detectable in the case of DP^* and FP^* , and slightly higher but still very low in the case of OA^* . Predictably, the impact of switching to deterministic interarrival times on OA^* was all the more sensitive as the time sensitivity was high.

- Increasing (respectively, decreasing) the valuation variability, however, did result in a significant increase (respectively, decrease) of $E[DP^*]$, $E[FP^*]$, and $E[OA^*]$. At least within the range of time discount rates considered, all three mechanisms seem to benefit from a heavier right tail of the valuation distribution, that is, from the occasional appearance of a high-valuation bidder.

- The suboptimality $S[FP^*]$ slightly improved under $10 \times \text{Beta}[0.54, 0.54]$ and slightly deteriorated under $10 \times \text{Beta}[1.84, 1.84]$. We infer that this follows from the U shape of $10 \times \text{Beta}[0.54, 0.54]$, which concentrates in a smaller region the valuations targeted by the optimal prices under DP^* , and therefore offers a reduced advantage to DP^* over FP^* compared to $10 \times \text{Beta}[1.84, 1.84]$ (which exhibits a bell shape)—thus, this result may not hold under different distributions.

- The suboptimality $S[OA^*]$ slightly improved under $10 \times \text{Beta}[0.54, 0.54]$ and slightly deteriorated under $10 \times \text{Beta}[1.84, 1.84]$ for $\beta = 0.1\%$, but slightly improved under both distributions when the interest rate was 0.5% or 1%. We do not have any satisfactory explanation for this intriguing observation, and conjecture that it may be an artefact of the particular shapes of the distributions considered. However, note that, as for the previous point, the absolute impact of this effect is very small.

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 some relatively restrictive assumptions about the market environment are imposed, 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 (9) and system (14), which gather in concise forms all the basic problem data. Remarkably, this result rationalizes somewhat the restriction to dynamic pricing policies assumed a priori in many studies, as well as the typical assumption that buyers are impatient. In addition, the analysis presented relies on the study of two generic dynamic mechanism properties, discreteness and stability, which we hope to be of interest to others researching this topic.

In our numerical study, we found the benefit of dynamic pricing over a well-chosen fixed posted-price mechanism to be relatively small in this environment. Simulation experiments also suggested that, according to the criteria of both expected discounted revenue and parameter 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, however, the significantly better parameter robustness of online auctions should in practice more than make up for their slightly lower predicted performance in terms of expected discounted revenue. Because we observed these findings to hold under several different distributional assumptions, we believe that they may suggest useful guidelines for practitioners when selling goods online.

Future research will hopefully permit solving the same mechanism design problem under a BNE formulation, possibly generalizing to our setting along the way the set of conditions for dominant strategy implementation provided in Mookherjee and Reichelstein (1992). Other extensions could include relaxing the unit-demand assumption (perhaps using some of the techniques developed for a static environment by Maskin and Riley 1989), and the assumption that all market participants share the same time-discount factor. Further work may also focus on an adaptive mechanism design problem, where the valuation distribution is not known initially, but rather progressively inferred by observing participants' actions. Finally, the theoretical question of which dynamic mechanisms are optimal in a finite-horizon setting seems particularly important and is, to the best of our knowledge, unanswered so far.

Appendix

A.1. Summary of Notation

We use the following vector notation throughout: When w is any of the symbols in $\{q(\cdot), y(\cdot)\}$ (respectively, $\{v, t, \varphi\}$), w^n refers to the vector with n components (w_1^n, \dots, w_n^n) (respectively, (w_1, \dots, w_n)). Also, if ψ and $\hat{\psi}$ are dynamic mechanisms, their associated detailed variables follow the same notational modification and will be referred to as (say) q_n, y_n, s_n and $\hat{q}_n, \hat{y}_n, \hat{s}_n$, respectively.

