Revenue Maximizing Mechanisms with Strategic Customers and Unknown Demand: Name-Your-Own-Price.

June 4, 2015

Abstract

A designer allocates several indivisible objects to a stream of randomly arriving agents. The long-lived agents are privately informed about their value for an object, and about their arrival time to the market. The designer learns about future arrivals from past arrivals, while agents strategically choose when to make themselves available for trade. We characterize revenue maximizing direct mechanism and show that it can be implemented by a "Name-Your-Own-Price" dynamic auction mechanism.

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 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])[1]. As a consequence, customers have an incentive to strategize by carefully timing their purchase[2]. Until recently the execution of such strategies was encumbered by the fact that pricing algorithms are 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[3]. 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). Note that 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. An important question is therefore to understand what are the precise consequences of such behavior, and how the existing pricing techniques could be adjusted in order to take it into account.

As Talluri and van Ryzin [2004] note in their excellent overview of the RM literature, the workhorse models in the scientific literature have "myopic" customers who buy a product as soon as the price drops below their willingness to pay. Moreover, the stochastic pattern of customer arrivals (in other words, the demand pattern) is assumed to be known to the designer, although, in practice, accurate demand forecasting in a changing environment is one of the main difficulties in RM.

In this paper we take a step towards closing the gap between the practical reality and theoretical research. We study a dynamic trading model where a designer allocates

[1]Booking is sometimes possible a year in advance !

[2]See for example Li, Granados and Netessine [2012] who empirically estimate the percentage of strategic customers in the airline industry.

[3]See Etzioni [2003] for the scientifc basis of this development.

several indivisible objects to a stream of randomly arriving, privately informed agents. Agents arrive according to a general Markov counting process. This class is particularly important, since it includes Poisson arrival process with unknown arrival rate[4] (in this case, the designer updates his beliefs about the arrival process and hence about the future demand depending on calendar time and the number of earlier arrivals). All agents are long lived, and each agent is privately informed about her value for an object, and about her 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. Thus, 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 her expected revenue.

We first look at the benchmark case with observable arrivals. In this case, 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 1).

In case of unobservable arrivals the mechanism should take into account the second dimension of the agents private information, namely their arrival time. The main result here (Theorem 3) is based on a monotonicity property of the revenue maximizing allocation: Agents who arrive earlier get the object earlier. We use this property - which is satisfied in our environment by the revenue maximizing allocation with observable arrivals - to show that the payment scheme that maximizes revenue under observable arrivals maximizes revenue even if arrivals are unobservable (as long as the arrival process is Markov!). Intuitively, early arrivals may be detrimental for the agents since, it makes the seller more optimistic about the future arrivals and may induce higher prices. Yet,

[4]Most of the revenue management literature assumes that consumers arrive according to Poisson arrivals with a known rate of arrivals.

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 an agent who delays her arrival and arrives at time t gets exactly the same expected value as an agent who truthfully arrives at t. Note that the agents here solve an optimal stopping problem: when to reveal their presence to the mechanism designer? Together with the fact that the designer's allocation 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, namely that the agent may misrepresent it only in one direction - by arriving or making herself available for trade later. Therefore, only one-directional deviations with respect to arrival time should be taken into account.

In the revenue-maximizing mechanism derived in our paper, the price paid by every buyer depends only on the information obtained before the declared arrival time of that buyer, and is known to him at the reporting time. The main distinctive feature of the optimal mechanism - the recall option - is quite realistic and is actually employed in practice. A leading example is the "Name your Own Price " mechanism and its variants used by many firms following Priceline.com's lead. In these schemes the customer names a price - this is the indirect version of our direct mechanism where the buyer names a price. If the price is above a posted price the buyer gets the object immediately. If it is below the seller decides, after a period of time (during which it potentially sells to others while observing other arrivals and demands,) whether to come back to this customer 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.[5]

We use the above results to show that the presence of long-lived agents that 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 5). In other words, we show that

[5]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).

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] for an excellent discussion of these issues). Most of the previous literature showed that strategic customers hurt seller revenue. The reason we find that strategic customers in our setup always increase the expected revenue of the seller is that we do not restrict the seller to posted-price mechanisms.

Finally, our analysis answers the question whether the seller can increase revenue by withholding information from potential buyers. This question is of high practical relevance: shall an airline inform buyers about the remaining number of seats or not? 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 number of items is low it reduces customers incentive to wait and thus increases the discounted revenue of the seller. When many items are available the opposite happens. In a recent study Yin, Aviv & Pazgal [2009] find that, with a known Poisson arrival process, two types of buyers and a seller who is restricted to posted price mechanisms with two prices, hiding the number of remaining items from potential buyers can increase expected revenue by up to 20%. In contrast to this result we show in Theorem 4 that if the seller is not restricted to a finite number of prices and to posted price mechanisms hiding information is never beneficial.

1.1 Related Literature

Complete information, continuous time, dynamic allocation problems with long-lived agents and with recall have been analyzed by, among others, Zuckerman [1988], Zuckerman [1986], Stadje [1991], and Boshuizen and Gouweleeuw [1993]. In these models, the planner is perfectly informed about the arrival process, and he also observes values and arrival times.

Gershkov and Moldovanu [2009a] and Gershkov and Moldovanu [2010a] analyze efficiency and revenue maximization in continuous-time optimal stopping frameworks where the agents are short-lived (thus there is no recall) and where the planner has several heterogenous objects. In a similar framework (but with discrete time), Gershkov and Moldovanu [2009b] and Gershkov and Moldovanu [2012] allow the planner to learn about the distribution of values from past observations[6]. Learning about values is akin to introducing direct informational externalities (i.e., interdependent values), and these authors show that the efficient implementation of the complete information optimal dynamic policy is only possible under strong assumptions about the learning process. They also characterize the incentive-efficient mechanism (second best) in this framework. In contrast to the case where there is learning via past observed values, a current paper uses a special physical property of the arrival times: agents can only lie in one direction, making themselves available for trade after they arrive, but not before.[7]

Mason and Välimäki [2011] focus on revenue maximizing, posted-price mechanisms in a model with one object and with stochastic and unobservable arrivals of short-lived buyers. The arrival process is Poisson with two possible rates. The present strategic effects of delaying arrivals do not arise in their model because the agents - who can only be served upon arrival - cannot manipulate the designer's belief about the underlying demand. Aviv and Pazgal [2005] consider myopic agents. In their model the arrival process is Poisson, and the belief about its rate follows a Gamma distribution with updated parameters (also in their model consumers cannot strategically affect the designer's beliefs).

