

# Revenue Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand

Alex Gershkov, Benny Moldovanu and Philipp Strack\*

November 14, 2016

## Abstract

We show that appropriate dynamic pricing strategies can be used to draw benefits from the presence of consumers who strategically time their purchase even if the arrival process is not known. In our model, a seller sells a stock of objects to a stream of randomly arriving long-lived agents. Agents are privately informed about their values, and about their arrival time to the market. The seller needs to learn about future demand from past arrivals. We characterize the revenue maximizing direct mechanism. While the optimal mechanism cannot be reduced to posted prices (and requires personalized prices), we also present a simple, *“learn and then sell”* mechanism that is able to extract a large fraction of the maximal revenue. In this mechanism the seller first charges a relatively low price that allows learning about the arrival process, and in a second stage, the seller charges the optimal posted price given the previously obtained information

## 1 Introduction

Revenue Management - broadly speaking, the study of the dynamic allocation of capacity and its pricing under uncertain, fluctuating demand - has been pioneered on an industrial scale by airline companies in the

---

\*We are grateful to the area editor Serguei Netessine, the associate editor and three referees for helpful comments. A previous version of this paper has been circulated under the title “Demand Uncertainty and Dynamic Allocation with Strategic Agents”. We also wish to thank Ron Lavi, Konrad Mierendorff, Alessandro Pavan, Andy Postlewaite, Ron Siegel, Juuso Toikka, Rakesh Vohra, Mike Whinston and Asher Wolinsky for helpful remarks. Moldovanu is grateful to the European Research Council and to the German Science Foundation for financial support. Gershkov is grateful to the Google Inter-University Center for Electronic Markets and Auctions, to Israel Science Foundation and German-Israeli Foundation. Gershkov: Department of Economics and Center for the Study of Rationality, The Hebrew University of Jerusalem, and School of Economics, University of Surrey, alexg@huji.ac.il; Moldovanu: Department of Economics, University of Bonn, mold@uni-bonn.de; Strack: Department of Economics, University of California, Berkeley, pstrack@berkeley.edu.

mid 70's. These practices have rapidly spread to the allocation of fixed capacities in hotel booking, freight transportation, car rentals and holiday resorts, the retail of seasonal and style goods (e.g., groceries and apparel), electricity generation, e-business (online advertising and broadcasting, allocation of bandwidth), and event management (sports, concerts, etc...).

The RM techniques utilized in practice yield frequent price fluctuations, as prices depend on a multitude of constantly changing variables such as time to take-off or major event, remaining capacity, demand forecasts, weather and so on. In particular, prices significantly change following unexpected demand shocks, as airlines, for example, have the time and ability to revise their estimates about residual demand (see Escobari [2012]). As a consequence, customers have an incentive to strategize by carefully timing their purchase.

Until recently the execution of careful purchase timing strategies by consumers was encumbered by the fact that pricing algorithms are relatively opaque, and by the absence of reliable historical data: strategizing customers were basically playing a lottery. This dramatically changed with the advent of price comparison websites such as Bing/Travel and Kayak who offer free advice (based on huge amounts of data and on Artificial Intelligence algorithms) about the timing of purchase: customers for a certain flight are advised whether to buy immediately or wait, together with an estimate of the probability of saving money by waiting (See Etzioni [2003] for the scientific basis of this development). Similar websites offer advice about the timing of purchase for durable goods.

The conventional wisdom is that consumer strategizing may hurt revenues (see, for example, Mantin and Rubin [2013] who estimate a 3% revenue loss on routes where information from Bing/Travel is available, or Soysal and Krishnamurthi [2012] who estimate a significant 11% loss in the market for fashion goods). Li, Granados and Netessine [2014] empirically estimate the percentage of strategic customers in the airline industry (a significant fraction is found) and carefully expose the existing trade-off between lowering prices in the future and losing demand by keeping prices high. Interestingly, the effect of strategic consumers is ambiguous in their study<sup>1</sup>.

With strategic consumers revenue may be potentially harmed both by the shift in demand to lower prices, and by the indirect effect on the seller's ability to learn about residual demand. Since the underlying pricing techniques observed by the empirical studies were not necessarily optimized with respect to the strategic timing of purchases, it is important to show -as we do here -that the presence of strategic consumers can never decrease revenue if the optimal pricing technique that accounts for their presence is used.

In this paper we study a dynamic trading model where a designer allocates several units to a stream of randomly arriving, privately informed agents. Agents arrive according to a general Markov counting

---

<sup>1</sup>Strategic customers are found to have a beneficial effect in the group buying model of Marinési, Girotra, Netessine [2016].

process. This class is particularly important, since it includes the Poisson arrival process with unknown arrival rate (see e.g. Presman [1990], and Example 4.2 in the present paper). In fact, almost the entire literature deals solely with this case.

All agents are long lived, and each agent is privately informed about his value for an object, and about his arrival time to the market (thus, private information is two-dimensional). As the designer may not be aware of the nature of the arrival process, this naturally leads to correlation between arrival times from the designers point of view. A major feature of our model is that the designer learns about future demand from past arrivals (for example, an airline may learn both from past sales but also by gathering information about “waiting” customers on platforms such as Bing/Travel). In turn, agents strategically choose when to make themselves available for trade, and in doing so they realize that they influence the designer’s beliefs, and hence, potentially, their own terms of trade. Both agents and designer discount the future, and the designer wishes to maximize his expected revenue.

We first look at the benchmark case with observable arrivals: private information pertains only to values and is one-dimensional. As in a static auction, the seller’s revenue is maximized by a policy that maximizes the expected discounted sum of virtual values (Theorem 6).

If arrivals are unobservable, mechanisms need to take into account the second dimension of the agents private information, namely their arrival time. The main result here (Theorem 11) is based on a monotonicity property of the revenue maximizing allocation: agents who arrive earlier get the object earlier. We use this to show that, as long as the arrival process is Markov, the payment scheme that maximizes revenue under observable arrivals maximizes revenue even if arrivals are unobservable. Intuitively, early arrivals may be detrimental for the agents since, it can make the seller more optimistic about the future arrivals and may induce higher prices. Yet, it does not imply that the agents necessarily would like to postpone their arrivals, since later arrivals also increase the risk of sales to other agents. The Markov property implies that, in any incentive compatible mechanism, an agent who delays his arrival and arrives at time  $t$  gets exactly the same expected utility as an agent with the same value who truthfully arrives at  $t$ .

Agents also solve here an optimal stopping problem: when to reveal their presence to the mechanism designer? Because the designer’s decision is itself the solution to a sequence of optimal stopping problems, this allows us to show that, under the above payment scheme, it is optimal for every agent to announce his presence immediately. This characterization uses a specific and important physical property of the private information about the arrival time: an agent may misrepresent it only in one direction, i.e., by arriving or making himself available for trade later. Therefore, only one-directional deviations with respect to arrival time should be taken into account.

Board and Skrzypacz [2015] recently demonstrated the optimality of *posted price* mechanisms (followed by an auction at the deadline in case of a finite horizon) for a dynamic environment with long-lived agents that arrive according to a Poisson process with a known, fixed arrival rate. The main conclusion of the present paper is that the need to learn about future demand renders this class of mechanisms sub-optimal: revenue maximization in the environment with learning generally requires mechanisms where prices are more involved than just posted prices. In particular, the price charged to an agent depends on the entire history (including his own arrival to the market) and needs thus to be personalized.

Nevertheless, complex direct mechanisms are rarely used in practice. Therefore, we also present a very simple, indirect *two-phase* mechanism that is able to extract a large fraction of the maximal revenue even if the number of objects is not very large. In this mechanism the seller first charges a relatively low price that allows him to learn about the arrival process. In a second stage, the seller charges the optimal fixed posted price given the information obtained in the first stage.

Our analysis also answers the question whether the seller can increase revenue by withholding information from potential buyers. For example, shall an airline inform buyers about the remaining number of seats? Shall a fashion store inform buyers when only few items of the current collection are left in store to discourage waiting for the end-of-season sale?

The effects of informing customers go in different directions. When the remaining stock is low it reduces customers incentive to wait and thus increases the discounted per unit revenue of the seller. When the stock is high, the opposite happens. Yin, Aviv & Pazgal [2009] find that, with a known Poisson arrival process, two types of buyers and a seller who is restricted to two posted prices, hiding the number of remaining items from potential buyers can increase expected revenue by up to 20%. In contrast, we show in Theorem 15 that, if the seller is not restricted to in his choice of mechanism, hiding information is never beneficial.

Finally, we use the above results to show that the presence of long-lived agents who strategize over the timing of their purchases (assuming arrivals are unobservable) yields a higher revenue than the one optimally obtained in the situation where agents are short lived and must buy immediately upon arrival (Corollary 16). In other words, we show that appropriate revenue management techniques can be used to overcome and even draw benefits from the presence of consumers who strategically choose their purchase time. This is particularly important in environments where learning about demand is relevant since in those settings the advantages of RM techniques over simpler “naive” strategies such as fixed pricing or pricing without belief updates are most pronounced (see Aviv and Pazgal [2005] and Aviv, Levin and Nediak [2009] for excellent discussions of these issues).

## 1.1 Literature Review

A pioneering study on revenue management with myopic customers and known (but stochastic) demand process is Gallego and van Ryzin [1994]<sup>2</sup>. There arriving customers buy an homogenous product if the current price is below their willingness to pay and leave the market otherwise. Newer studies generalize some features of the workhorse model, e.g. Gershkov and Moldovanu [2009a] analyze revenue maximization and inventory planning with short-lived agents and with several heterogenous objects.

In practice, there are many instances where agents are long-lived and strategize over purchase time, and where the seller observes sales and needs to make an accurate demand forecasting in a complex, changing environment. The combination of these features are the focus of the present paper. The first attempts to incorporate a randomly arriving stream of non-myopic, long-lived agents appeared in the search literature, e.g., Zuckerman [1988], Zuckerman [1986], Stadje [1991], and Boshuizen and Gouweleeuw [1993]. There the planner is perfectly informed about the arrival process, the values and arrival times. Thus, pricing does not play a role, and the focus is on the allocation procedure. In a pricing model, Aviv and Pazgal [2008] quantify the losses incurred by assuming that agents are myopic when, instead, they are forward-looking. The ensuing RM literature displayed various ingenious methods to adjust the pricing techniques developed for myopic consumers in order to take into account strategic timing of purchases. Notable examples are: Gallego, Kou and Phillips [2008] (option on callable seats with possible reallocation), Elmaghraby et. al [2009] (selling with future reservations), Jerath, Nettekine, Verarghavan [2009] (opaque sales through an intermediary, so that consumers do not know which brand they get), Osdachyi and Vulcano [2010] (combination of temporal price discrimination with rationing at the low price), Su [2007] (exploitation of differences in patience among customers' patience), Marinesi, Girotra, Nettekine [2016] (threshold discounts for group purchasing) Chen and Farias [2015] (limits on the magnitude of price drops) and Borgs et al. [2014] (varying capacity in a finite horizon model with deterministic demand). It is fair to say that none of these papers attempts a full mechanism design analysis where optimality is analytically established with respect to a large class of pricing procedures.