Market Environment

K	number of identical items for sale
t_n	arrival time of bidder n
$x \sim t_n - t_{n-1}$	bidder interarrival time
$\lambda \equiv 1/E[x]$	bidder arrival rate
α	time-discount factor
$\beta \equiv 1/\alpha - 1$	interest rate
$\mathcal{G}(z) = E[z^x]$	transform of bidder interarrival time
$h(t) \equiv E[\alpha^{x-t} \mid x > t]$	conditional discount function

v_n	valuation of bidder n
$V = [v, \bar{v}]$	support of valuation distribution
$f(\cdot), F(\cdot)$	p.d.f. and c.d.f. of valuation distribution
$j(v) \equiv v - (1 - F(v))/f(v)$	virtual value function
$\varphi_n = (v_n, t_n)$	type of bidder n
$\varphi = \{(v_1, t_1), (v_2, t_2), \dots\}$	bidder arrival stream
Φ	set of possible values for φ
$\varphi^n = \{(v_1, t_1), \dots, (v_n, t_n)\}$	types of n first bidders
Φ^n	set of possible values for φ^n
Φ^{n^*}	set of possible values for (φ^n, t) with $t \geq t_n$
$\bar{\varphi}^n = \{(v_{n+1}, t_{n+1}), \dots\}$	types of bidders $n + 1$ and above (tail)
$u_n(v)$	random utility of bidder n when $v_n = v$
$u_n(v, v')$	random utility of bidder n when reporting v' instead of $v_n = v$
$U_0(\cdot)$	seller's expected utility

Mechanism

$\psi = (q^n, y^n)_{n \geq 1}$	dynamic mechanism
(q^n, y^n)	allocation and payment vectors after n arrivals
(q_i^n, y_i^n)	allocation to and payment from of the i th bidder after n arrivals
(q_n, y_n)	final allocation to and payment from the n th bidder
τ_n	allocation epoch for the n th bidder (possibly $+\infty$)
s_n	number of arrivals when the n th bidder receives an item (possibly $+\infty$)
n_k	buyer to whom the k th item is sold

Miscellaneous

\mathbb{R}^+	nonnegative real numbers
\mathbb{N}^*	positive integers
1_A	indicator function for set A
$\Gamma(a, z)$	incomplete gamma function ratio ($\equiv \int_0^z t^{a-1} e^{-t} dt / \int_0^\infty t^{a-1} e^{-t} dt$)
$A \times B$	2-cartesian product of sets A and B
A^n or $(A)^n$	n -cartesian product of set A
$N(p)$	positive geometric r.v. with parameter $F(p)$

A.2. Proof of Proposition 1

PROOF OF PROPOSITION 1. Define the r.v.s s'_k and τ'_k for $1 \leq k \leq K$ as the number of arrived bidders and time epoch at the k th item sale, respectively. Mathematically,

$$s'_k \equiv \inf \left\{ n \in \mathbb{N}^*: \exists t < t_{n+1} \text{ such that } \sum_{i=1}^n q_i^n(\varphi^n, t) \geq k \right\} \quad (20)$$

and $\tau'_k \equiv \inf \{ t \in \mathbb{R}^+: \sum_{i=1}^{s'_k} q_i^{s'_k}(\varphi^{s'_k}, t) \geq k \}$. Let n_k be the buyer to whom the k th item is sold, so that $(s_{n_k}, \tau_{n_k}) =$

(s'_k, τ'_k) and $q_{n_k}^n(\varphi^n, t) = 1 \Leftrightarrow n \geq s'_k$ and $t \geq \tau'_k$. Now 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), define the set $I(t_n) = \{n_j: 1 \leq j \leq k(t_n)\}$ (bidders to whom these k items have been sold), and consider the tail $\bar{\varphi}^n = \{(v_{n+1}, t_{n+1}), \dots\}$ of the stochastic process φ . From the principle of dynamic optimality, $(q^m, y^m)_{m \geq n}$ maximizes $E_{\bar{\varphi}^n}[\sum_{j=k(t_n)+1}^K \alpha^{t_j} y_{n_j} \mid (\varphi_i)_{\{i: 1 \leq i \leq n\} \setminus I(t_n)}]$ 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. Now denote 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 $|I| \leq K$: By definition, $J(\varphi^{n-1}, (v_n, t), I)$ is the optimal expected discounted revenue-to-go upon the n th arrival at time t when the arrival process to date is $(\varphi^{n-1}, (v_n, t))$ and the set of winning bidders immediately before time t is given by I . Because of the seller's time-sensitivity model, translating by $t' - t$ a revenue stream exactly amounts to multiplying its discounted value by $\alpha^{t'-t}$; otherwise, the bidders already arrived that are available for a sale from time t and t' on (respectively) are exactly the same, and the (translated) type distributions for all the bidders not yet arrived are also the same. Therefore, $\Theta(t) \equiv J(\varphi^{n-1}, (v_n, t), I)$ satisfies for each φ^{n-1} , v_n , and I the functional equation $\alpha^{-t} \Theta(t) = M$, where M is independent of t , which proves the existence of a function M such that $J(\varphi^{n-1}, (v_n, t), I) = \alpha^t M(\varphi^{n-1}, v_n, I)$.