Only a small literature considers the effect of patient, strategic customers. Su [2007] determines the revenue maximizing policy for a monopolist selling a fixed supply to a deterministic flow of agents that differ in valuations and patience, and hence have different incentives to wait for sales. Aviv and Pazgal [2008] also consider patient buyers,

[6]Babaioff, Blumrosen, Dughmi and Singer [2011] analyze posted price mechanism with observable arrivals, and evaluate the loss of the revenue due to lack of knowledge of the distribution of the values.

[7]Green and Laffont [1986] pioneered the study of (static) mechanism design problems where particular deviations from truth-telling are not feasible. Other aspects of these (static) environments were analyzed in Deneckere and Severinov [2008], Bull and Watson [2004], [2007], Ben-Porath and Lipman [2012], Kartik, Tercieux [2012].

but restrict the monopolist seller to choosing two prices, independently of the past history of sales. They show that the presence of patient buyers may be detrimental for revenue.

Closer in spirit to the present paper, Gallien [2006] analyzes the revenue maximizing procedure in a continuous time model where the agents are long-lived, and where arrivals are private information. He restricts attention to arrivals that are governed by special, commonly known arrival processes where recall is never used by the complete information stopping policy. In these processes the 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. In other words, Gallien's solution coincides with the one where arrivals are observable, and where agents are short lived (see Albright [1977] and Gershkov and Moldovanu [2009a]).

Board and Skrzypacz [2015] characterize revenue maximization in a finite horizon model with several objects and 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. Thus, there are no informational externalities, and learning from arrivals does not play a role in their model. Roughly speaking, allocative externality payoffs (modified to maximize virtual values instead of values) maximize revenue in their case. Similar to our Corollary 5, Board and Skrzypacz show that, in their environment, strategic customers increase the sellers revenue.[8]

Garrett [2013] also analyzes revenue maximization in a model where arrivals are governed by a known, fixed process. In his setting the values of the agents evolve over time, which yields price fluctuations. Contrasting the present model, there are no informational externalities between the agents since the evolution processes are independent across agents. His analysis builds upon the work of Pavan, Segal and Toikka [2014] who analyzed incentive compatibility and derived optimal mechanisms in very general environments where all the agents are present at all periods but where information arrives over time.

Besbes and Lobel [2012], Pai and Vohra [2008] and Mierendorff [2010a] analyze revenue maximization in a discrete time, finite horizon framework where the arriving agents are

[8]Board and Skrzypacz's [2015] paper and the present article where written in parallel, and independently of each other.

privately informed about values, and about a deadline by which they need to get the object. The distribution of the number of arrivals in each period is known to the designer.

A recent literature 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] and Mierendorff [2010b] construct generalizations of the Clarke-Groves-Vickrey/ Arrow-D'Aspremont-Gerard Varet mechanisms for various environments where either the population or the available information changes over time[9]. It is important to note, for example, that the independence assumptions made by Bergemann and Välimäki do not hold in our model since arrivals are correlated and unobservable[10]. Moreover, the type of mechanisms used by Athey and Segal are excluded here by the requirement that prices cannot depend on information arriving after the allocation time.

Gershkov et. al [2014] analyze the properties and the implementation of the efficient dynamic allocation within a more general arrival model. They find that the welfare maxizing allocation can often not be implemented 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 agents' arrivals are described by a time-inhomogeneous Markov counting process $\{\mathcal{N}(t), t \geq 0\}$ where $\mathcal{N}(t)$ is a random variable representing the number of arrivals up to time t. This means that the probability measure over arrivals after time t depends on the time t, and on the number of prior arrivals.[11] 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 $\bar{T} < \infty$ no more

[9] An earlier literature starting with Dolan [1978] has dealt with similar questions in the more restricted environment of queueing/scheduling. For example, Kittsteiner and Moldovanu [2005] study efficient dynamic auctions in a continuous time queueing framework where agents arrive according to a Poisson process and have private information about needed processing times.

[10] This would correspond to correlations in the times where values jump upwards (from zero) in their model, and hence to correlations of values.

[11] The Markov case is of particular importance, as all commonly used models of learning about the arrival rate of a Poisson process are covered by it (see e.g. Presman [1990], and Example 3.3 in the present paper). In fact, almost the entire literature deals solely with this case.

arrivals occur, i.e., where $\mathcal{N}(t)$ stays constant for any $t \geq \bar{T}$. Since arrivals are described by Markov counting processes, the designer's beliefs about future arrivals may evolve over time, and may depend on the number of past arrivals.

Each agent's private information is two-dimensional: the arrival time $a \geq 0$ and the value $v \geq 0$ he gets if allocated an object. In other words, we assume that the designer does not observe agents' arrivals. If an agent arrives at time a, gets the object at time $\tau \geq a$ and pays p at time $t' \in [a, \tau]$, then her utility is given by

$$e^{-r\tau}v - e^{-rt'}p$$

where $r \geq 0$ is the discount factor. We assume that an item cannot be reallocated after an initial assignment. We denote by $\{k(t), t \geq 0\}$ the number of remaining objects at time t.

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, and with continuous density f. We assume that each v_i has a finite mean, and a finite variance. We make a standard assumption that the virtual valuation $v - \frac{1-F(v)}{f(v)}$ is increasing in v. We also assume that, for each agent, his arrival time is independent of his value. This allows us to focus on the information revealed by manipulating arrivals, as opposed to information revealed by manipulating values.[12] A designer maximizes his expected discounted revenues. That is, a payment of p at time τ' generates utility of $e^{-r\tau'}p$. Potentially, every agent may pay more than once, however, since 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 payment and utilities, in which every agent pays only once.