Closer in spirit to our paper, Gallien [2006] analyzes revenue maximization in a model where the agents are long-lived, and where arrivals are private information. He focuses on a commonly known arrival process where sales always occur upon arrival. The resulting optimal pricing scheme is time-independent, thus strategizing in the time dimension and the ensuing learning - which are the focus of our paper - do not play any role<sup>3</sup>.

---

<sup>2</sup>As Talluri and van Ryzin [2004] note in their excellent book, this has been for many years the workhorse model

<sup>3</sup>In other words, Gallien's solution coincides with the one where arrivals are observable, and where agents are short lived (see

Another closely related paper is Board and Skrzypacz [2015]. They characterize revenue maximization in a model with patient agents where arrivals are described by a Poisson process with a known, fixed arrival rate. A main assumption is that current arrivals are independent of past ones. Therefore, learning from past arrivals about the future demand does not play a role in their model. Roughly speaking, allocative externality payments (modified to maximize virtual values instead of values) maximize revenue in their case. The main structural elements of the revenue maximizing mechanism are posted prices, which is also the case in the setting with short lived agents. In contrast, the implication of learning is that mechanisms that post prices do not maximize revenue. Similar to our Corollary 16, Board and Skrzypacz show that, in their environment, strategic customers increase the sellers revenue.<sup>4 5</sup>

The other main topic of our paper, the trade-offs between dynamic pricing and earning while learning about demand, has also been emphasized in the recent RM literature. For example, Bora Keskin and Zeevi [2014] and Besbes and Zeevi [2015] look at learning in dynamic pricing models with deterministic demand whose parameters are not known to the seller. Closer to our paper, Mason and Välimäki [2011] focus on posted-price mechanisms in a model with one object and with stochastic, unobservable arrivals of short-lived buyers. Aviv and Pazgal [2005] also consider myopic agents arriving according to a Poisson process whose rate follows a Gamma distribution. The arrival processes in both previously mentioned papers are special cases of the class considered here, but the strategic effects of delaying arrivals do not arise in their models because the agents - who can only be served upon arrival - cannot manipulate the designer's belief about the underlying demand. Araman and Caldentey [2009] allow for patient buyers and numerically evaluate the consequences<sup>6</sup>. Again, it is fair to say that no general mechanism design analysis was attempted yet.

A different type of dynamic model - leading to a completely different contracting environment - was pioneered by Courty and Li [2000]: it has all agents present at all periods. but their information arrives over time. Recent mechanism design analyses of revenue maximization in such frameworks include Kakade, Lobel and Nazerzadeh [2013] and Pavan, Segal and Toikka [2014]. While the setup analyzed of Kakade, Lobel and Nazerzadeh [2013] and Pavan, Segal and Toikka [2014] is quite general, it does not cover our Albright [1977] and Gershkov and Moldovanu [2009a]).

<sup>4</sup>Board and Skrzypacz's [2015] paper and the present article were written in parallel, and independently of each other.

<sup>5</sup>It is also instructive to compare their results to models with myopic buyers, where auctions take place at each period, e.g. as in Vulcano, van Ryzin and Magalarias [2002]. Besbes and Lobel [2012], Pai and Vohra [2008] and Mierendorff [2010a] analyze revenue maximization in models where arriving agents are privately informed about values, and about a deadline by which they need the object. The distribution of the number of arrivals in each period is known to the designer and there is no learning in these models.

<sup>6</sup>As in the paper by Mason and Valimäki, the arrival process is Poisson with two possible rates

case where arrival times are private information and customers report them strategically.<sup>7</sup> Moreover, in those papers the designer can not learn from past demand, about future demand since private information is independently distributed. In contrast, in our paper learning (i.e. correlation between arrivals) is a main feature of the analysis, and extending their type of analysis to our setup seems very hard.

Finally there is a related literature that looks at efficient dynamic mechanism design, e.g. Bergemann and Välimäki [2010], Cavallo, Parkes and Singh [2010], Parkes and Singh [2003], Said [2012], Athey and Segal [2013]. It is important to note, for example, that the independence assumptions in Bergemann and Välimäki do not hold in our model since arrivals are correlated and unobservable. Moreover, the mechanisms used by Athey and Segal are excluded here by the requirement that prices cannot depend on information arriving after the allocation time. In a companion paper, Gershkov et. al [2014], we find that it is often impossible to implement the welfare maximizing allocation if the process is not Markov.

## 2 The Model

A designer endowed with  $k \geq 1$  indivisible, identical objects faces a stream of randomly arriving agents in continuous time. The probability measure over agents' arrival times is generated by a time-inhomogeneous Markov counting process  $(\mathcal{N}(t))_{t \geq 0}$  where  $\mathcal{N}(t) \in \mathbb{N}$  is the random variable representing the number of arrivals up to time  $t$ . The arrival time of the  $i$ 's agent is then given by

$$a_i = \inf\{t \in \mathbb{R}_+ : \mathcal{N}(t) \geq i\}.$$

The Markov assumption implies that the probability measure over arrivals after time  $t$  only depends on  $t$  and on the number of agents  $\mathcal{N}(t)$  who arrived prior to time  $t$ , but not on when precisely they arrived. Moreover, we assume, that there exists a bounded rate function  $\lambda : \mathbb{R}_+ \times \mathbb{N} \rightarrow \mathbb{R}_+$  such that the conditional probability of  $m \geq 1$  agents arriving depends only on the time  $t$ , and on the number of prior arrivals:

$$\lim_{\Delta \searrow 0} \frac{1}{\Delta} \mathbb{P} \left[ \mathcal{N}(t + \Delta) = \mathcal{N}(t) + m \mid (\mathcal{N}(s))_{s \leq t} \right] = \mathbf{1}_{\{m=1\}} \lambda(t, \mathcal{N}(t)).$$

This assumption ensures that the probability that multiple agents arrive at the same time equals zero. The time horizon is potentially infinite, but the framework is rich enough to embed the finite horizon case by considering arrival processes where, after some time,  $T < \infty$  no more arrivals occur, i.e., where for any  $n$ ,

---

<sup>7</sup> It is in principle possible to embed our model in the informational structure studied by Kakade, Lobel and Nazerzadeh [2013] and by Pavan, Segal and Toikka [2014]: in their language, this would require each agent's valuation to jump up from zero at a single random point in time (arrival in our model) But, our information at time zero is two-dimensional (jump time and jump size) while these authors require the private information at time zero to be one-dimensional.

$t \geq T$  the arrival rate is zero  $\lambda(t, n) = 0$ . Since arrivals are described by Markov counting processes, the designer's beliefs about future arrivals may evolve over time, and may also depend on the number of past arrivals.

Each agent's private information is two-dimensional: the arrival time  $a_i \geq 0$  and the value  $v_i \geq 0$  he gets if allocated an object. In other words, we assume that the designer observes neither agents' arrivals nor their values. Each agent demands exactly one object. We denote by  $\tau_i$  the (random) time the designer allocates an object to the  $i$ -th agent and set  $\tau_i = \infty$  if agent  $i$  does not receive an object.

If agent  $i$  arrives at time  $a_i$ , gets the object at time  $\tau_i \geq a_i$  and pays  $p_i$  at time  $\tau' \in [a_i, \tau_i]$ , then his discounted utility is given by

$$e^{-r\tau_i}v_i - e^{-r\tau'}p_i$$

where  $r \geq 0$  is the discount factor. Note that this specification leads to the same preferences (and hence to the same behavior) as those obtained by assuming that an agent discounts the future relative to his arrival time, i.e.,

$$e^{-r(\tau_i - a_i)}v_i - e^{-r(\tau' - a_i)}p_i.$$

Consequently, the revenue maximizing mechanism for both specifications is the same.

The agents' values are represented by I.I.D. random variables  $v_i$  on the support  $[0, \bar{v}]$  where  $\bar{v} \leq \infty$ , with common c.d.f.  $F : [0, \bar{v}] \rightarrow [0, 1]$  and with continuous, strictly positive density  $f : [0, \bar{v}] \rightarrow \mathbb{R}_+$ . We assume that each  $v_i$  has a finite mean  $\mu \geq 0$ , and a finite variance. We make a standard assumption that the virtual valuation  $v - \frac{1-F(v)}{f(v)}$  is increasing in  $v$ . While we allow for arrival times to be correlated between agents, we assume that valuations are independent. We also assume that, for each agent, his arrival time is independent of his value.<sup>8</sup>

The designer maximizes his expected discounted revenue. That is, a payment of  $p$  at time  $\tau'$  generates utility of  $e^{-r\tau'}p$  to the designer. Potentially, every agent may pay more than once. However, because the agent has the same discount factor as the seller, if some agent pays more than once, there exists another mechanism that generates the same expected discounted payment and utilities where every agent pays only once. If we denote by  $\tau'_i$  the time agent  $i$  makes a payment of  $p_i$  the discounted revenue of principal is given by

$$\sum_{i=1}^{\infty} e^{-r\tau'_i}p_i.$$

---

<sup>8</sup>Independence allows us to focus on learning from arrivals, as opposed to learning from values. Assuming dependence between the arrival times and values generates correlation among the values. This creates distinct, additional complications, as explored in Gershkov and Moldovanu [2009b] and Gershkov and Moldovanu [2012].



## 2.1 Direct Revelation Mechanisms

We look for the revenue maximizing procedure within a very general class of mechanisms. The designer can commit to any game played by the agents in order to determine who receives an object and at which price. Commitment is a crucial assumption. Special cases are auctions at predetermined times and posted price mechanisms.

The Revelation Principle for dynamic environments (see Myerson [1986]) implies that we can restrict attention to direct mechanisms where each agent reports his value and arrival time, and where the mechanism specifies a probability distribution over times when each agent gets an object, and a payment. Without loss of generality, we can restrict attention to mechanisms where each agent reports his type upon arrival, e.g., the time of the report coincides with the arrival time. This holds since the equilibrium outcome of any mechanism where at least one agent reports his type after his arrival, can be replicated by another mechanism and equilibrium where all agents reports their types upon arrival. Another implication of the revelation principle to the dynamic environments is that, in the direct mechanisms, no information should be revealed about the previous arrivals and reports. Intuitively, minimizing the information revealed to each agent reduces the available contingent deviations from truthtelling, and relaxes the incentive compatibility constraints for that agent. Therefore, each agent  $i$  choose a time  $\tilde{a}_i$  at which to report his arrival to the mechanism and a message to sent to the mechanism  $m_i \in [0, \bar{v}]$ .