Now applying 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 $\tau'_K = t$ between the n th and $(n + 1)$ 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

$$\begin{aligned} \alpha^t y &\geq E_{\varphi_{n+1}}[J(\varphi^n, \varphi_{n+1}, I_{K-1}) \mid t_{n+1} > t] \\ &= E_{(v_{n+1}, t_{n+1})}[\alpha^{t_{n+1}} V(\varphi^n, v_{n+1}, I_{K-1}) \mid t_{n+1} > t] \\ &\quad \text{(from the result just proven)} \\ &= E_{t_{n+1}}[\alpha^{t_{n+1}} \mid t_{n+1} > t] E_{v_{n+1}}[V(\varphi^n, v_{n+1}, I_{K-1})] \\ &\quad \text{(from the independence of } v_{n+1} \text{ and } t_{n+1}\text{).} \end{aligned} \quad (21)$$

Because the resulting inequality holds for $t \geq \max(t_n, \tau'_{K-1})$, invoking now the hypothesis that $E_{t_{n+1}}[\alpha^{t_{n+1}-t} \mid t_{n+1} > t]$ 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 $\tau'_K = \tau'_{K-1}$. However, note that in the latter case the same reasoning can be applied to $y = y_{n_K} + y_{n_{K-1}}$,

$t = \tau'_k = \tau'_{k-1}$, and I_{k-2} (and so on), it therefore follows 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 that $k_1 > 1$ (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 $\tau'_k = t_{s'_k}$ 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 $k_l = 1$, 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. \square

A.3. Proof of Theorem 1

PROOF OF THEOREM 1. Extending the definition of a stable mechanism, a mechanism is said to be 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 $v^n = \hat{v}^n$,

$$\begin{cases} \psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n) & \text{if } n \leq N, \\ \psi^n(\varphi^n) = \psi^n(\hat{\varphi}^n) & \text{if } n > N \text{ and } t_l = \hat{t}_l \text{ for } N < l \leq n. \end{cases} \quad (22)$$

Note that the notions of stability and N -stability for all $N \geq 1$ coincide. The proof begins with the following lemma, then proceeds to show by induction that for all $N \geq 1$, there exists a mechanism $\tilde{\psi}$ that is optimal and N -stable.

LEMMA 4. Let $(\psi^n)_{n \geq 1}$ be an N -stable mechanism feasible for (3), and let $(g^n)_{n \geq 1}$ be a family of mappings onto $(\Phi^n)_{n \geq 1}$ such that

(i) $\Delta^n(g^n(\varphi^n)) = \Delta^n(\varphi^n)$ for all $n \geq 1$ and $\varphi^n \in \Phi^n$, where $\Delta^n(\varphi^n) \equiv v^n$;

(ii) $\Lambda^n(g^m(\varphi^m)) = \Lambda^n(g^n(\varphi^n))$ for all $m \geq n \geq N + 1$, where $\Lambda^n(\varphi^m) \equiv t^n$; and

(iii) $\psi^m(g^m(\varphi^m)[v'_n, n]) = \psi^m(g^m(\varphi^m)[v'_n, n])$ for each $m \geq n \geq 1$, $\varphi^m \in \Phi^m$, and $v'_n \in V$.