We denote by τ_i the (random) time the designer allocates an object to the i-th agent (set $\tau_i = \infty$ if agent i does not receive an object) and by τ_i' the time agent i pays. A vector τ of such allocation times is feasible if no agent receives an object before he arrives, $a_i \leq \tau_i$, if no agent pays before his arrival, $a_i \leq \tau_i'$ and if at most k objects are allocated $\sum_{i=1}^{\infty} 1_{\{\tau_i < \infty\}} \leq k$. Denote the set of feasible allocations by $\mathcal{T}$ and by p_i the payment of

[12] Assuming dependence between the arrival times and values will generate correlation between the values. This may create additional complications in implementation, as explored in Gershkov and Moldovanu [2009b] and Gershkov and Moldovanu [2012].

player i. The revenue maximizing principal aims at maximizing

$$\mathbb{E}\left[\sum_{i=1}^{\infty} e^{-r\tau_i'} p_i\right]$$

over the set of feasible and incentive compatible allocations $\mathcal{T}$.

2.1 Direct Revelation Mechanisms

The Revelation Principle for dynamic environments (see Myerson [1986])[13] implies that we can restrict attention to mechanisms where each agent reports his value and arrival time, and where the mechanism specifies a probability of getting the object and a payment as a function of the reported value, reported arrival time and the time of the report. Moreover, without loss of generality, we can restrict attention to direct mechanisms where each agent reports his type upon arrival, e.g., the time of the report coincides with the arrival time.[14]. To maximize revenue, we could restrict attention to mechanisms where agents do not observe the history[15]. But, we conduct our analysis in a different way: we first assume that all the agents observe the entire history of reports, and we later show that, assuming this information structure, does not affect the seller's expected revenue. That is, we show that the seller can never generate any additional revenue from hiding information from the agents. Since recall may be employed by the optimal policy, an allocation and a payment to an agent can be conditioned also on information that accrues between the arrival of that agent and the allocation time.

We shall assume below that, upon arrival, the agents observe the entire history of the reports made by previous agents. Nevertheless, we show that revelation of previous information by the designer does not affect the ex-ante expected revenue. Therefore, the optimal mechanism obtained here under the full disclosure is optimal even if the designer

[13]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].

[14]The equilibrium outcome of any mechanism where at least one agent reports his type (value and the arrival time) after his arrival, can be replicated by another mechanism and equilibrium where all agents reports their types upon arrival.

[15]Intuitively, minimizing the information revealed to each agent reduces the available contingent deviations from truthtelling, and therefore relaxes the incentive compatibility constraints for that agent.

could hide from the agents historical information revealed by earlier arriving agents.

A direct mechanism specifies at every time t and for every agent that reported an arrival at that or any earlier time the probability of getting the object and a payment at t. We assume that agents leave the mechanism as soon as they obtain an object, and thus we do not allow an agent's payment to depend on information that accrues after her allocation:

Assumption 1 (Agents' Exit Condition) *The payment P_i made by agent i only depends on information accruing before the time she gets allocated an object, τ_i.*

Assumption 1 is identical to the efficient exit condition made in Bergemann and Välimäki [2010]. We discuss the effect of this assumption in later stages. This property is very appealing as in many applications uncoupling of the physical and monetary parts, and conditioning the payment on the information arrived after the physical allocation is not always realistic, and we abstract from it here.

Assumption 2 (Ex-Post Individual Rationality) *We restrict attention to ex-post individual rational mechanisms, where, after every history, the equilibrium utilities of all agents are non-negative.*[16]

2.2 Observable Arrivals

Let us first briefly consider the benchmark case of welfaremaximizing designer who observes both the agents' arrivals and their values for the object, so that agents have no private information. Our environment is then equivalent to a standard continuous-time search model with perfect recall. Since the main focus here is on the implementation of the revenue maximizing dynamic allocation (or, equivalently, the implementation of the classical welfare maximizing policy), we assume that a unique optimal policy in the complete information model exists.

[16] Assumption 2 excludes schemes a la Cremer and McLean which could otherwise be used to extract correlated arrival times at no cost.

We denote by $\tau^\diamond \in \mathcal{T}$ the policy maximizing the expected discounted sum of virtual valuations:

$$\tau^\diamond \in \arg\max_{\tau\in\mathcal{T}} \mathbb{E}\left[\sum_{i\in\mathbb{N}} e^{-r\tau_i}\left(v_i - \frac{1-F(v_i)}{f(v_i)}\right)\right].$$

It follows from the dynamic programming principle that the optimal stopping policy $\tau^\diamond$ is deterministic. As $\tau^\diamond$ depends only on the realized sequence of arrival times $a = (a_1, a_2, \ldots)$ and valuations $v = (v_1, v_2, \ldots)$ and hence we use the notation $\tau^\diamond(a, v)$. To shorten notation we denote by $h(t)$ a history at time t, i.e.

$$h(t) = \left((a_j)_{j\leq\mathcal{N}(t)}, (v_j)_{j\leq\mathcal{N}(t)}\right).$$

Finally, we denote by $\mathbb{E}[\cdot \mid h(t)]$ the probability measure conditional on the history $h(t)$.

As an intermediate step we consider now the case where values are private information, but where arrival times are observable. We characterize below the optimal allocation and the transfers that implement it.[17]

Theorem 1 (Revenue Maximizing Policy) *When arrivals are observable, the policy $\tau^\diamond$ is implementable using the following payments*

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

charged upon allocation. In addition, this mechanism is revenue maximizing.

Proof. See Appendix. ■

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

The proof of the above theorem is based on a few observations. First, we show that monotonicity of the stopping time with respect to agent's valuation is a necessary and sufficient condition for implementation. Monotonicity implies that a higher value of agent i must lead to a smaller expected discounted time of getting the object. Second, we show that, in the optimal mechanism, hiding information about earlier arrivals and values does not increase the seller's expected revenues. Finally, we show that the stopping rule $\tau^\diamond$ maximizes the seller's expected revenue.