A direct mechanism  $(\tau, p, \tau')$  specifies the allocation times  $(\tau_i(\tilde{a}_i, m_i))_{i \in \mathbb{N}}$ , payments  $(p_i(\tilde{a}_i, m_i))_{i \in \mathbb{N}}$  and the payment times  $(\tau'_i(\tilde{a}_i, m_i))_{i \in \mathbb{N}}$  as a function of times of reports and messages sent. Since in the optimal policy a sale can occur later than an actual arrival, the allocation and payment to an agent can also be conditioned on information that accrues between the arrival time of that agent and the allocation time. Denote by

$$h(t) = ((\tilde{a}_j, m_j) : \tilde{a}_j \leq t)$$

the history of reports up to time  $t$  : this consists of the reported values  $m_j$  and the times of the reports  $\tilde{a}_j$ . Finally, we denote by  $\mathbb{E}[\cdot \mid h(t)]$  the probability measure conditional on the history  $h(t)$ .

A vector  $(\tau_i)_{i \in \mathbb{N}}$  of allocation times is *feasible* if no agent receives an object before he reported his arrival,  $\tilde{a}_i \leq \tau_i(\tilde{a}, m)$ , or pays before his arrival,  $\tilde{a}_i \leq \tau'_i(\tilde{a}, m)$ , and if at most  $k$  objects are allocated  $\sum_{i=1}^{\infty} 1_{\{\tau_i(\tilde{a}, m) < \infty\}} \leq k$ . Furthermore, as objects can not be reallocated  $\tau_i(\tilde{a}, m)$  must be constant in all  $\tilde{a}_j, m_j$  such that  $\tilde{a}_j > \tau_i$ .

**Assumption 1 (Individual Rationality)** *We restrict attention to ex-post individual rational mechanisms, where, after every history, the equilibrium utilities of all agents are non-negative.*

Assumption 1 ensures that no agent makes a payment larger than his discounted utility from consuming the good. This rules out selling lotteries where the agent makes an up-front payment and randomly receives the good later. Furthermore, Assumption 1 excludes random payment schemes a la Cremer and McLean [1988], where agents guess other agents' reports - this could, in principle, be used to extract information about correlated arrival times at no cost.

Since the agents and the designer have the same discount factor, we can assume without loss of generality that the payment times coincide with allocation times: a *direct* mechanism  $(\tau, p)$  specifies then allocation times  $(\tau_i(\tilde{a}_i, m_i))_{i \in \mathbb{N}}$  and payments  $(p_i(\tilde{a}_i, m_i))_{i \in \mathbb{N}}$  as a function of reported arrivals and messages sent.

## 2.2 Incentive Compatibility

We assume that no agent can report his arrival before he truly arrives  $\tilde{a}_i \geq a_i$ .<sup>9</sup> A strategy  $\tilde{a}_i, m_i$  for agent  $i$  is a mapping from his private type  $(a, v_i)$  into a stopping/reporting time  $\tilde{a}_i$  and a value  $m_i$ . Denote by  $\mathbb{E}_{i,t}[\cdot]$  agent  $i$ 's expectations conditional on his private information  $(a_i, v_i)$ , and assume that all other agent's report their arrival immediately,  $\tilde{a}_j = a_j$ , and report their values truthfully,  $m_j = v_j$ .

*Bayesian incentive compatibility* requires that it is optimal for every agent to report his arrival immediately, and also to report truthfully his value. That is, for any  $i$ , it holds:

$$(a_i, v_i) \in \arg \max_{\tilde{a}_i \geq a_i, m_i} \mathbb{E}_i \left[ e^{-r\tau'_i(\tilde{a}, m)} v_i - e^{-r\tau'_i(\tilde{a}, m)} p_i(\tilde{a}, m) \right].$$

## 3 Observable Arrivals

In this section we analyze the designer's problem under the simplifying assumption that he observes all arrivals. Thus, the agents' private information only pertains to values. However, we still assume that agents do not observe the arrivals and reports of previously arrived agents. Our first result shows that, although hiding information from buyers can facilitate the implementation of certain allocations, it does not affect the expected revenue from allocations that are implementable via other information structures. That is, if some allocation is incentive compatible even if some information is revealed, then it generates the same expected revenue as the mechanism that implements the same allocation and reveals no information<sup>10</sup>.

<sup>9</sup>Although the so called "revelation principle" need not hold in settings where some deviations from truth-telling are unfeasible for certain types, this principle does hold for our case of unilateral deviations in the time dimension - see Theorem 1 and Example 5.a.2 in Green and Laffont [1986].

<sup>10</sup>While our original definition of incentive compatibility corresponds to the environment where every agent only knows his arrival time and value, the definition straightforwardly extends to the present environment: it requires that, for any observed signal about the past history, it is optimal for every agent to report his arrival immediately and his value truthfully.

**Theorem 1 (Generalized Revenue Equivalence)** *Suppose that arrivals are observable to the principal, and that, at the time of his arrival  $a_i$ , each agent  $i$  observes his valuation  $v_i$  and a signal  $s_i$  which is (weakly) less informative than observing the entire prior history of arrivals and reported values  $\{(\tilde{a}_j)_{j \leq \mathcal{N}(a_i)}, (m_j)_{j \leq \mathcal{N}(a_i), j \neq i}\}$ .<sup>11</sup> Furthermore, assume that this signal is also observable to the principal. Then, the expected revenue of the principal in an incentive compatible mechanism that implements the allocation  $\tau = (\tau_i)_{i \in N}$  such that every agent with a valuation of zero gets a utility of zero equals*

$$\mathbb{E} \left[ \sum_{i \in N} e^{-r\tau_i(a, v, s)} J(v_i) \right]$$

where  $J(v_i) := v_i - \frac{1-F(v_i)}{f(v_i)}$ .

**Proof.** As the signal  $s_i$  is observable to the principal, he can potentially condition the allocation  $\tau_i$  and the payment  $P$  on it. To reflect this dependence in our notation we write  $\tau_i(v, a, s)$ . By the Envelope Theorem, the expected payoff of agent  $i$  in any incentive compatible mechanism when arriving at time  $a_i$  with a value of  $v_i$  and a signal  $s_i$  is given by

$$\mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i(a, (z, v_{-i}), s)} dz \mid s_i, v_i, a_i \right] = \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i(a, (z, v_{-i}), s)} dz \mid s_i, a_i \right],$$

where the above equality follows as future arrivals and valuations are independent of all past valuations by assumption. It follows from the law of iterated expectations that the ex-ante (i.e., before seeing his signal and arrival time) expected payoff (or information rent) to agent  $i$  is given by

$$\mathbb{E} \left[ \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i(a, (z, v_{-i}), s)} dz \mid s_i, a_i \right] \right] = \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i(a, (z, v_{-i}), s)} dz \right].$$

The outer expectation on the left-hand side is over the signals and arrival times the agent could observe, while the inner expectation is over arrival times and valuations of all agents. The expectation on the right-hand side is over valuations, arrivals and signals of all agents, and for the rest of the proof we shall only use this expectation. The expected utility agent  $i$  assigns to the mechanism before knowing his arrival time,

---

<sup>11</sup>In particular, this implies that the signal is independent of  $v_i$ , the valuation of player  $i$ ,  $v_i$ . Moreover, we assume that the signal  $s_i$  also contains information on  $i$ 's own arrival  $a_i$ .

valuation or signal is thus given by

$$\begin{aligned}
\mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i(a,(z,v_{-i}),s)} dz \right] &= \int_0^{\bar{v}} \mathbb{E} \left[ \int_0^x e^{-r\tau_i(a,(z,v_{-i}),s)} dz \right] f(x) dx \\
&= \mathbb{E} \left[ \int_0^{\bar{v}} f(x) \int_0^x e^{-r\tau_i(a,(z,v_{-i}),s)} dz dx \right] \\
&= \mathbb{E} \left[ \int_0^{\bar{v}} (1 - F(x)) e^{-r\tau_i(a,(x,v_{-i}),s)} dx \right] \\
&= \mathbb{E} \left[ \frac{1 - F(v_i)}{f(v_i)} e^{-r\tau_i(a,v,s)} \right].
\end{aligned}$$

At the last step we used integration by parts. Subtracting the expected value of the agents from the total surplus generated by the mechanism  $\sum_i e^{-r\tau_i(a,v,s)} v_i$ , we obtain that the expected revenue in any incentive compatible mechanism equals

$$\mathbb{E} \left[ \sum_i e^{-r\tau_i(a,v,s)} \left( v_i - \frac{1 - F(v_i)}{f(v_i)} \right) \right].$$

■

Define  $W_k(t)$  as the highest expected discounted sum of virtual values  $J(v_i)$  the designer can achieve by optimally choosing allocation times  $(\tau_1, \tau_2, \dots)$  and such that, after time  $t$ , it allocate at most  $k$  objects to the agents who did not receive an object prior to time  $t$  :

$$W_k(h(t)) = \sup_{\tau, \tau_i \geq t} \mathbb{E} \left[ \sum_i e^{-r(\tau_i - t)} J(v_i) \mid h(t) \right].$$

As the arrival rate function  $\lambda(n, t)$  is uniformly bounded by a constant  $\bar{\lambda} < \infty$ , and as the designer can always ignore arrivals, an upper bound is what the designer can achieve when the arrival rate is constant and equal to  $\bar{\lambda}$ . In this case, the revenue of the designer and hence  $W_k$  is bounded by the revenue which is generated if the designer has infinitely many objects and if he gives each agent an object at a price equal to that agent's value

$$W_k(h(t)) \leq \mathbb{E} \left[ \sum_i e^{-r a_i} v_i \right] = \mathbb{E}[v_i] \sum_{i=1}^{\infty} \left( \frac{\lambda}{\lambda + r} \right)^i = \frac{\lambda \mu}{r}.$$

As  $W_k(h(t))$  is bounded from above, it is well defined for all  $k, h(t)$ . Note that  $W_k$  constitutes an upper bound on the revenue that is achievable in any incentive compatible mechanism with  $k$  objects to allocate.

Let  $\kappa(t)$  be the number of remaining objects the designer can allocate  $\kappa(t) := k - \sum_i 1_{\{\tau_i < t\}}$ . We define the policy  $\tau^\diamond$  which allocates an object to agent  $i$  at the earliest time at which it is optimal to do so<sup>12</sup>

$$\tau_i^\diamond = \inf \{ t : J(v_i) \geq W_{\kappa(t)}(h(t)) - W_{\kappa(t)-1}(h(t)) \text{ and } \kappa(t) \geq 1 \}.$$

<sup>12</sup>If there are two agents for whom an allocation is optimal at the same time, the one with the higher value gets allocated the object first. We follow the convention that  $\inf \emptyset = \infty$ .