The mechanism $(\psi^n)_{n \geq 1} \equiv (\psi^n \circ g^n)_{n \geq 1}$ is then also N -stable and feasible for (3).

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

$$(q_n(g(\varphi)), y_n(g(\varphi)), s_n(g(\varphi))) = (\dot{q}_n(\varphi), \dot{y}_n(\varphi), \dot{s}_n(\varphi)). \quad (23)$$

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

$$\begin{aligned} & \alpha^{t_{s_n(g(\varphi))}}(v_n q_n(g(\varphi)) - y_n(g(\varphi))) \\ & \geq \alpha^{t_{s_n(g(\varphi)[v'_n, n])}}(v_n q_n(g(\varphi)[v'_n, n]) - y_n(g(\varphi)[v'_n, n])), \\ & \alpha^{t_{s_n(\varphi)}}(v_n \dot{q}_n(\varphi) - \dot{y}_n(\varphi)) \\ & \geq \alpha^{t_{s_n(g(\varphi)[v'_n, n])}}(v_n q_n(g(\varphi)[v'_n, n]) - y_n(g(\varphi)[v'_n, n])) \\ & = \alpha^{t_{s_n(\varphi)[v'_n, n]}}(v_n \dot{q}_n(\varphi[v'_n, n]) - \dot{y}_n(\varphi[v'_n, n])), \end{aligned} \quad (24)$$

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. \square

Let $N \in \{0, 1, \dots\}$ and $\psi = (\psi^n)_{n \geq 1}$ be a discrete, N -stable solution to (3). All feasible mechanisms are defined by extension to be 0-stable, so that the base of the induction is 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 $E[\sum_{l=k(\varphi^N)+1}^K \alpha^{t'_l} y_{n_l} \mid \varphi^{N+1}]$ among all sequences $(\psi^j)_{j > N}$ such that $(\psi^j)_{1 \leq j \leq N} \times (\psi^j)_{j > N}$ is feasible for (3). 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 $k(\hat{\varphi}^N) = k(\tilde{\varphi}^N) = k < K$,

$$\begin{aligned} & \alpha^{-\hat{t}_{N+1}} E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \hat{\varphi}^{N+1} \right] \\ & = \alpha^{-\tilde{t}_{N+1}} E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \tilde{\varphi}^{N+1} \right]. \end{aligned} \quad (25)$$

By contradiction, assume instead that

$$\begin{aligned} & \alpha^{-\hat{t}_{N+1}} E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \hat{\varphi}^{N+1} \right] \\ & < \alpha^{-\tilde{t}_{N+1}} E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \tilde{\varphi}^{N+1} \right]. \end{aligned} \quad (26)$$

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

$$\begin{cases} \pi^n(t^n) = t^n & \text{for } n \leq N, \\ \pi^n(t^n) = (\tilde{t}^{N+1}, t_{N+2} - t_{N+1} + \tilde{t}_{N+1}, \dots, \\ \quad t_n - t_{N+1} + \tilde{t}_{N+1}) & \text{for } n \geq N + 1. \end{cases} \quad (27)$$

From Lemma 4, the mechanism $\dot{\psi} = \psi \circ g$ is feasible, yet

$$\begin{aligned} E \left[\sum_{l=k+1}^K \alpha^{t'_l} \dot{y}_{n_l} \mid \hat{\varphi}^{N+1} \right] & = \alpha^{\hat{t}_{N+1} - \tilde{t}_{N+1}} E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \tilde{\varphi}^{N+1} \right] \\ & > E \left[\sum_{l=k+1}^K \alpha^{t'_l} y_{n_l} \mid \hat{\varphi}^{N+1} \right]. \end{aligned} \quad (28)$$