[17] Proposition 8 in the Appendix shows an analogous result for welfare maximizing mechanisms.

3 Unobservable Arrivals and Implementation via the Name-Your-Own-Price Mechanism

It is not at all clear that the revenue-maximizing allocation derived in Theorem 1 remains implementable also in case that the arrivals are unobservable. Recall that in case of a Markov arrival process, the probability measure on arrivals after time t, depends on the time t and on the number of arrivals prior to time t. Assume that all the agents but some arbitrary player i report truthfully their arrival time, while player i that arrives at a_i consider reporting $a_i' > a_i$. In both cases, after time a_i' the designer will have the same beliefs about the future arrivals. Therefore, one may think that agents' incentives about arrival time are trivially satisfied. This however is not true, since in case of arrival at time a_i rather then at a_i', the designer becomes more optimistic (in the time interval between a_i and a_i') about the future arrivals. Hence she may charge higher price than in case of arrival at time $a_i' > a_i$. Therefore, arriving later may decrease the expected payment. The disadvantage of this deviation is that it may decrease the chance of getting the object in case of arrivals of additional agents between a_i and a_i'. Nevertheless, we show that the payments characterized in Theorem 1 continue to generate correct incentives even if arrivals are unobservable.

In the optimal direct mechanisms agents are supposed to report their values. Direct mechanisms are somewhat less appealing in practice, and therefore we suggest another, more intuitive indirect mechanism. This mechanism has many common features with the "Name Your Own Price" mechanism, pioneered by Priceline.com. That scheme was introduced in 2000, and today Priceline.com generates annual revenues of 6.8 billion US$ and employs 9000 people. There a buyer specifies a price he is willing to pay for a product (e.g., an airline ticket or hotel room). The seller then decides after a period of time (during which it potentially sells to others while observing other arrivals and demands) whether to come back to this customer and sell at the named price.[18]

[18]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.

Definition 2 (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. (1) 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*

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

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[19].

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 next theorem shows that the allocation and the payment scheme characterized in Theorem 1 are implementable also in case of unobservable arrivals and thus revenue maximizing. Furthermore, we show that the dynamic bidding mechanism with the allocation rule $\tau^\diamond(\hat{a}, \hat{v}(b))$ generates the same expected revenue as the optimal mechanism.

Theorem 3 (Unobservable Arrivals and Name-Your-Own-Price)

1. *The payment defined in Theorem 1 implements policy $\tau^\diamond$ under unobservable arrivals, and this mechanism is revenue maximizing.*

2. *The Name-Your-Own-Price mechanism with the allocation rule $\tau^\diamond(\hat{a}, \hat{v}(b))$ is revenue maximizing.*

Proof. See Appendix. ■

Note that if an agent gets an object immediately all agents with higher valuations also get the object immediately. By incentive compatibility, all agents who get the object immediately must pay the same price. Let us denote this price by $p(h(t))$. As any agent who makes a bid of $p(h(t))$ in the indirect mechanism gets the object immediately, we can interpret $p(h(t))$ as a posted price (or buy it now price). Thus, the dynamic bidding

[19] Note, that this allocation rule depends only on the times the players made bids and the bids they submitted.

mechanism proposed in this paper is an extension of a posted-price mechanism where agents who value the good less than the posted price have the chance to make lower bids and be recalled later.

The first result above uses the fact that it impossible for the agent to "fool" the designer about the latter's continuation value. At the time the agent reports his arrival, the designer knows his true continuation value since this only depends on the information that the agent arrived, but not on the precise arrival time. The result is, however, non-trivial since the agent can still manipulate the price by changing the reported arrival time. For example, in a learning setting, 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 were 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, the monotonicity of the allocation yields monotonicity of the expected utility, while prices (given the allocation) again play a secondary role, and are determined by the 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.

Note also that in the direct mechanism an agent who decides when to report his arrival solves an optimal stopping problem about her arrival time! Proving that truthful reporting of arrival times is incentive compatible under the payment scheme of Theorem 1 is difficult because no explicit solution $\tau^{\diamond}$ to the designer's (complete information)

optimization problem 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. In the proof of Theorem 3 we first show that monotonicity in arrivals is sufficient to ensure truthfulness even when arrivals are unobservable. This monotonicity requires that agents who arrive earlier get the object earlier. In a second step we show how the fact that the designer's allocation decision is itself the solution of an optimal stopping problem ensures the necessary monotonicity in arrival times, so that it is indeed optimal for every agent to announce his presence immediately.

3.1 Information Disclosure

In dynamic environments 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. This observation is puzzling since we often see that sellers disclose 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 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 preannounces two prices: premium and postseasonal 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 which holds independently of the information revealed to the agent:

Theorem 4 (Generalized Revenue Equivalence) *Suppose that arrivals are unobservable to the principal, and that each agent i observes at the time of her arrival a_i a signal s_i which 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 from Theorem 3-1.*

Proof. See Appendix. ■

The intuitive interpretation of Theorem 4 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 4 also shows that, in the case of observable arrivals, the principal cannot increase revenue by hiding any information from agents: as an agent forms correct expectations about his chance of getting an object, he can not be "fooled" on average. If some information is concealed from an agent, then the agent and the principal have different beliefs about the the future arrivals, and hence different probability measures in calculating the discounted time of getting the object. However, the agent has the "correct" ex-ante beliefs about the arrivals. Since the principal cares about the ex-ante revenues, he is willing to adopt the agent's beliefs.

3.2 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 situations: In the first situation, buyers arrive and decide whether to buy or not, and then leave the mechanism immediately. In the second situation, 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 here 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.

Corollary 5 (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.[20]

Proof. See Appendix. ■

Board and Skrypacz [2015] reached the same conclusion using deterministic posted price mechanisms in a model where arrivals follow a 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.