By definition,  $\tau^\diamond$  is a stopping time and it allocates at most  $k$  objects. Hence,  $\tau^\diamond$  is a feasible allocation time that maximizes the expected discounted sum of virtual valuations:

$$\tau^\diamond \in \arg \max_{\tau \in \mathcal{T}} \mathbb{E} \left[ \sum_{i \in \mathbb{N}} e^{-r\tau_i} J(v_i) \right].$$

As  $\tau^\diamond$  depends only on the realized sequence of arrival times  $a = (a_1, a_2, \dots)$  and valuations  $v = (v_1, v_2, \dots)$  we use the notation  $\tau^\diamond(a, v)$ .

From the previous Theorem, it follows that if policy  $\tau^\diamond$  is implementable (that is, if there exists a payment scheme  $p$ , such that the mechanism  $(\tau^\diamond, p)$  is incentive compatible) when arrivals are observable, then this mechanism is optimal. Similarly to static revenue equivalence, this result emphasizes the underlining allocation rule, while the pricing scheme plays only a secondary role. It also illustrates that the revealed information is of secondary importance as long as the underlying allocation is implementable. In other words, unless the allocation changes, the seller cannot increase his expected utility by hiding/ misreporting his private information.

We first characterize the implementable allocation rules if all the information regarding the reports of the previous agents is revealed. We then show that  $\tau^\diamond(a, v)$  is implementable. Similarly to the static environment, implementability hinges on a notion of monotonicity: for any history  $h(a_i)$  preceding a given, reported arrival of agent  $i$ , increasing the reported valuation  $v_i$  decreases the expected discounted time for obtaining an object  $\mathbb{E}[e^{-r\tau_i} \mid h(a_i)]$ . Based on this observation, our main result in this Section shows that the allocation  $\tau^\diamond(a, v)$  is indeed revenue maximizing. In addition to implementing  $\tau^\diamond(a, v)$ , the associated payment scheme generates zero utility to an agent with zero valuation.

**Definition 2** Consider two histories  $h(a_i)$  and  $h'(a_i)$  that differ only in the valuation of agent  $i$ , such that  $v'_i$ , the valuation of  $i$  in  $h'(a_i)$ , is larger than  $v_i$ , the valuation in  $h(a_i)$ . An allocation rule  $\tau$  is monotone with respect to valuations if for any agent  $i$ , and for any two such histories it holds that

$$\mathbb{E} \left[ e^{-r\tau_i(a, (v'_i, v_{-i}))} \mid h'(a_i) \right] \geq \mathbb{E} \left[ e^{-r\tau_i(a, (v_i, v_{-i}))} \mid h(a_i) \right].$$

**Theorem 3 (Monotone Allocations are Implementable)** Assume that all the information regarding the reports of the agents is revealed. Assume further that arrivals are observable. If the allocation  $\tau$  is implementable, then it is monotone with respect to valuations. Conversely, if the allocation  $\tau$  is monotone with respect to valuations, then  $\tau$  can be implemented using a payment paid at the time of allocation  $\tau_i$ :

$$P_i(a, v) = \mathbf{1}_{\{\tau_i(a, v) < \infty\}} \frac{\mathbb{E} \left[ \int_0^{v_i} (e^{-r\tau_i(a, v)} - e^{-r\tau_i(a, (z, v_{-i}))}) dz \mid h(a_i) \right]}{\mathbb{E} \left[ e^{-r\tau_i(a, v)} \mid h(a_i) \right]}. \quad (1)$$

Although the time when agent  $i$  gets the object depends on the history up to this time (in particular, it depends on the arrival times and values of the agents that arrive after the arrival of  $i$ ), his payment only depends only on the information available up to his arrival.

Therefore, to show that the allocation rule  $\tau^\diamond(a, v)$  is implementable, we have to show that it is monotone with respect to valuations. Instead of directly showing the monotonicity of  $\tau^\diamond(a, v)$ , we will use the monotonicity of another, closely related allocation rule, and the assumption that the virtual value function  $J(v)$  is increasing in  $v$ . Denote by  $\tau^* \in T$  the *welfare maximizing* policy under complete information

$$\tau^* \in \arg \max_{\tau \in T} \mathbb{E} \left[ \sum_{j \in \mathbb{N}} e^{-r\tau_j} v_j \right].$$

By the same argument which we gave for virtual valuations, the realized optimal allocation depends only on the realized sequence of arrival times  $a = (a_1, a_2, \dots)$  and of valuations  $v = (v_1, v_2, \dots)$  and we use the notation  $\tau^*(a, v)$ .

The next Proposition shows that a payment equal to the expected allocative externality (conditional on the information available to the agent at arrival), divided by the expected discounted allocation time implements the efficient (welfare maximizing) allocation. Despite the possible correlation in arrival times, which implicitly determine whether the value for the object at a certain period is positive or not in the formulation of Bergemann and Välimäki [2010], their *dynamic pivot* mechanism implements the efficient allocation in the case of observable arrivals since, conditional on the observable arrivals, the agents' values are independent.

**Proposition 4 (Pivotal Payment)** *The payment scheme*

$$P_i(a, v) = \mathbf{1}_{\{\tau_i(a, v) < \infty\}} \frac{\mathbb{E} \left[ \sum_{j \neq i} (e^{-r\tau_j^*(a, (0, v_{-i}))} - e^{-r\tau_j^*(a, v)}) v_j \mid h(a_i) \right]}{\mathbb{E} \left[ e^{-r\tau_i^*(a, v)} \mid h(a_i) \right]} \quad (2)$$

*implements the efficient dynamic allocation policy  $\tau^*$ . The resulting mechanism is ex-post individually rational.*

**Corollary 5 (Monotonicity in Valuations)** *The stopping time that maximizes the expected discounted sum of virtual valuations  $\tau^\diamond$  is monotone in valuations.*

Finally, we now have all the necessary tools in order to prove the main result of this section:

**Theorem 6 (Revenue Maximizing Policy)** *When arrivals are observable, the policy  $\tau^\diamond$  is implementable using the following payments, which are charged upon allocation:*

$$P_i(a, v) = \mathbf{1}_{\{\tau_i^\diamond(a, v) < \infty\}} \frac{\mathbb{E} \left[ \int_0^{v_i} (e^{-r\tau_i^\diamond(a, v)} - e^{-r\tau_i^\diamond(a, (z, v_{-i}))}) dz \mid h(a_i) \right]}{\mathbb{E} \left[ e^{-r\tau_i^\diamond(a, v)} \mid h(a_i) \right]}. \quad (3)$$

Consequently, this mechanism is revenue maximizing.

**Proof.** By Corollary 5, the policy  $\tau^\diamond$  is monotone with respect to valuations, and thus implementable by Theorem 3. By the definition of  $\tau^\diamond$  as the virtual valuation maximizing allocation, and by Theorem 1, the revenue in any other implementable mechanism is lower than the revenue in this mechanism. ■

## 4 Unobservable Arrivals

### 4.1 The Revenue Maximizing Direct Mechanism

Is the revenue-maximizing allocation derived in Theorem 6 implementable also in case where arrivals are unobservable? Proving that truthful reporting of arrival times is incentive compatible under the payment scheme of Theorem 6 is difficult because no explicit solution  $\tau^\diamond$  to the designer's allocation problem when he observes valuations and arrivals is known. Hence, it is not possible to directly verify that, under the proposed optimal stopping time  $\tau^\diamond$ , agents have incentives to truthfully report their arrival times. We overcome this difficulty by connecting the incentives of an agent to report a later arrival with the designer's incentives to delay the allocation of the object. Note that in the direct mechanism an agent who decides when to report his arrival also solves an optimal stopping problem about his arrival time!

**Definition 7 (Monotonicity in Arrivals)** *A vector of deterministic allocation times  $(\tau_1, \tau_2, \dots)$  is monotone in arrival times if and only if agents who arrive earlier get the object earlier, i.e. for all  $i$ ,  $a_i < \tilde{a}_i$  and all  $a_{-i}, v \in \mathbb{R}_+^\infty$*

$$\tau_i((a_i, a_{-i}), v) \leq \tau_i((\tilde{a}_i, a_{-i}), v).$$

We will show that monotonicity in arrivals is sufficient to ensure that incentive compatibility of a mechanism with observable arrivals will imply incentive compatibility of the same mechanism with unobservable arrivals. We first characterize the agents' utilities in all mechanisms that are incentive compatible with observable arrivals.

**Proposition 8** *Assume that the payment  $\mathbf{P}$  implements  $\tau$  under observable arrivals. The utility of agent  $i$  equals*

$$U_i(h(a_i)) = \mathbb{E} \left[ \int_0^v e^{-r \tau_i(a, (z, v_{-i}))} dz \mid h(a_i) \right] + U_i((a_1, \dots, a_i), (v_1, \dots, v_{i-1}, 0)).$$

**Lemma 9** *Fix an arbitrary mechanism  $(\tau, \mathbf{P})$ . An agent  $i$  who misreports his arrival  $\tilde{a}_i > a_i$  obtains the same expected utility as if he would have arrived at time  $\tilde{a}_i$  and used the same strategy to report his value.*

An immediate consequence of Lemma 9 is that if agent  $i$  misreports his arrival it is still optimal for him to truthfully report his value  $v_i$ .

**Corollary 10** *Assume that the payment  $\mathbf{P}$  implements  $\tau$  under observable arrivals. It is optimal for every agent  $i$  to report his value  $v_i$  truthfully even when he misreports his arrival.*

The next theorem uses the previous results to show that if a mechanism is incentive compatible with observable arrivals it is also incentive compatible with unobservable arrivals as long as 1) agents who do value the object at zero obtains an equilibrium utility of 0; and 2) agents who arrive earlier receive the object earlier.

**Theorem 11** *Consider a vector of allocation times  $(\tau_1, \tau_2, \dots)$  and a vector of payments  $\mathbf{P}$  such that:*

- 1 *the payment  $\mathbf{P}$  implements  $\tau$  under observable arrivals*
- 2  *$\tau$  is monotone in arrival times*
- 3  *$(\tau, \mathbf{P})$  yields zero utility for every agent with zero value.*

*Then  $\mathbf{P}$  implements  $\tau$  under unobservable arrivals.*

**Proof.** By Corollary 10 the agent will report his value truthfully in any mechanism that is incentive compatible under observable arrivals. Hence, we restrict attention to misreports of arrival times. Let us denote by  $\tilde{h} = ((\tilde{a}_i, a_{-i}), v)$  the history of the game where agent  $i$  arrives at time  $\tilde{a}_i$ . By (3) the designer does not make any transfer to the lowest type, and thus - by Proposition 8 - the utility of agent  $i$  if she arrives at time  $\tilde{a}_i$  equals

$$U_i(\tilde{h}(\tilde{a}_i)) = \mathbb{E} \left[ \int_0^v e^{-r\tau_i((\tilde{a}_i, a_{-i}), (z, v_{-i}))} dz \mid \tilde{h}(\tilde{a}_i) \right].$$

By Lemma 9 the equilibrium utility of agent  $i$  who arrived at time  $\tilde{a}_i$  equals the equilibrium utility of an agent who arrived at time  $a_i$ , but misreported his arrival time to be  $\tilde{a}_i$ . Hence, the expected utility of agent  $i$  when reporting her arrival to be  $\tilde{a}_i$  is given by

$$\begin{aligned} \mathbb{E} \left[ U_i(\tilde{h}(\tilde{a}_i)) \mid h(a_i) \right] &= \mathbb{E} \left[ \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i((\tilde{a}_i, a_{-i}), (z, v_{-i}))} dz \mid \tilde{h}(\tilde{a}_i) \right] \mid h(a_i) \right] \\ &= \mathbb{E} \left[ \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i((\tilde{a}_i, a_{-i}), (z, v_{-i}))} dz \mid h(\tilde{a}_i) \right] \mid h(a_i) \right] \\ &= \mathbb{E} \left[ \int_0^{v_i} e^{-r\tau_i((\tilde{a}_i, a_{-i}), (z, v_{-i}))} dz \mid h(a_i) \right]. \end{aligned} \quad (4)$$



The first step follows since, by the Markov property, the probability measure only depends on the number of arrivals prior to time  $\tilde{a}_i$ , which is the same in the history  $\tilde{h}(\tilde{a}_i)$  and  $h(\tilde{a}_i)$ . The last step is the law of iterated expectations.