The first equality above follows because $\{t_n\}$ is assumed to follow a renewal process, therefore by construction of g , the distribution of $\{\pi^n(t^n)\}_{n \geq N+1}$ conditional on $t^{N+1} = \hat{t}^{N+1}$ is exactly the same as that of $\{\tilde{t}^n\}_{n \geq N+1}$ conditional on \tilde{t}^{N+1} . Therefore, by construction of $\dot{\psi} = \psi \circ g$ and because $\hat{v}^{N+1} = \tilde{v}^{N+1}$, the distribution of the future allocation and payments resulting from mechanism $\dot{\psi}$ conditional

on $\widehat{\varphi}^{N+1}$ is exactly the same as that resulting from ψ conditional on $\widehat{\varphi}^{N+1}$, except that their actual timing has been translated by $\widehat{t}_{N+1} - \widetilde{t}_{N+1}$; the following inequality follows from the contradiction hypothesis. However, this inequality is strict and $(\psi^j)_{j \leq N} = (\widetilde{\psi}^j)_{j \leq N}$, which contradicts the optimality of $(\psi^j)_{j > N}$ for $\max_{(\psi^j)_{j > N}} E[\sum_{l=k(\widehat{\varphi}^N)+1}^K \alpha^{t_{l'}^j} y_{n_l} | \widehat{\varphi}^{N+1}]$ stated above, and proves (25).

Now define 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

$$\left\{ \begin{array}{l} \pi^n(t^n) = t^n \quad \text{for } n \leq N, \\ \pi^{N+1}(t^{N+1}) = (t^N, \widehat{t}_{N+1}(t_N)), \\ \quad \text{where } \widehat{t}_{N+1}(t_N) = t_N + E[x], \\ \pi^n(t^n) = (t^N, \widehat{t}_{N+1}(t_N), t_{N+2} - t_{N+1} + \widehat{t}_{N+1}(t_N), \dots, \\ \quad t_n - t_{N+1} + \widehat{t}_{N+1}(t_N)) \\ \quad \text{for } n \geq N+2. \end{array} \right. \quad (29)$$

Note that the mechanism $\widetilde{\psi}$ defined by $\widetilde{\psi}^n = \psi^n \circ g^n$ is $(N+1)$ -stable, and it is easy to check that it satisfies the conditions of Lemma 4, so that it is feasible for (3). Besides, for any $\varphi^{N+1} \in \Phi^{N+1}$ such that $k(\varphi^N) = k < K$, $\widetilde{k}(\varphi^N) = k$ because ψ is N -stable from the induction hypothesis and

$$\begin{aligned} & E \left[\sum_{l=1}^K \alpha^{t_{l'}^j} \widetilde{y}_{\widehat{n}_l} | \varphi^{N+1} \right] \\ &= E \left[\sum_{l=1}^k \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] + E \left[\sum_{l=k+1}^K \alpha^{t_{l'}^j} \widetilde{y}_{\widehat{n}_l} | \varphi^{N+1} \right] \\ &= E \left[\sum_{l=1}^k \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] \\ &\quad + \alpha^{t_{N+1} - \widehat{t}_{N+1}(t_N)} E \left[\sum_{l=k+1}^K \alpha^{t_{l'}^j} \widetilde{y}_{\widehat{n}_l} | g^{N+1}(\varphi^{N+1}) \right] \\ &= E \left[\sum_{l=1}^k \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] \\ &\quad + \alpha^{t_{N+1} - \widehat{t}_{N+1}(t_N)} E \left[\sum_{l=k+1}^K \alpha^{t_{l'}^j} y_{n_l} | g^{N+1}(\varphi^{N+1}) \right] \\ &= E \left[\sum_{l=1}^k \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] + E \left[\sum_{l=k+1}^K \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] \\ &= E \left[\sum_{l=1}^K \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right], \end{aligned} \quad (30)$$

where the first equality follows from the fact that $\alpha^{t_{l'}^j} y_{n_l}$ is \mathcal{F}_N -measurable for $l \leq k$ and $\widetilde{\psi}^n = \psi^n$ 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 ψ ; and the fourth

from (25). Equation (30) and the fact that $\widetilde{\psi}$ is N -stable imply, finally, that

$$\begin{aligned} U_0(\widetilde{\psi}) &= E \left[E \left[\sum_{l=1}^K \alpha^{t_{l'}^j} \widetilde{y}_{\widehat{n}_l} | \varphi^{N+1} \right] \right] \\ &= E \left[E \left[\sum_{l=1}^K \alpha^{t_{l'}^j} y_{n_l} | \varphi^{N+1} \right] \right] \\ &= U_0(\psi). \end{aligned} \quad (31)$$