Remark 1 *In the analysis above we assumed (Agents' Exit Condition) that the payment made by agent i only depends on information accruing before the time she gets the object. However, this assumption is without loss of generality and the seller cannot increase his expected revenue even if he could condition the payment for an object on the information obtained after allocating that object.*

We now illustrates the above results with an example of Poisson arrival process with unknown rate of arrivals.

3.3 Example: Learning 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

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

and note that the posterior expected arrival rate $\lambda(t)$ only depends on time t and 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

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

[20] We assume here that the seller is able to distinguish between the long- and short-lived agents.

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.

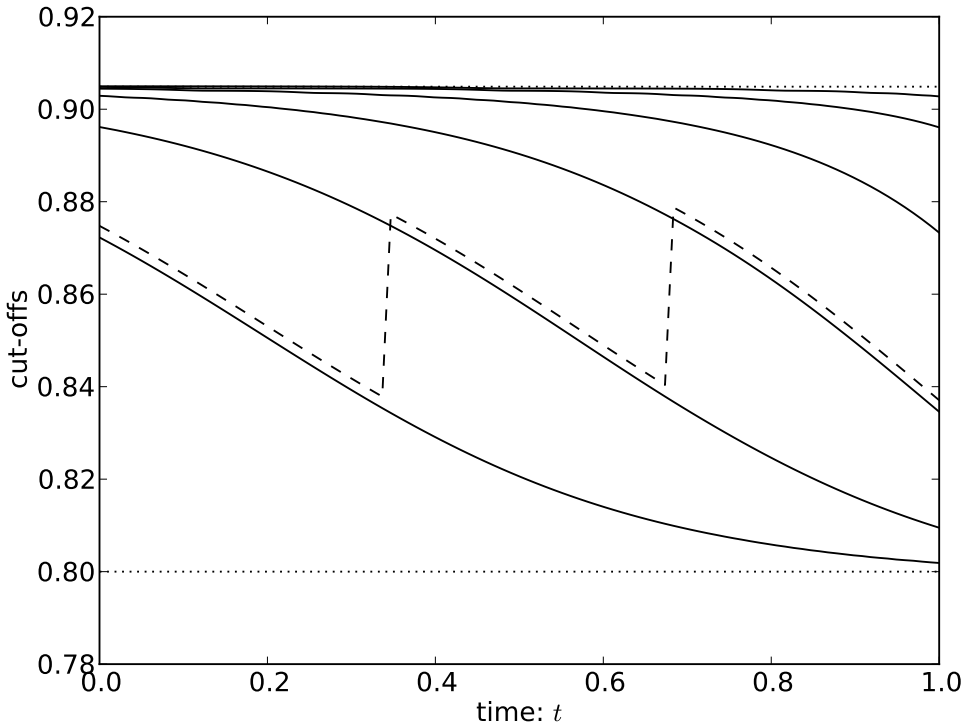

Figure 1: Numerical solution of the optimal cutoffs as a function of time. The cutoffs are welfare maximizing when the valuations are uniform distributed on $[0.5, 1]$ and revenue maximizing when valuations are uniform distributed on $[0, 1]$. Depicted here are the cutoffs for zero to five arrivals in the single object situation with a prior that assigns probability to the arrival rates 1 and 5. Higher cutoffs correspond to more arrivals. The deadline equals $\bar{T} = 1$ and the exponential discount rate is given by $r = 0.05$. 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. The dotted lines mark the optimal constant cutoff's when the arrival rate is known to be 1 or 5. The expected arrival rate at time zero equals two. The dashed line is an example path where at time 1/3 and 2/3 an agent arrives and the optimal cutoff jumps up.

Note that 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 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 together with a fire-sale auction at the deadline (if any) (see Board and Skrypacz [2015]).

In our case with learning, the value obtained by a fixed continuation strategy is higher when the arrival rate is higher, and thus the optimal policy allocates the object later. Thus, as only higher types acquire the object, the optimal cut-off and price both jump up after every arrival (for an illustration see Figure 1)[21]. In practice the arrival of the agent might be a physical arrival in a store, or a visit on a website selling the good.

Ignoring the opportunity of learning may cause a significant loss in revenue: For example in the setup of Figure 1 setting the optimal constant cut-off of 0.854 (i.e., optimal for the time zero expected arrival rate of 2) yields an expected welfare/revenue of 0.577. This represents a loss of approximately 19% compared to the optimal optimal policy that generates here a welfare/revenue of 0.711.

4 Conclusion

We have analyzed dynamic allocation and pricing in a continuous time, discounted model where 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 planner learns along the way about future arrivals. Besides the theoretical interest of extending the static mechanism design paradigm to a basic dynamic allocation problem, we see the main applications of our model and methods to revenue management techniques in complex situations where capacity is limited, where agents can strategically choose the time of their purchases, and where the underlying stochastic nature of the demand pattern must be learned along the way. We offer tools and insights that allow providers of Revenue Management and of consumer search tools to extract the benefits from such situations.

[21] In addition, with several objects, cutoffs and prices jump after each sale since supply becomes smaller.

5 Appendix

5.1 Proof of Theorem 1

The proof of Theorem 1 is based on a couple of results interesting on their own. As an intermediate step we provide a characterization of a mechanism that implements the welfare maximizing (the dynamically efficient) allocation.

Definition 6 *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 $h(a_i)$. A stopping time τ is* monotone with respect to valuations *if for any agent i, and for any two such histories it holds that*

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

The next theorem shows that the monotonicity with respect to valuations is crucial for implementation.