As  $\tau_i((\tilde{a}_i, a_{-i}, (z, v_{-i})))$  increases in  $\tilde{a}_i$  by assumption (2) the utility of agent  $i$  given in (4) decreases in the reported arrival time  $\tilde{a}_i$ . As agent  $i$  can not report his arrival before he arrives, i.e.  $\tilde{a}_i \geq a_i$  it follows that it is optimal to report the arrival immediately, i.e.  $\tilde{a}_i = a_i$ . ■

Theorem 11 can be now used to directly prove that the revenue maximizing allocation  $\tau^\diamond$  is implementable even if arrivals are unobservable.

**Theorem 12** *The payment defined in Theorem 6 implements policy  $\tau^\diamond$  under unobservable arrivals, and this mechanism is revenue maximizing.*

**Proof.** The payments defined in Theorem 3 for allocation policy  $\tau^\diamond$  leave an agent with a valuation of zero with utility of zero. Thus, it only remains to prove that the optimal policy is monotone with respect to arrivals. Recall that the probability measure over future arrivals depends only on the number of past arrivals. Hence, when agent  $i$  arrives later at time  $\tilde{a}_i > a_i$ , the principal uses the same continuation strategy as he would have used conditional on not allocating the object to agent  $i$  in  $[a_i, \tilde{a}_i)$ , i.e.  $\tau_i^\diamond(a, v) \geq \tilde{a}_i$  implies that  $\tau_i^\diamond((a_i, a_{-i}), v) = \tau_i^\diamond((\tilde{a}_i, a_{-i}), v)$ . Consequently we have that,

$$\tau_i^\diamond(a, v) \leq \max\{\tilde{a}_i, \tau_i^\diamond(a, v)\} = \max\{\tilde{a}_i, \tau_i^\diamond((\tilde{a}_i, a_{-i}), v)\} = \tau_i^\diamond((\tilde{a}_i, a_{-i}), v)$$

where the last step follows since no agent can get an object before arriving. ■

The agent can manipulate the designer's beliefs about the latter's continuation value only in a very specific way: as soon as the agent reports his arrival, the designer knows his own true continuation value since this only depends on the information that the agent arrived, but not on the precise arrival time. The agent can still manipulate the price by changing the reported arrival time. For example, the designer's continuation value may be decreasing over time if no agent arrives. Thus, an agent who announces his presence later will pay a lower price (conditional on no other agent arriving in between). But, delaying arrival also entails the risk of another agent with a higher valuation arriving beforehand. Roughly speaking, our above result shows that the risk of not getting the object offsets the benefits of lower future prices.

In the static case, the expected utility in any incentive compatible mechanism is specified by the allocation (up to a constant) while prices are set, for a given allocation, by the incentive compatibility constraints. In the dynamic setting, the discounted stopping times replace the static probabilities of getting the object. Therefore, monotonicity of the allocation with respect to values yields monotonicity of the expected utility,

while prices are determined by incentive compatibility with respect to values. The agent who reports type  $(a', v)$  where  $v$  is his true value, gets the utility of type  $(a', v)$  independently of his true arrival time  $a \leq a'$  (up to the effect of discounting), since, in both cases, the designer will have the same beliefs about the arrivals posterior to time  $a'$ . Monotonicity of the allocation with respect to arrivals together with incentive compatibility with respect to values imply that a later arrival necessarily decreases the expected utility. In addition, discounting makes such a deviation even less profitable.

Finally note that, in our revenue maximizing mechanism, the payment  $P_i$  made by agent  $i$  only depends on information accruing before the time he gets allocated an object,  $\tau_i$ . In particular, our mechanism automatically satisfies the *exit condition* made in Bergemann and Välimäki [2010] whose consequence here would be that agents must leave the mechanism as soon as they obtain an object

## 4.2 Example: Learning about the Arrival Rate of a Poisson Process

Let  $(\mathcal{N}(t))_{t \in \mathbb{R}_+}$  be a Poisson process with unknown arrival rate  $\kappa \in \{l, h\} \subset \mathbb{R}_+$ . Define the posterior arrival rate process

$$\lambda(t) = \mathbb{E}[\kappa \mid h(t)],$$

and note that the posterior expected arrival rate  $\lambda(t)$  only depends on the time  $t$  and on the number of arrivals before  $t$ ,  $\mathcal{N}(t)$ , i.e.  $\lambda(t) = \lambda(t, \mathcal{N}(t))$ . More precisely, by Bayes rule, we have that

$$\begin{aligned} \lambda(t, n) &= l + (h - l) \mathbb{P}[\kappa = h \mid \mathcal{N}_t = n] \\ &= l + (h - l) \frac{\mathbb{P}[\kappa = h] \mathbb{P}[\mathcal{N}_t = n \mid \lambda = h]}{\mathbb{P}[\kappa = h] \mathbb{P}[\mathcal{N}_t = n \mid \kappa = h] + \mathbb{P}[\kappa = l] \mathbb{P}[\mathcal{N}_t = n \mid \kappa = l]} \\ &= l + \frac{\frac{\lambda(0)-l}{h-l} \exp(-ht) \frac{(ht)^n}{n!} (h-l)}{\frac{\lambda(0)-l}{h-l} \exp(-ht) \frac{(ht)^n}{n!} + \frac{h-\lambda(0)}{h-l} \exp(-lt) \frac{(lt)^n}{n!}} \\ &= l + \frac{h-l}{1 + \frac{h-\lambda(0)}{\lambda(0)-l} \exp((h-l)t - n \log(\frac{h}{l}))}. \end{aligned}$$

We can easily introduce a deadline  $\bar{T} \in \mathbb{R}_+$  after which the designer cannot allocate the good: this is done by simply setting  $\lambda(t, n) = 0$  for all  $t \geq \bar{T}$  and all  $n \in \mathbb{N}$ . As no agent arrives after time  $\bar{T}$ , and as the designer discounts the future, it will never be optimal to allocate an object after the deadline.

The expected posterior arrival rate continuously decreases if there is no arrival, and it jumps up at the time of each arrival. As a consequence, the optimal allocation policy with learning is very different from the optimal policy without learning. To see this, consider, for example, the case of a single object. Without learning the optimal policy is given by a constant cutoff (i.e., only types with values above a constant get the object), together with a fire-sale auction at the deadline (if any) – see Board and Skrypacz [2015]. In our




Figure 1: Numerical solution of the revenue maximizing cutoffs as a function of time for valuations uniformly distributed on  $[0, 1]$ . Depicted here are the cutoffs for zero to five arrivals in the single object case (red) and the the two object case (blue). The prior assigns probability only to the arrival rates 2 and 10. The initial expected arrival rate equals 4. Higher cutoffs correspond to more arrivals. The deadline equals  $\bar{T} = 1$  and the exponential discount rate is given by  $r = 0.095$ . At the deadline  $\bar{T}$  there is an auction with reserve price equal to 0.5 and the agent, who arrived with the highest valuation above 0.5 gets the object.

case with learning, the value obtained per object by a fixed continuation strategy is higher when the arrival rate is higher, and thus the optimal policy allocates the object to agents with higher values if the arrival rate is higher. Thus, as only higher types acquire the object, both the optimal allocation cut-off and price jump up after every arrival (for an illustration see Figure 1). In addition, when there are several objects for sale, both cutoffs and prices jump after each sale since supply becomes smaller.

Ignoring the opportunity of learning may cause a significant loss in revenue: For example in the setup of Figure 1 setting the optimal allocation constant cut-off of 0.854 (i.e., optimal for the time zero expected arrival rate of 4) yields an expected revenue of 0.577. This represents a loss of approximately 19% compared to the optimal policy that generates here a revenue of 0.711.

### 4.3 Changing Distribution of Values

We now extend our main result to a setting where the agents' distribution of values is allowed to change over time<sup>13</sup>. Here, the values of agents arriving at time  $t$  are represented by I.I.D. random variables  $v_{ti}$  on the support  $[0, \bar{v}_t]$  where  $\bar{v}_t \leq \infty$ , with common c.d.f.  $F_t : [0, \bar{v}_t] \rightarrow [0, 1]$  and with continuous, strictly positive density  $f_t : [0, \bar{v}_t] \rightarrow \mathbb{R}_+$ .<sup>14</sup>

**Assumption 2** *The distributions  $F_t$  are ordered by the hazard rate order, i.e., if  $t_1 \leq t_2$  then*

$$\frac{f_{t_1}(v)}{1 - F_{t_1}(v)} \geq \frac{f_{t_2}(v)}{1 - F_{t_2}(v)} \text{ for all } v.$$

The hazard rate order is a widely used stochastic order that is stronger than first order stochastic dominance. Intuitively, the assumption says that later arrivals are likely to have higher values. This scenario seems to fit well, for example, the market for airline seats where business travelers tend to arrive nearer to the deadline.

**Proposition 13** *Assume that Assumption 2 holds. Denote by  $\hat{\tau}(a, v)$  the policy maximizing the expected discounted sum of virtual valuations for a given vector of arrivals and values  $(a, v)$ :*

$$\hat{\tau}(a, v) \in \arg \max_{\tau \in \mathcal{T}} \mathbb{E} \left[ \sum_{i \in \mathbb{N}} e^{-r\tau_i} \left( v_i - \frac{1 - F_{a_i}(v_i)}{f_{a_i}(v_i)} \right) \right].$$

*Then this policy is implementable, and the mechanism that implements it such that the agent with zero value gets an expected utility of zero is revenue maximizing.*

Compared to the scenario where the distribution of values is fixed over time, agents who arrive later obtain an object later and there are now even less incentives to deviate in the arrival time dimension. In the opposite scenario where hazard rates are increasing over time, the virtual value would increase over time and the virtual value maximizing allocation could be increasing in allocation time. Thus,  $\tau^\diamond$  need not remain implementable under unobservable arrivals.