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

A.4. Proof of Lemma 2

PROOF OF LEMMA 2. Note that $u_n(v', v) = u_n(v') + \alpha^{t_{s_n[v'] - t_n}(v - v')} q_n[v']$, so that (IC) is equivalent to

$$u_n(v) \geq u_n(v') + \alpha^{t_{s_n[v'] - t_n}(v - v')} q_n[v']. \quad (32)$$

This implies that

$$\begin{aligned} \alpha^{t_{s_n[v'] - t_n}(v - v')} q_n[v'] &\leq u_n(v) - u_n(v') \\ &\leq \alpha^{t_{s_n[v] - t_n}(v - v')} q_n[v], \end{aligned} \quad (33)$$

which in turn implies (IP1). Thus, $\alpha^{t_{s_n[v] - t_n}(v - v')} q_n[v]$ is increasing in v for every realization of φ ; that function is therefore Riemann integrable, and (33) also implies (IC1). Finally, note that (IR1) is a special case of (IR). Conversely, the integrand in the right-hand side of (IC1) is nonnegative because $q_n[v] \in \{0, 1\}$, so that $u_n(v_n) \geq u_n(\underline{v})$, and (IR) thus follows from (IR1). Now take $v_n \geq v' \in V^2$; (IC1) and (IP1) together imply that

$$\begin{aligned} u_n(v_n) &= u_n(v') + \int_{v'}^{v_n} \alpha^{t_{s_n[v] - t_n}(v - v')} q_n[v] dv \\ &\geq u_n(v') + \int_{v'}^{v_n} \alpha^{t_{s_n[v'] - t_n}(v - v')} q_n[v'] dv \\ &= u_n(v') + \alpha^{t_{s_n[v'] - t_n}(v_n - v')} q_n[v'], \end{aligned} \quad (34)$$

which has already been shown to be equivalent to (IC). The case $v_n < v'$ follows analogously after writing instead

$$\begin{aligned} u_n(v_n) &= u_n(v') - \int_{v_n}^{v'} \alpha^{t_{s_n[v] - t_n}(v - v')} q_n[v] dv \\ &\geq u_n(v') - \int_{v_n}^{v'} \alpha^{t_{s_n[v'] - t_n}(v - v')} q_n[v'] dv, \end{aligned} \quad (35)$$

which also implies (IC), and concludes the proof. \square

A.5. Proof of Proposition 2

PROOF OF PROPOSITION 2. From the definition of $u_n(v_n)$ and (IC1), we have

$$\begin{aligned} \alpha^{t_{s_n} - t_n} y_n &= \alpha^{t_{s_n} - t_n} v_n q_n - u_n(v_n) \\ &= \alpha^{t_{s_n} - t_n} v_n q_n - u_n(\underline{v}) - \int_{\underline{v}}^{v_n} \alpha^{t_{s_n[v] - t_n}(v - \underline{v})} q_n[v] dv. \end{aligned} \quad (36)$$

Taking expectations, using (4), and dividing by $\mathcal{G}(\alpha)^n$ yields

$$E[\mathcal{G}(\alpha)^{s_n} y_n] = E[\mathcal{G}(\alpha)^{s_n} v_n q_n] + \mathcal{G}(\alpha)^n \cdot E\left[\int_{\underline{v}}^{v_n} \alpha^{t_{s_n[v]} - t_n} q_n[v] dv\right] - \mathcal{G}(\alpha)^n E[u_n(\underline{v})]. \quad (37)$$