Theorem 7 (Monotone Allocations are Implementable) *Assume that arrivals are observable. If the allocation τ is implementable then it is monotone with respect to valuations. Conversely, if the allocation τ is monotone with respect to valuations, then τ can be implemented using a payment paid at the time of allocation τ_i:*

$$P_i(a,v) = \mathbf{1}_{\{\tau_i(a,v)<\infty\}} \frac{\mathbb{E}\left[\int_0^{v_i}\left(e^{-r\tau_i(a,v)} - e^{-r\tau_i(a,(z,v_{-i}))}\right)\mathrm{d}z \mid h(a_i)\right]}{\mathbb{E}\left[e^{-r\tau_i(a,v)} \mid h(a_i)\right]}. \tag{2}$$

Proof. Under observable arrivals, any agent reports his type (valuation) at just one point in time, and the agent's incentive problem is analogous to a static one. The expected utility of agent i who has valuation v_i, arrives at a_i and reports truthfully equals

$$U_i(h(a_i)) = \mathbb{E}\left[e^{-r\tau_i}(v_i - P_i(h(a_i)) \mid h(a_i)\right].$$

By the standard envelope argument, in every mechanism where the agent reports his value v_i truthfully, it holds that

$$\frac{\partial U_i(h(a_i))}{\partial v_i} = \mathbb{E}\left[e^{-r\tau_i(a,v)} \mid h(a_i)\right]$$

The result follows immediately from the standard static analysis (see Myerson [1981]) by noting that $\mathbb{E}[e^{-r\tau_i(a,(v_i,v_{-i}))} \mid h(a_i)]$ plays here the role of the probabilistic allocation function in Myerson's analysis. ■

For future reference, we denote by $\tau^\star \in \mathcal{T}$ the welfare maximizing policy under complete information

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

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

We now show that the dynamically efficient allocation is implementable by a mechanism where each agent pays the expected externality he imposes on other current, and on future agents. 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]), in the case of observable arrivals the dynamic pivot mechanism implements the efficient allocation in the case of observable arrivals since, conditional on the observable arrivals, the agents' values are independent.

Let $P_i(a, v)$ denote the payment charged to agent i at time τ_i when he gets an object, as a function of the arrival times and valuations reported by all agents. The next Proposition shows that a payment equal to the expected externality conditional on the information available to the agent at arrival divided by the expected discounted allocation time implements the socially efficient allocation:

Proposition 8 (Pivotal Payment) *The payment scheme*

$$P_i(a, v) = \mathbf{1}_{\{\tau_i(a,v)<\infty\}} \frac{\mathbb{E}\left[\sum_{j\neq i}(e^{-r\tau_j^\star(a,(0,v_{-i}))} - e^{-r\tau_j^\star(a,v)})v_j \mid h(a_i)\right]}{\mathbb{E}\left[e^{-r\tau_i^\star(a,v)} \mid h(a_i)\right]} \tag{3}$$

implements the efficient dynamic allocation policy τ^*. The resulting mechanism is ex-post individually rational. Moreover, the efficient dynamic allocation policy τ^* is monotone with respect to valuations.

Proof. We prove that the mechanism is incentive compatible if each agent can observe all arrivals and valuations of agents that arrived prior to herself. First note that the payment P_i only depends on arrivals and valuations of agents who arrived prior to i, and

thus is fixed for any future transaction that involves i. The expected value of agent i when he arrives at time a_i and reports $(a_i, \hat{v}_i)$ equals

$$\begin{aligned}
&\mathbb{E}\Big[e^{-r\tau_i^\star(a,(\hat{v}_i,v_{-i}))}\,(v_i - P_i(a,(\hat{v}_i,v_{-i})))\mid h(a_i)\Big]\\
=&\mathbb{E}\Big[e^{-r\tau_i^\star(a,(\hat{v}_i,v_{-i}))}\mid h(a_i)\Big]\,(v_i - P_i(a,(\hat{v}_i,v_{-i})))\\
=&\mathbb{E}\Big[e^{-r\tau_i^\star(a,(\hat{v}_i,v_{-i}))}v_i + \sum_{j\neq i}(e^{-r\tau_j^\star(a,(\hat{v}_i,v_{-i}))} - e^{-r\tau_j^\star(a,(0,v_{-i}))})v_j \mid h(a_i)\Big]\\
=&\mathbb{E}\Big[\sum_{j\in\mathbb{N}} e^{-r\tau_j^\star(a,(\hat{v}_i,v_{-i}))}v_j \mid h(a_i)\Big] - \mathbb{E}\Big[\sum_{j\neq i} e^{-r\tau_j^\star(a,(0,v_{-i}))}v_j \mid h(a_i)\Big]
\end{aligned}$$

Note that only the first part of the last expression above depends on the report of the agent: this exactly equals the value of the principal when agent i arrives with valuation v_i and when the policy $\tau^\star(a,(\tilde{v}_i,v_{-i}))$ is used. By definition, this is less than the value when the optimal policy $\tau^\star(a,(v_i,v_{-i}))$ is used, and thus it is optimal for the i to report v_i truthfully. In this case the expected utility of agent i is exactly equal to the difference between the values of the principal if agent i arrives with valuation either v_i or zero, which is positive, again by optimality. As the payment is fixed, this implies that $P_i(a,v)$ is less than the valuation of the agent v_i, and thus that the mechanism is individually rational. Monotonicity follows from the previous construction and from Theorem 7. ■

Corollary 9 (Monotonicity in Valuations) *The stopping time which maximizes the expected discounted sum of virtual valuations $\tau^\diamond$ is monotone in valuations.*

Proof. Proposition 8 shows that the welfare maximizing policy is monotone in valuations. Replacing valuations by virtual valuations shows that, in the policy maximizing the expected discounted sum of virtual valuations, agent $i's$ expected discounted probability of getting an object is increasing in her virtual valuation $v_i - \frac{1-F(v_i)}{f(v_i)}$. Finally, the result follows from the monotonicity of the virtual valuation function. ■

The next theorem specifies the expected revenue in any incentive compatible mechanism.

Theorem 10 (Generalized Revenue Equivalence) *Suppose that arrivals are observable to the principal, and that each agent i observes at the time of her arrival a_i a signal s_i which is (weakly) less informative than observing the prior history of arrivals and*

reported values $\{(a_j)_{j\leq\mathcal{N}(a_i)},\ (v_j)_{j\leq\mathcal{N}(a_i),j\neq i}\}$.[22] *Furthermore, assume that this signal is observable to the principal. Than, 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*

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

Proof. As the signal s_i is observable to the principal, he could potentially condition the allocation τ_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 is given by

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

It follows from the law of iterated expectations that the ex-ante (i.e., before seeing his signal) expected payoff (or information rent) to agent i is given by

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

The outer expectation on the left-hand side is over the signals 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, and for the rest of the proof we shall only use this expectation. In the last step we use the usual integration by parts argument to get revenue equivalence;

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

[22] 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.