### 4.4 The Non-Markov Case

We conclude this Section by illustrating, via an example, that the Markov assumption is crucial for our result; in its absence, the optimal mechanism with unobservable arrivals may generate strictly less revenue vis-a-vis the case with observable arrivals.

<sup>13</sup>We are grateful to an anonymous referee for suggesting this extension.

<sup>14</sup>As above, we assume that each  $v_{ti}$  has a finite mean and a finite variance, and we continue to make the standard assumption that the virtual valuations functions  $v - \frac{1 - F_t(v)}{f_t(v)}$  are increasing in  $v$  for each  $t$ .

Assume that there is one object, and that the times between two consecutive arrivals (*inter-arrival times*) are I.I.D. The designer believes that all inter-arrival times distribute either uniformly on the interval  $[a, b] = [1, 2]$ , or uniformly on the interval  $[a, b] = [2, 3]$ , and assigns equal probabilities to each possibility. In both cases the distribution of the agents' values is taken to be uniform on the interval  $[0, 1]$ .

If the designer knows the inter-arrival distribution and if he observes arrivals, the optimal mechanism is given by a fixed posted price  $p_{[a,b]}$  that is the unique solution to

$$\frac{2p}{\int_p^1 (v-p)dv} = \frac{\phi_{[a,b]}(r)}{1 - \phi_{[a,b]}(r)} \quad (5)$$

Here  $\phi_{[a,b]}(r)$  is Laplace transform of the uniform inter-arrival distribution on support  $[a, b]$  (see Zuckerman [1986]). Since the uniform distribution on  $[2, 3]$  first-order stochastically dominates the uniform distribution on  $[1, 2]$ , we obtain  $\phi_{[1,2]}(r) > \phi_{[2,3]}(r)$ . Since the left hand side of equation (5) increases in  $p$ , we also obtain  $p_{[1,2]} > p_{[2,3]}$ .

If the designer does not know the arrival process but observes arrivals, he can replicate the same allocation as above with the same posted prices. That is, the designer sets a posted price of  $p_{[1,2]}$  for the time interval  $[0, 2]$ . If at least one agent arrives by time  $t = 2$ , but the object was not sold, the price remains  $p_{[1,2]}$  until it is sold out. If no agent arrives until  $t = 2$ , the posted price drops to  $p_{[2,3]}$ , and remains fixed until it is sold at that price.<sup>15</sup> Note that, in this case, the payoff to the lowest type is zero.

Assume now that the designer knows neither the arrival process, nor does he observe arrivals. To show that the designer cannot extract information about arrivals without paying something for it, we need to show that there is no mechanism that implements the same allocation as above, such that the payoff to the lowest type is also zero. Assume by contradiction that such a mechanism exists. Then, the charged prices must be the same, as otherwise some agents would misreport their valuations. Consider then an agent who arrives before time  $t = 2$  and has value  $v \in (p_{[2,3]}, p_{[1,2]})$ . If he truthfully reports his type, he obtains utility of zero. But, if he delays his report by one unit of time, he obtains the object with probability 1 (since no one can arrive in the meantime) and pays a price below his value. In other words, such a deviation is profitable. Hence, if arrivals are unobservable, the designer must pay a fee to early arrivals, or he needs to distort the allocation away from the one that was optimal under observable arrivals. In both cases, revenue must decrease.

---

<sup>15</sup>In this simple example the time of the first arrival completely reveals the underlying arrival process.

## 5 Extensions

### 5.1 The Name Your Price Mechanism

In the optimal direct mechanisms agents report their values: this feature makes such mechanisms less suitable for practical purposes. Nevertheless, the main distinctive feature of the optimal mechanism - the recall option - is employed in practice. A leading example is the "Name Your Own Price (NYOP)" mechanism and its variants used by many firms following Priceline.com's lead. NYOP was introduced in 2000, and today Priceline.com generates annual revenues of 6.8 billion US\$ and employs 9000 people<sup>16</sup>. In this schemes the customer names a price upon arrival. If the price is above a posted price the buyer gets the object immediately; if it is below the seller decides whether to come back to this customer (after a period of time) and sell at the named price. In this sense, the Name Your-Own-Price mechanism is an extension of a posted price mechanism which allows to sell to low value customers later.<sup>17</sup>

Since the price paid by every buyer in our direct mechanism depends only on the information obtained before the declared arrival time of that buyer (and is thus known to the buyer at the reporting time), it is possible to implement our optimal scheme in an indirect mechanism that has some common features with the NYOP: each agent submits a bid upon arrival, and, at each point in time, the seller decides which of the submitted bids, if any, to satisfy. This indirect mechanism has an equilibrium where every agent submits a bid equal to the optimal expected discounted payment conditional on getting the object, i.e., the payment given by (3). Given these bids, the seller deduce the buyers' values, and then uses the same allocation rule (for the deduced values) as in the direct mechanism.

**Definition 14 (Name Your Own Price Mechanism)** *Every agent decides on a bid and on the time of submitting his bid. At every point in time, the seller decides which of the previously submitted bids to accept (if any). As the function  $P_i$  defined in Eq. (3) is increasing in  $v$  we can define the function  $\hat{v}(b)$  implicitly and recursively in  $i$  for every sequence of reported submission times  $\hat{a}$  by*

$$\hat{v}_i(b, \hat{a}) = \min_{v_i} \{b_i = P_i(\hat{a}, (v_i, \hat{v}_{-i}(b)))\},$$

---

<sup>16</sup>Later similar schemes were adopted by additional companies in other industries (e.g., Chiching.com offers using Name Your Price for local services). In 2005 EBay introduced a design option called "Best Offer" that allows potential buyers to submit an offer to a seller that can accept or reject. ScoreBig offers Name Your Own Ticket Price for sport events and concerts.

<sup>17</sup>Priceline accepts bids within a prespecified, bounded time period. The maximal time after which priceline accepts a bid is 1 hour for national flights and 24 hours for international flights (see <http://www.priceline.com/InformationCenter/html/faq.htm#quest8>).

where we take the minimum over the empty set to be zero. Using those inferred valuations, the seller uses the allocation rule  $\tau^\diamond(\hat{a}, \hat{v}(b, \hat{a}))$ , where  $\hat{a}$  are the submission times<sup>18</sup>.

Recall that  $P_i(\hat{a}, (v_i, \hat{v}_{-i}(b)))$  depends only on the information up until  $\hat{a}$ . Therefore, the recursion is with respect to  $i$ . The dynamic bidding mechanism with the allocation rule  $\tau^\diamond(\hat{a}, \hat{v}(b))$  generates the same expected revenue as the optimal mechanism.

The main difference between the above proposed mechanism and the one used by Priceline.com is the period during which the bids are responded by the seller. The one used by Priceline.com uses an a-priori bounded time limit for the seller to respond to the buyers' bids.

While a NYOP mechanism is practical, it is unrealistic to expect that it can replicate the outcome of the optimal direct mechanism. The reason is that both seller and buyers are subjected there to a heavy computational burden: the buyers need to aggregate complex information into bids, and the seller needs to make complex inferences (inverting bids into values) from the history of arrivals and the ex-ante information. In the next subsection we present a much simpler mechanism that is asymptotically optimal.

## 5.2 A Simple, Asymptotically Optimal Mechanism

In this section we offer a very simple mechanism such that the seller observes only executed sales, and that is often able to achieve a substantial fraction of the optimal revenue. This mechanism consists of two phases: a *learning phase* of length  $T > 0$  and a subsequent *selling phase*. In the learning phase, the seller sets a relatively low price for the objects: this allows him to accumulate information regarding the agents' arrival process. Given the information acquired during the learning phase, the seller chooses a constant, posted price mechanism at the second (selling) phase.

Let us illustrate the workings of this scheme in a simple example. Assume that the arrival process is Poisson, either with a higher rate  $h$ , or with a low rate  $l < h$ . Each of the two rates has the same probability of occurring. The seller is facing several trade-offs: 1. Setting a lower price at the learning stage enhances the learning since more arrivals are identified, but decreases the expected revenues from these sales. 2. Extending the learning phase improves the precision of the seller's information regarding the arrival process, but leaves less units to be sold at potentially higher prices during the second phase.

We analyze the mechanism backwards: we first analyze the second, selling phase; after deriving the optimal posted price at the second phase, we analyze the learning phase. For the sake of simplicity of illustration, we assume here that the agents values distribute uniformly on the  $[0, 1]$  interval.

---

<sup>18</sup>Note, that this allocation rule depends only on the times the players made bids and the bids they submitted.

### 5.2.1 The Selling Phase

If arrivals are governed by a Poisson process with known rate, the optimal mechanism with short lived agents consists of a posted price that depends on the remaining number of objects. Assume that the seller chooses either the optimal posted price schedule given a high arrival rate, or the optimal posted price schedule given the low rate.

Denote by  $y_j$  the minimal value (or *cutoff*) of a buyer who gets an object if there are  $j$  objects left. For a Poisson arrival process with rate  $\lambda$  and discount rate  $r$ , the revenue for any vector of cutoffs  $y = (y_1, y_2, \dots, y_n)$  can be recursively calculated:

$$R_j(y, \lambda) = \frac{(y_j + R_{j-1}(y, \lambda))(1 - y_j)}{\frac{r}{\lambda} + (1 - y_j)}.$$

where  $j$  is the number of the remaining objects. The revenue maximizing cutoffs are, again recursively, given by

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

Both above results follow from Gershkov and Moldovanu, 2015 (Chapter 3.4). While these authors considered only short-lived agents, it follows from the assumption of an infinite horizon that the optimal prices are an increasing function of the remaining number of objects, and hence increase over time. Thus, when those prices are used, it is optimal for even long lived agents to only buy upon arrival. Consequently, the optimal policy for long-lived agents is identical to the optimal policy with short lived agents, as derived in Gershkov and Moldovanu, 2015.

The expected revenue from  $j$  objects is then given by

$$R_j = \frac{\lambda [1 - F(y_j)]^2}{r f(y_j)}.$$

For a uniform distribution of values we obtain:

$$y_j = 1 + \frac{r}{\lambda} - \frac{r}{\lambda} \sqrt{1 + \frac{\lambda}{r} \left(1 + \frac{\lambda}{r} (1 - y_{j-1})^2\right)}.$$

$$R_j = \frac{r}{\lambda} \left(2 + \frac{\lambda}{r} \left(1 + \frac{\lambda}{r} (1 - y_{j-1})^2\right) - 2 \sqrt{1 + \frac{\lambda}{r} \left(1 + \frac{\lambda}{r} (1 - y_{j-1})^2\right)}\right)$$

### 5.2.2 The Learning Phase

With the above expressions at hand we can now analyze the first phase of the mechanism: the seller sets a constant price equal to the lowest possible posted price in the selling phase. There are two complementary reasons for this conclusion:



1. The low price does not create incentives for arriving buyers to wait for the price to decrease, guaranteeing that it is optimal to buy upon arrival, independent of other buyers behavior<sup>19</sup>;
2. The lower price enhance learning, since it induces a larger set of buyers to make a purchase, and hence to reveal their arrival.