The second term of the right-hand side of (37) can be written as

$$\begin{aligned} \mathcal{G}(\alpha)^n E\left[\int_{\underline{v}}^{v_n} \alpha^{t_{s_n[v]} - t_n} q_n[v] dv\right] &= \mathcal{G}(\alpha)^n \int_{\underline{v}}^{\bar{v}} E\left[\int_{\underline{v}}^{v_n} \alpha^{t_{s_n[v]} - t_n} q_n[v] dv \mid v_n\right] f(v_n) dv_n \\ &= \mathcal{G}(\alpha)^n \int_{\underline{v}}^{\bar{v}} \left(\int_{\underline{v}}^{v_n} E[\alpha^{t_{s_n[v]} - t_n} q_n[v]] dv\right) f(v_n) dv_n \\ &= \int_{\underline{v}}^{\bar{v}} \left(\int_{\underline{v}}^{v_n} E[\mathcal{G}(\alpha)^{s_n[v]} q_n[v]] dv\right) f(v_n) dv_n \\ &= \left[\left(\int_{\underline{v}}^{v_n} E[\mathcal{G}(\alpha)^{s_n[v]} q_n[v]] dv\right) F(v_n)\right]_{\underline{v}}^{\bar{v}} \\ &\quad - \int_{\underline{v}}^{\bar{v}} E[\mathcal{G}(\alpha)^{s_n[v]} q_n[v]] F(v_n) dv_n \\ &= \int_{\underline{v}}^{\bar{v}} E[\mathcal{G}(\alpha)^{s_n[v]} q_n[v]] (1 - F(v)) dv \\ &= \int_{\underline{v}}^{\bar{v}} E\left[\frac{1 - F(v)}{f(v)} \mathcal{G}(\alpha)^{s_n[v]} q_n[v]\right] f(v) dv \\ &= E\left[\frac{1 - F(v_n)}{f(v_n)} \mathcal{G}(\alpha)^{s_n} q_n\right], \end{aligned} \quad (38)$$

where the first equality above follows from the law of total probability; the second from Fubini's theorem; the third from (4); the fourth from an integration by parts; the fifth from the fact that $F(\bar{v}) = 1$ and $F(\underline{v}) = 0$; and the last one from the law of total probability. A substitution in (37) yields

$$E[\mathcal{G}(\alpha)^{s_n} y_n] = E\left[\left(v_n - \frac{1 - F(v_n)}{f(v_n)}\right) \mathcal{G}(\alpha)^{s_n} q_n\right] - \mathcal{G}(\alpha)^n E[u_n(\underline{v})]. \quad (39)$$

Substituting (39) into (5) and again applying the monotone convergence theorem completes the proof. \square

A.6. Proof of Proposition 3

PROOF OF PROPOSITION 3. Using the previous remark that $R_k = E[\mathcal{G}(\alpha)^{N(p_k)}](p_k + R_{k-1})$, we write $p_k = \arg \max_{p \in V} \mathcal{T}(p, R_{k-1})$, where $\mathcal{T}: V \times [0, +\infty) \rightarrow R$ is defined by

$$\mathcal{T}(p, y) = \frac{1 - F(p)}{1 - \mathcal{G}(\alpha)F(p)}(p + y).$$

Because \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 $\arg \max_{p \in V} \mathcal{T}(p, y)$ 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 . \square

A.7. Proof of Proposition 4

PROOF OF PROPOSITION 4. The existence of $\lim_{k \rightarrow +\infty} p_k$ is immediate as the sequence $(p_k)_{k \geq 1}$ is decreasing (from Proposition 3) and it is bounded from below (by \underline{v}). It follows from the last equation in (14) that for $k \geq 1$, $R_k = \sum_{i=1}^k j(p_i)$. However, 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

$$\begin{cases} N(p_k) \geq 1 \\ p_k \leq \bar{v} \end{cases} \text{ implies } R_k \leq \mathcal{G}(\alpha)(\bar{v} + R_{k-1});$$

this implies, in turn, $R_k \leq \mathcal{G}(\alpha)\bar{v}/(1 - \mathcal{G}(\alpha))$ 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 infinite sum with general term $j(p_k)$, this entails $\lim_{k \rightarrow +\infty} j(p_k) = 0$ and, by continuity, $j(p^*) = 0$. \square

Acknowledgments

The author thanks Garrett van Ryzin and two anonymous referees, as well as Stephen C. Graves, for many insightful comments and suggestions about an earlier version of this paper. He is grateful to Andreea Cepoiu for her help with the numerical experiments, and to the MIT Undergraduate Research Opportunities Program for its funding. Rakesh Vohra, Martin Lariviere, Paul Milgrom, and Donald Rosenfield suggested relevant references, while Jay Sethuraman, Leonid Kogan, Lawrence Wein, 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.