Thus, we have that the revenue of the principal in any incentive compatible mechanism equals

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

■

We now have all the necessary tools in order to prove Theorem 1:

Proof of Theorem 1. By Corollary 9, the policy $\tau^\diamond$ is monotone with respect to valuations, and thus implementable by Theorem 7. By the definition of $\tau^\diamond$ as the virtual valuation maximizing allocation, and by Theorem 10, the revenue in any other implementable mechanism is lower than the revenue in this mechanism. ■

5.2 Proof of Theorem 3

The main property conducive to implementation under unobservable arrivals is the monotonicity of a stopping time with respect to arrivals:

Definition 11 (Monotonicity in Arrivals) *A deterministic stopping time τ is monotone in the 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}, $v \in \mathbb{R}_+^\infty$*

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

Theorem 12 *Consider a vector of deterministic and Markovian*[23] *stopping times τ and a vector of payments $\mathbf{P}$ such that:*

1 the payment $\mathbf{P}$ implements τ under observable arrivals

2 τ is monotone with respect to arrivals

3 $(\tau, \mathbf{P})$ leaves every agent with a value of zero with a utility of zero

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

[23]Markovian: the stopping decision at every point in time depends only on calendar time and on the number of arrivals, but not on their precise timing.

Proof. The expected utility of agent i who has valuation v_i, arrives at a_i and reports truthfully equals

$$U_i(h(a_i)) = \mathbb{E}\left[e^{-r\tau_i}(v_i - P_i(a,v)) \,|\, h(a_i)\right] .$$

By the standard envelope argument, in every mechanism where the agent reports his value v_i truthfully it holds that

$$\frac{\partial U_i(h(a_i))}{\partial v_i} = \mathbb{E}\left[e^{-r\tau_i(a,v)} \,|\, h(a_i)\right] .$$

Thus, the payoff of the agent in any incentive compatible mechanism (in the valuation dimension) equals

$$\begin{aligned} U_i(h(a_i)) &= \int_0^v \mathbb{E}\left[e^{-r\tau_i(a,(z,v_{-i}))} \,|\, h(a_i)\right] \mathrm{d}z + U_i((a_1,\ldots,a_i),(v_1,\ldots,v_{i-1},0)) \\ &= \mathbb{E}\left[\int_0^v e^{-r\tau_i(a,(z,v_{-i}))}\mathrm{d}z \,|\, h(a_i)\right] + U_i((a_1,\ldots,a_i),(v_1,\ldots,v_{i-1},0)) . \end{aligned}$$

By (3) the designer does not make any transfer to the lowest type and thus the second summand equals zero. We show that no agent i has an incentives to misreport his arrival time as $\tilde{a}_i \geq a_i$. As the stopping time and the process are Markov, conditional on not having allocated the object between a_i and $\tilde{a}_i$, the stopping time τ does not change if the agent deviates and reports his arrival at $\tilde{a}_i$. Consequently, conditional on the object not being allocated between a_i and $\tilde{a}_i$, an agent who misreports his arrival time as $\tilde{a}_i > a_i$ and his valuation as $\tilde{v}_i \neq v_i$ receives the same payoff as the agent who truly arrived at that time $\tilde{a}_i$ with valuation v_i and misreported his valuation to be $\tilde{v}_i$. Hence, the optimality of truthful reporting (1) implies that even if he deviates by misreporting his arrival time it will be optimal for him to report his valuation truthfully. Let us denote by $\tilde{h}$ the counterfactual histories where agent i arrived at time $\tilde{a}_i$. The expected utility of reporting the arrival at the (stopping) time $\tilde{a}_i$ equals:

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

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 just the law of iterated expectations. The loss in expected payoff for agent i from reporting the arrival at time $\tilde{a}_i \geq a_i$ instead of a_i is thus given by:

$$\begin{aligned}
& U_i(h(a_i)) - \mathbb{E}\left[U_i(\tilde{h}(\tilde{a}_i)) \,|\, h(a_i)\right] \\
= \; & \mathbb{E}\left[\int_0^v e^{-r\tau_i(a,(z,v_{-i}))}\mathrm{d}z \,|\, h(a_i)\right] - \mathbb{E}\left[\int_0^{v_i} e^{-r\tau_i((\tilde{a}_i,a_{-i}),(z,v_{-i}))}\mathrm{d}z \,|\, h(a_i)\right] \\
= \; & \mathbb{E}\left[\int_0^v e^{-r\tau_i(a,(z,v_{-i}))} - e^{-r\tau_i((\tilde{a}_i,a_{-i}),(z,v_{-i}))}\mathrm{d}z \,|\, h(a_i)\right] \\
\geq \; & 0\,.
\end{aligned}$$

Here, the last step follows from (2), i.e. from the monotonicity of the allocation with respect to arrival times. ■

Theorem 12 can be used directly to prove that the revenue maximizing allocation $\tau^\diamond$ is implementable even if arrivals are unobservable.

Proof of Theorem 3. **1.** It follows from the dynamic programming principle that the virtual valuation maximizing policy is Markov. Furthermore, the payments defined in Theorem 7 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 an agent arrives later, the principal uses the same continuation strategy as he would have used conditional on not allocating the object to the agent, 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 gets an object before she arrives.

2. Recall that the price paid by agent i at the time of allocation in Theorem 1 depends only on the information obtained before the arrival of that agent and thus $P_i(a,v)$ is a feasible bidding strategy in the indirect mechanism.

First, we show that bidding $P_i(a,v)$ at the time of arrival is incentive compatible, and second that it implements the same allocation as the revenue maximizing, direct mechanism. By the revenue equivalence result proved above, the indirect mechanism generates the same expected revenue to the seller.