After selling  $k$  objects in a learning phase of length  $T$  at a price  $y$ , the posterior probability of the high arrival rate equals

$$\left[1 + (l/h)^k e^{(h-l)(1-y)T}\right]^{-1}.$$

In the selling phase the seller uses the optimal prices for the high (low) arrival rate if the above posterior probability is above (below) a threshold that depends on the number of remaining objects. The optimal threshold is such that the seller is indifferent between using the high cutoff  $y_h$  or the low cutoff  $y_l$ , and it is given by the ratio of losses from choosing the wrong cutoff:

$$\frac{[R_k(y_l, l) - R_k(y_h, l)]}{[R_k(y_h, h) - R_k(y_l, h)] + [R_k(y_l, l) - R_k(y_h, l)]}.$$

As a consequence, it is optimal for the seller to choose the high price in the selling phase if and only if the number of sold objects  $k$  satisfies

$$\frac{R_k(y_l, l) - R_k(y_h, l)}{R_k(y_h, h) - R_k(y_l, h)} \leq \left(\frac{h}{l}\right)^k e^{-(h-l)(1-y)T}.$$

Table 5.2.2 and Figure 5.2.2 shows a calculated example with two possible arrival rates, 1 and 20. The first column (x-Axes) presents the number of total available objects. The second column (light blue line) is a benchmark that shows the fraction of the obtained revenue in a mechanism without any learning as a share of the revenue in the scenario where the seller knows the arrival rate and sets posted prices optimally. The revenue without learning is calculated assuming that the seller is choosing the best among the posted prices - the one that corresponds to either a rate of 1 or a rate of 20. The third column (dark blue line) shows the fraction of the total revenue obtained under the optimal learning phase as a share of the revenue from the scenario with a known arrival rate. The fourth and last column (red line) shows the optimal length of the learning phase.

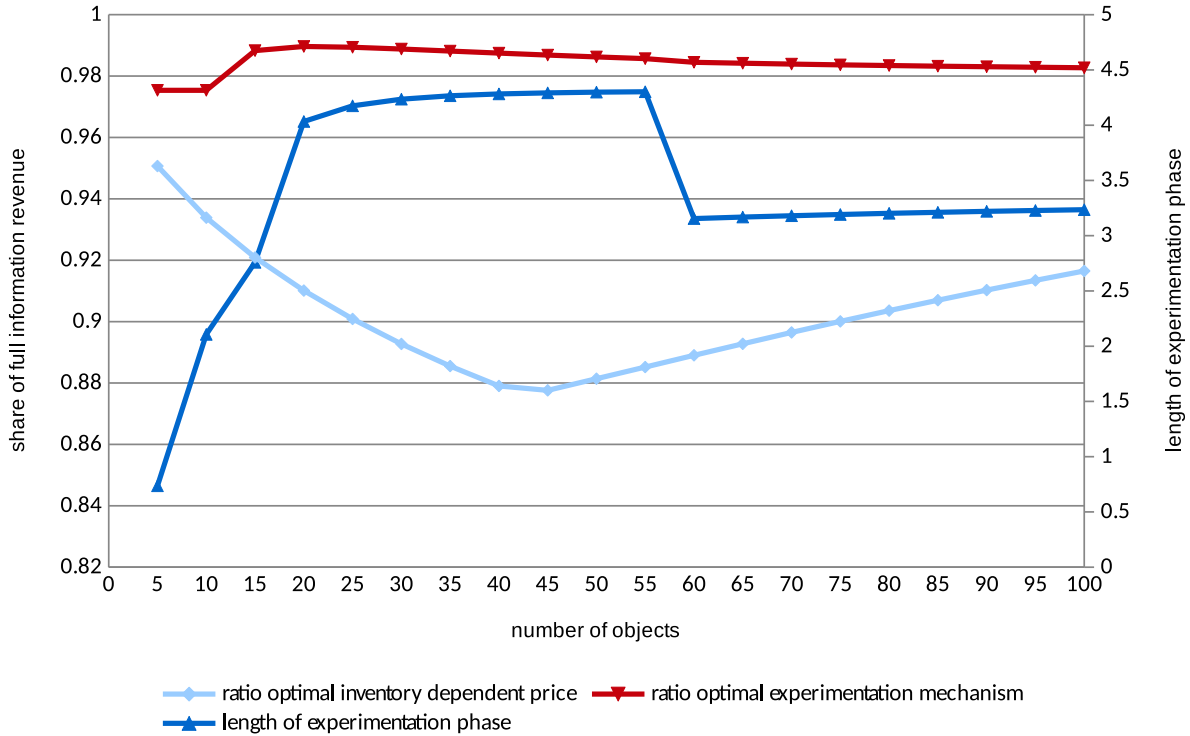
Note that the revenue of the fully optimal mechanism must be between the one achieved in the optimal two-phase mechanism, and the one in an environment with a known arrival rate. With more than 5 objects, the optimal two-phase mechanism is able to get at least 98% of the revenue from the hypothetical

---

<sup>19</sup>Formally, the equilibrium solution concept here is in dominant strategies.

n	No Experimentation	Experimentation	Length
5	0.950683	0.975354	0.735997
10	0.933874	0.985022	2.10662
15	0.920898	0.988287	2.75851
20	0.910127	0.989637	4.03345
25	0.900868	0.989386	4.17518
30	0.892741	0.988813	4.23548
35	0.885515	0.988151	4.26613
40	0.879031	0.987485	4.28286
45	0.877597	0.986845	4.29262
50	0.881384	0.986244	4.29865
55	0.88521	0.985682	4.30253
60	0.889024	0.984474	3.15505
65	0.892789	0.984174	3.16854
70	0.896482	0.983903	3.18082
75	0.900087	0.983657	3.19204
80	0.903591	0.983433	3.2023
85	0.906989	0.983229	3.21167
90	0.910275	0.983041	3.22022
95	0.913448	0.982868	3.22802
100	0.916508	0.982708	3.23512

Table 1: The ratio of the revenue the planer could obtain when knowing the arrival rate and the revenue he obtains by using posted prices (column 2) or the simple mechanism introduced in Section (column 3). The fourth column shows the optimal length of the experimentation phase the first column shows the number of objects.



environment (and hence from the optimal mechanism). It performs substantially better than the posted price mechanism without a learning phase.

Finally, note that our two-phase mechanism is asymptotically efficient for the case where the number of objects goes to infinity. To see this, observe that, if the number of the objects goes to infinity, the posted price mechanism that charges each agent the fixed, static monopoly price is asymptotically optimal.

### 5.3 The Role of Information Disclosure

The seller can use various information disclosure policies, e.g. about the available stock of objects, or about the previous arrivals. Yin et al [2009] showed that hiding information about the remaining stock is revenue enhancing, as it places “more pressure” on the agents with high values. But we often observe disclosure of information about the number of remaining objects, or about the timing of earlier reservations (for example, Expedia reveals information about the number of the remaining seats in the current price category, and about the time of the last booking in suggested hotels). Given that this information disclosure is voluntary, it is reasonable to think that it should not decrease the seller’s expected revenues. Our next result shows that, in our environment, disclosing information to the agents does not affect the ex-ante expected revenue of the designer. The difference in results stems from the imposed restriction on the selling mechanisms in Yin et al [2009]. They analyzed a class of mechanisms in which the seller pre-announces two prices: premium and

post-seasonal prices. In particular, more sophisticated pricing policies that depend on the available stock, or on some additional information such as previous arrivals, are a-priori excluded. Our next result formally proves a revenue equivalence property that holds independently of the information revealed to the agents:

**Theorem 15 (Generalized Revenue Equivalence)** *Suppose that arrivals are unobservable to the principal, and that each agent  $i$  observes at the time of his arrival  $a_i$  a signal  $s_i$  that is (weakly) less informative than observing the prior history of arrivals and reported values  $\{(a_j)_{j \leq \mathcal{N}(a_i)}, (v_j)_{j \leq \mathcal{N}(a_i), j \neq i}\}$ . Then, it is revenue maximizing for the principal to disclose all information to the agents, and to use the mechanism described in Theorem 12.*

The intuitive interpretation of Theorem 15 is that, as the signal  $s_i$  is observable to the principal, he needs not to pay an information rent for it, and thus his revenue is not distorted by it. More interestingly, Theorem 15 also shows that, in the case of observable arrivals, the principal cannot increase revenue by hiding any information from agents. If some information is concealed from an agent, then the agent and the principal have different beliefs about the future arrivals, and hence different probability measures in calculating the discounted time of getting an object. However, after seeing his signal, the agent has correct updated beliefs about the arrivals. The principal only cares about revenue from an ex-ante perspective, which are completely determined by the agents' ex-ante expected allocation. More precisely, given a fixed allocation of the objects, the seller's expected revenue is given by the total social value generated from the allocation less the agents' expected utilities. Using the law of iterated expectations, the agent's ex-ante expected utility from participating in a mechanism where a more precise signal  $s_i$  is revealed equals the agent's ex-ante expected utility from participating in the mechanism where no information is revealed and where the agent can condition his strategy on arrival time and value only.

#### 5.4 The Effect of Strategic Agents on Revenue

As mentioned in the introduction, a major issue in the recent literature on applied revenue management has been to quantify the cost of strategic arrivals. This literature compares two scenarios: In the first, buyers arrive and decide whether to buy or not, and then leave the mechanism immediately. In the second, buyers strategically time their purchase. The following proposition shows that, at the optimum, the obtained revenue is always higher when customers strategically time their purchase! This holds because the seller is able to better intertemporally price-discriminate by “keeping in store” customers that may myopically not buy at a time where prices are too high, and then disappear. This benefit accrues here despite the fact that learning about demand may be potentially disrupted by consumers' strategic behavior. We assume below that the

seller is either in a scenario with short-lived consumers or in a scenario with long-lived customers that can time their buying decision, and that he knows which scenario is the relevant one.

**Corollary 16 (Strategic)** *The revenue in the optimal mechanism for short-lived agents (that only buy upon arrival) is lower than the revenue in the optimal mechanism for agents who can strategically time their buying decision, and whose arrivals are unobservable.*

Board and Skrypacz [2015] reached the same conclusion using deterministic posted price mechanisms in a simpler model where arrivals follow a known Poisson process. In our model the same result holds true despite the effect of later arrivals on the designer's beliefs, continuation value and pricing scheme.

It is important to note that the Markov assumption about the arrival process is crucial for the above conclusion. The optimal policy with observable arrivals in the example of Section 4.4 only allocates objects upon arrival. Hence it follows that it is also implementable if agents are short-lived, and therefore the revenue there is strictly lower if agents can strategically time their purchase !