References

- Arnold, M., S. Lippman. 2001. The analytics of search with posted prices. *Econom. Theory* 17 447–466.
- Bergemann, D., S. Morris. 2001. Robust mechanism design. Working paper, Department of Economics, Yale University, New Haven, CT.
- Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control*, Vol. 1. Athena Scientific, Belmont, MA.

- Besanko, D., W. L. Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* **36**(5) 555–567.
- Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5** 203–229.
- Bulow, J., J. Roberts. 1989. The simple economics of optimal auctions. *J. Political Econom.* **97**(5) 1060–1090.
- Chow, Y. S., H. Robbins, D. Siegmund. 1971. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Company, Boston, MA.
- Chung, K.-S., J. C. Ely. 2002. Ex-post incentive compatible mechanism design. Working paper, Department of Economics, Northwestern University, Evanston, IL.
- Chung, K.-S., J. C. Ely. 2003. Foundations of dominant strategy mechanisms. Working paper, Department of Economics, Northwestern University, Evanston, IL.
- Das Varma, G. D., N. Vettas. 2001. Optimal dynamic pricing with inventories. *Econom. Lett.* **72**(3) 335–340.
- Elfving, G. 1967. A persistency problem connected with a point process. *J. Appl. Probab.* **4** 77–89.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* **49**(10) 1287–1309.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.
- Karlin, S. 1962. Stochastic models and optimal policy for selling an asset. K. J. Arrow, S. Karlin, H. Scarf, eds. *Studies in Applied Probability and Management Science*. Stanford University Press, Stanford, CA, 148–158.
- Kincaid, W. M., D. A. Darling. 1963. An inventory pricing problem. *J. Math. Anal. Appl.* **7** 183–208.
- Klemperer, P. 1999. Auction theory: A guide to the literature. *J. Econom. Surveys* **3**(13) 227–286.
- Laffont, J.-J., E. Maskin. 1982. *The Theory of Incentives: An Overview*, Chapter 2. Cambridge University Press, Cambridge, UK, 31–94.
- Lucking-Reiley, D. 2000. Auctions on the Internet: What's being auctioned, and how? *J. Indust. Econom.* **48**(3) 227–252.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.
- Maskin, E., J. Riley. 1984. Optimal auctions with risk averse buyers. *Econometrica* **52**(6) 1473–1518.
- Maskin, E., J. Riley. 1989. Optimal multi-unit auctions. F. Hahn, ed. *The Economics of Missing Markets, Information and Games*. Oxford University Press, New York, 312–335.
- McAfee, R. P., J. McMillan. 1988. Search mechanisms. *J. Econom. Theory* **44** 99–123.
- Mookherjee, D., S. Reichelstein. 1992. Dominant strategy implementation of Bayesian incentive compatible allocation rules. *J. Econom. Theory* **56** 378–399.
- Mortensen, D. T. 1986. Job search and labor market analysis. O. Ashenfelter, R. Layard, eds. *Handbook of Labor Economics*, Vol. 2. North-Holland, Amsterdam, The Netherlands, 849–919.
- Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* **6** 58–73.
- Riley, J., R. Zeckhauser. 1983. Optimal selling strategies: When to haggle, when to hold firm. *Quart. J. Econom.* **98**(2) 267–289.
- Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- Vulcano, G., G. van Ryzin, C. Maglaras. 2002. Optimal dynamic auctions for revenue management. *Management Sci.* **48**(11) 1388–1407.
- Wang, R. 1993. Auctions versus posted-price selling. *Amer. Econom. Rev.* **83**(4) 838–851.