Incentive compatibility: Suppose that there exist a bid $b \neq P_i(a, v)$ that increases the expected utility of buyer i in the indirect mechanism. If the agent submits a bid at time a which is inconsistent with any value that arrives at time a, he never gets the object. This gives him the lowest possible payoff of zero, and thus such a bid is never a profitable deviation. If the bid b is consistent with some valuation $\hat{v}$, buyer i can profitably deviate already in the direct mechanism by reporting the valuation $\hat{v}_i(b, \hat{a})$ inferred by the seller from this bid. Similarly, the existence of a profitable deviation in the arrival time dimension in the indirect mechanism implies the existence of a profitable deviation in the direct mechanism. Since the direct mechanism is incentive compatible, it is optimal for every agent in the indirect mechanism to submit immediately upon arrival a bid equal to the payment $P_i(a, v)$ given in Theorem 1.

Revenue Maximization: Since the stopping policy is given by $\tau^\diamond$, and since the virtual values are monotone, the revenue maximizing allocation has a cutoff property: for every history, values above some cutoff get the object. If there is a unique valuation that is consistent with the bid made by the agent at time a_i the seller infers buyer i's valuation perfectly from her bid. If two different values in the direct mechanism are supposed to get the object with the same probability and to pay the same price, the designer is indifferent between them and he may apply stopping policy $\tau^\diamond$ with respect to any of these values. Therefore, the indirect mechanism implements an allocation which leads to the same revenue as the revenue maximizing, direct mechanism. ■

5.3 Remaining Proofs

Proof of Theorem 4. An upper bound for the revenue of the principal is the revenue he could obtain if he would be able to observe the agents arrivals, which is given by Theorem 10. In this case his revenue is maximized by disclosing all information to the agents, and by using the virtual valuation maximizing mechanism. Due to Theorem 3-1 the same revenue can be obtained by disclosing all information to the agents even when arrivals are unobservable. ■

Proof of Corollary 5. First, note that the revenue in the optimal mechanism where agents only buy upon arrival is lower than the revenue in the optimal mechanism when

buyers are long-lived and arrivals are observable. This holds since any mechanism that is implementable with short-lived agents is also implementable with long-lived agents and yields the same revenue by Theorem 10. By Theorem 3 this revenue is unchanged if arrivals are unobservable and if the process is Markov. ■

References

[1977] Albright S.C.: "Optimal Sequential Assignments with Random Arrival Times", *Management Science* **21** (1), 60-67.

[2013] Athey, S. and Segal, I.: "An Efficient Dynamic Mechanism", *Econometrica* **81** (6), 2463-2485.

[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

[2011] Babaioff, M., Blumrosen, L., Dughmi, S. and Singer. Y.: "Posting Prices with Unknown Distributions," *Innovations in Computer Science* (ICS 2011).

[2012] Ben-Porath, E. and Lipman, B.: "Implementation with Partial Provability," *Journal of Economic Theory* **147**, 1689-1724.

[2010] Bergemann, D. and Välimäki, J.: "Efficient Dynamic Auctions", *Econometrica* **78**(2), 771–789.

[2012] Besbes, O. and Lobel, I.: "Inter-temporal price discrimination: Structure and Computation of Optimal Policies", discussion paper.

[2015] Board, S. and Skrzypacz, A.: "Revenue Management with Forward Looking Buyers", *Journal of Political Economy*, forthcoming .

[1993] Boshuizen, F. and Gouweleeuw, J.: "General Optimal Stopping Theorems for Semi-Markov Processes", *Advances in Applied Probability* **24**, 825-846.

[2004] Bull, J. and Watson, J.: "Evidence Disclosure and Verifiability," *Journal of Economic Theory* **118**, 1-31.

[2007] Bull, J. and Watson, J.: "Hard Evidence and Mechanism Design," *Games and Economic Behavior* **58**, 75-93.

[2010] Cavallo, R., Parkes, D. and Singh, S.: "Efficient Mechanisms with Dynamic Populations and Dynamic Types," discussion paper, Harvard University

[2008] Deneckere, R., and Severinov, S.: "Mechanism Design with Partial State Verifiability," *Games and Economic Behavior* **64**, 487-513.

[1978] Dolan, R.J.: "Incentive Mechanisms for Priority Queueing Problems." *Bell Journal of Economics* **9** (2), 421-436.

[2012] Escobari, D.: "Dynamic Pricing, Advance Sales and Aggregate Demand Learning in Airlines.", *Journal of Industrial Economics* **60**(4), 697–724.

[2003] Etzioni, O., Knoblock, C., Tuchinda, 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.

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

[2013] Garrett, D.: "Incoming Demand with Private Uncertainty," discussion paper, Toulouse School of Economics.

[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.

[2010a] Gershkov, A. and Moldovanu, B.: "Efficient Sequential Assignment with Incomplete Information", *Games and Economic Behavior* **68**(1), 144-154.

[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, Universiy of Bonn

[1986] Green, J. and Laffont, J.J.: "Partially Verifiable Information and Mechanism Design", *Review of Economic Studies* **53**, 447-456.

[2012] Kartik, N., and Tercieux, O.: "Implementation with Evidence," *Theoretical Economics* **7**, 323-355.

[2005] Kittsteiner, T., and Moldovanu, B.: "Priority Auctions and Queue Disciplines that Depend on Processing Time", *Management Science* **51**(2), 236-248.

[2004] Lavi, R. and Nisan, N.: "Competitive Analysis of Incentive Compatible On-line Auctions," *Theoretical Computer Science* **310(1)**, 159-180.

[2012] Li, J., Granados, N., and Netessine, S.: "Are Consumers Strategic ? Structural Estimation from the Air-Travel Industry", working paper, INSEAD

[2013] Mantin, B. and Rubin, R.: "Transaction Prices and Strategic Consumer Behavior: Empirical Evidence from the Airline Industry", discussion paper, University of Waterloo

[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

[2010b] Mierendorff, K.: "The Dynamic Vickrey Auction", discussion paper, University of Bonn

[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

[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] Stadje, 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.

[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.