## 6 Conclusion

We have conducted the first, full mechanism design analysis of the the revenue maximizing mechanism in a model where the unobservable arrivals are governed by a Markov counting process, and where agents are privately informed both about values and arrival times. Since arrivals may be correlated, the seller learns along the way about future arrivals, while agents strategically time their purchases.

Because of the effect of learning while earning, the revenue maximizing mechanism does not generally reduces to posted prices, and needs to charge personalized prices. But it is possible to implement it in an indirect mechanism that resembles the well-known Name-Your-Own-Price mechanism: each agent submits a bid upon arrival, and, at each point in time, the seller decides which of the submitted bids, if any, to satisfy. In our framework, the optimal strategy in a NYOP mechanism imposes a relatively high computational burden on the buyers. Thus, we also presented a two-phases "learn and then sell" mechanism that is asymptotically optimal. In this mechanism the seller first charges a relatively low price that allows him to learn about the arrival process. In a second stage, the seller charges the optimal fixed posted price given the information obtained in the first stage.

An important insight from our analysis is that the presence of long-lived, strategic agents yields here a higher revenue than the one that can be optimally obtained in the situation where agents must buy immediately upon arrival. In other words, we show that appropriate revenue management techniques can be used to overcome and even draw benefits from the presence of strategic consumers.

Finally, our analysis also answers the question whether the seller can increase revenue by withholding information from potential buyers: if the seller uses the optimal mechanism (which, recall does not generally reduce to posted prices), hiding information is never beneficial.

## References

- [1977] Albright S.C.: "Optimal Sequential Assignments with Random Arrival Times", *Management Science* **21** (1), 60-67.
- [2009] Araman, V. and Caldentey,: "Dynamic Pricing for Nonperishable Products with Demand Learning", *Operations Research* **57**(5), 1169-1188.
- [2013] Athey, S. and Segal, I.: "An Efficient Dynamic Mechanism", *Econometrica* **81** (6), 2463-2485.
- [2009] Aviv, Y., Levin, Y. and Nediak, M.: "Counteracting Strategic Consumer Behavior in Dynamic Pricing Systems", in *Consumer Driven Demand and Operations Management*, eds. Tang. C and Netessine, S.
- [2008] Aviv, Y. and Pazgal, A.: "Optimal Pricing of Seasonal Products in the Presence of Forward-Looking Consumers", *Manufacture & Service Operations Management* **10**, 339–359.
- [2005] Aviv, Y. and Pazgal, A.: "Pricing of Short Life cycle Products through Active Learning", discussion paper , Olin School of Business
- [2015] Bergemann, D. and Strack, P.: "Dynamic Revenue Maximization: A Continuous Time Approach", *Journal of Economic Theory* ??
- [2010] Bergemann, D. and Välimäki, J.: "Efficient Dynamic Auctions", *Econometrica* **78**(2), 771–789.
- [2012] Besbes, O. and Lobel, I.: "Intertemporal price discrimination: Structure and Computation of Optimal Policies", discussion paper.
- [2015] Besbes, O. and Zeevi, A.: "On the (Surprising) Sufficiency of Linear Models for Dynamic Pricing with Demand Learning" , *Management Science* **61**(4):723-739.
- [2015] Board, S. and Skrzypacz, A.: "Revenue Management with Forward Looking Buyers", *Journal of Political Economy*, forthcoming

- [2014] Bora Keskin, N. and Zeevi, A. Dynamic Pricing with an Unknown Demand Model: Asymptotically Optimal Semi-Myopic Policies. **Operations Research** **62**(5):1142-1167.
- [2014] Borgs, C., Candogan, O., Chayes, J., Loebel, I., and Nazerzadeh, H.: "Optimal Multiperiod Pricing with Service Guarantees", *Management Science* **60**(7), 1792-1811
- [1993] Boshuizen, F. and Gouweleeuw, J.: "General Optimal Stopping Theorems for Semi-Markov Processes", *Advances in Applied Probability* **24**, 825-846.
- [2010] Cavallo, R., Parkes, D. and Singh, S.: "Efficient Mechanisms with Dynamic Populations and Dynamic Types," discussion paper, Harvard University
- [2015] Chen, Y., and Farias, V.: "Robust Dynamic Pricing with Strategic Customers", *16th ACM Conference on Economics and Computation (EC)*, 2015
- [2000] Courty, P., and Li, H. : "Sequential Screening", *Review of Economics Studies* **67**(4), 697-717
- [1988] Cremer, J., and McLean, R.: "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica*, **56**(6), 1247-1257.
- [2009] Elmaghraby, W., Lipman, S., Tang, C. , and Yin, R.: "Pre-Announced Pricing Strategies with Reservations", *Production and Operations Management* **18**(4), 381-401
- [2012] Escobari, D.: "Dynamic Pricing, Advance Sales and Aggregate Demand Learning in Airlines.", *Journal of Industrial Economics* **60**(4), 697-724.
- [2007] Esö, P., and Szentes, B.: "Optimal Information Disclosure in Auctions and the Handicap Auction", *The Review of Economic Studies* **74**(3), 705-731
- [2003] Etzioni, O., Knoblock, C., Turchinda, R. and Yates, A.: "To Buy or Not to Buy: Mining Airfare Data to Minimize Ticket Purchase Price", in *Proceeding KDD '03, Ninth ACM SIGKDD international conference on knowledge discovery and data mining*, 119-128, ACM: New York, NY.
- [1994] Gallego, G. and van Ryzin, G.: "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons", *Management Science* **40**(8), 999-1020
- [2008] Gallego, G. Kou, S., and Phillips, R.: "Revenue Management of callable products", *Management Science* **54**(3) , 550-564.

- [2006] Gallien, J.: "Dynamic Mechanism Design for Online Commerce", *Operations Research* **54** (2), 291-310.
- [2009a] Gershkov, A. and Moldovanu, B.: "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach", *American Economic Journal - Microeconomics* **2**, 168-198.
- [2009b] Gershkov, A. and Moldovanu, B.: "Learning About The Future and Dynamic Efficiency", *American Economic Review* **99**(4), 1576-1588.
- [2012] Gershkov, A. and Moldovanu, B.: "Optimal Search, Learning, and Implementation", *Journal of Economic Theory* **147**, 881-909.
- [2014] Gershkov, A., Moldovanu, B. and Srack, P.: "Efficient Dynamic Allocation with Strategic Arrivals", discussion paper, University of Bonn
- [2015] Gershkov, A. and Moldovanu, B.: "Dynamic Allocation and Pricing: A Mechanism Design Approach," *The MIT Press*.
- [1986] Green, J. and Laffont, J.J.: "Partially Verifiable Information and Mechanism Design", *Review of Economic Studies* **53**, 447-456.
- [2009] Jerath, K, Netessine, S., and Veeraraghavan, S: "Selling to Strategic Consumers: Opaque Selling Strategies" in *Consumer Driven Demand and Operations Management*, eds. Tang. C and Netessine, S.
- [2010] Jerath, K., Netessine, S., and Veeraraghavan, S: "Revenue Management With Strategic Customers: Last-Minute Selling and Opaque Selling", *Management Science* **56**(3), 430-448.
- [2013] Kakade, S., Lobel, I. , and Nazerzadeh, H.: "Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism", *Operations Research* **56**(4), 837-854
- [2014] Li, J., Granados, N., and Netessine, S.: "Are Consumers Strategic ? Structural Estimation from the Air-Travel Industry", *Management Science* **60**(9), 2114-2137
- [2013] Mantin, B. and Rubin, R.: "Transaction Prices and Strategic Consumer Behavior: Empirical Evidence from the Airline Industry", discussion paper, University of Waterloo
- [2016] Marinesi, S, Girotra, K. and Netessine, S.: "The Operational Advantages of Threshold Discounting Offers", *INSEAD Working Paper No. 2016/41/TOM*



- [2011] Mason, R. and Välimäki, J.: "Learning About The Arrival of Sales", *Journal of Economic Theory* **146**, 1699-1711.
- [2010a] Mierendorff, K.: "Optimal Dynamic Mechanism Design with Deadlines", discussion paper, University of Bonn, forthcoming in *Journal of Economic Theory*
- [1986] Myerson, R.: "Multistage Games with Communication," *Econometrica* 54 (2), 323-358
- [1981] Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operation Research* **6**, 58-73
- [2010] Osdachiy, N., and Vulcano, G.: "Selling with Binding Reservations in the Presence of Strategic Consumers" , *Management Science* ??
- [2008] Pai, M. and Vohra, R.: "Optimal Dynamic Auctions", discussion paper, Northwestern University
- [2003] Parkes, D.C., and Singh, S.: "An MDP-Based Approach to Online Mechanism Design", *Proceedings of 17th Annual Conference on Neural Information Processing Systems (NIPS 03)*
- [2014] Pavan, A., Segal, I. and Toikka, J.: "Dynamic Mechanism Design: A Myersonian Approach", *Econometrica*, forthcoming.
- [1995] Pratt, J.W., Raiffa, H. and Schlaifer, R.: *Statistical Decision Theory*, Cambridge: The MIT Press
- [1990] Presman, Ernst L.: "Poisson version of the two-armed bandit problem with discounting", *Theory of Probability and Its Applications*, **35**, 307–317
- [1983] Ross, S.M.: *Stochastic Processes*, Wiley: New York.
- [2012] Said, M.: "Auctions with Dynamic Populations: Efficiency and Revenue Maximization", *Journal of Economic Theory*, forthcoming.
- [2012] Soysal, G. and Krishnamurthi, L.: "Demand dynamics in the Seasonal Goods Industry: An Empirical Analysis", *Marketing Science* **31**(2), 293-316.
- [1991] Stadge, W.: "A new Continuous-Time Search Model", *Journal of Applied Probability* **28**, 771-778.
- [2007] Su, Xuanming: "Inter-temporal Pricing with Strategic Customer Behavior", *Management Science* **53**(5), 726–741,
- [2004] Talluri, K.T., and Van Ryzin, G.: *The Theory and Practice of Revenue Management*, Springer: New York.

- [2002] Vulcano, G., van Ryzin, G., and Maglaras, C.: "Optimal Dynamic Auctions for Revenue Management", *Management Science* **48**(11), 1388-1407
- [2009] Yin, R., Aviv, Y., Pazgal, A., and Tang, C.: "Optimal Markdown Pricing: Implications of Inventory Display Formats in the Presence of Strategic Customers", *Management Science* **55**(8), 1391-1408.
- [1986] Zuckerman, D.: "Optimal Stopping in a Continuous Search Model", *Journal of Applied Probability* **23**, 514-518
- [1988] Zuckerman, D.: "Job Search with General Stochastic Offer Arrival Rates", *Journal of Economic Dynamics and Control* **12**, 679-684.