

Efficient Dynamic Allocation with Strategic Arrivals

Alex Gershkov, Benny Moldovanu and Philipp Strack*

10.1.2015

Abstract

We analyze dynamic allocations in a model with uncertain demand and with unobservable arrivals. The planner learns along the way about future demand, but strategic agents, who anticipate this behavior, strategically choose the timing of their arrivals. We examine the conditions under which the complete information, dynamically efficient allocation is implementable, and characterize the necessary payments that control the ensuing allocative and informational externalities.

1 Introduction

We study the following continuous-time optimal stopping problem: A planner endowed with an indivisible object faces a sequence of agents who randomly arrive over time, according to a general counting process which allows correlations in arrival times. Agents are long lived, and each agent is privately informed about her value for an object and about her arrival time. Thus, private information is two-dimensional, and recall is possible. Both agents and planner discount the future, and the planner - who may not be perfectly informed about the nature of the arrival process - wishes to maximize the expected value of the allocation. We ask whether the complete information optimal stopping rule is implementable in the present informational setting via individually rational monetary transfers that do not depend on events posterior to the physical allocation ("online"). We also identify scenarios where the monetary scheme involves/does not involve transfers between the planner and agents that do not get the object.

The main innovation lies in the combination of several features:

*A previous version of this paper has been circulated under the title "Incentive Compatible Optimal Stopping". We wish to thank Ron Lavi, Konrad Mierendorff, Alessandro Pavan, Andy Postlewaite, Ron Siegel, Rakesh Vohra, Mike Whinston and Asher Wolinsky for helpful remarks. Moldovanu is grateful to the European Research Council, the German-Israeli Foundation and to the German Science Foundation for financial support. Gershkov is grateful to the Google Inter-University Center for Electronic Markets and Auctions, the German-Israeli Foundation and the Israel Science Foundation for financial support. 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.

1. 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, agents may strategically choose when to make themselves available for trade, and private information is two-dimensional.
2. The planner may not be aware of the nature of the arrival process, which allows here for correlations in arrival times. Thus, the planner is able to learn about future arrivals (and thus about future demand) from past arrivals.
3. Since our agents can optimally choose their arrival time, they will take into account the effect of their current actions on the planner’s belief about the future, which in turns affects the terms of trade and the probability of trade.

Besides the theoretical interest of extending the static mechanism design paradigm to a classical dynamic allocation problem, we see the main applications of our model and methods to dynamic pricing questions in situations where capacity is limited, demand is random and where agents can strategically choose the time of their purchases. The strategic effects of such timing decisions have been mostly neglected by the existing literature since the standard assumption is that buyers are short-lived (see review below). If agents’ arrivals are observable, it is well known that the dynamically efficient policy can be implemented by requiring payments equal to the imposed **allocative** externality. In contrast, our main new insights can be summarized as follows:

1. If arrivals are unobservable, allocative externality payments are not sufficient to induce agents to trade as soon as possible, leading to delay and inefficiencies (see Example 1). The reason is that, by choosing when to reveal themselves to the mechanism, strategic agents can manipulate the planner’s beliefs about the arrival process in a way that induces more advantageous terms of trade, e.g. lower prices. This happens when later arrivals make the designer more pessimistic about the future. In other words, an earlier arrival may provide valuable information about the arrival process (and hence about demand) by allowing the designer to charge the right kind of payment in the future. Therefore transfers that lead to truthful revelation of information must also take into account this additional **informational** externality. In order to deal with the informational externality, we introduce a non-negative subsidy that is paid irrespective of physical trading, and that is a decreasing function of arrival times¹. Together with the allocative externality payment discussed above, this subsidy implements the complete information efficient dynamic policy (Proposition 2). The subsidy has the flavor of an ”advanced booking” discount, but here the discount is independent of a physical transaction.
2. In contrast to the case where there is learning via past observed values (see Gershkov and Moldovanu [2009]), a positive implementation result can nevertheless be obtained here because of 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 before². But practical applications of this positive result hinge on a verifiability

¹Non-negativity is required by the individual rationality constraint.

²Green and Laffont [1986] pioneered the study of (static) mechanism design problems where particular deviations from truth-telling are not feasible.

assumption: the planner needs to identify those agents that falsely claim to be available for trade just in order to claim the subsidy. For example, such identification may work via proofs of liquidity or the possession of a certain necessary technology. When such identification is impossible it makes sense to restrict attention to mechanisms that couple monetary transfers to physical transactions (*winner-pays mechanism*). Proposition 7 shows that no individually rational winner-pay scheme can implement the efficient allocation in a natural setting where later arrival of some agent makes the designer more pessimistic about future arrivals.

3. Proposition 3 and Proposition 2 identify large and important classes of arrival processes - renewals with a known distribution of inter-arrival times, non-homogenous Poisson processes, as well as any pure birth process - where the complete information efficient policy can be implemented via winner-pays mechanism even if arrivals are not observable.

The paper is organized as follows: In Section 2 we present the model with two-dimensional private information. In Section 3 we illustrate the paper's main ideas in a simple, but typical example where the arrival process is known to be a renewal, but where the precise inter-arrival time distribution is not known. In Section 4 we introduce direct revelation mechanisms. In Section 5 we briefly describe the expected externality scheme that implements the efficient allocation for the case where values are private, but arrivals are observable. We then show that this scheme must be augmented by a subsidy for early arrivals in order to implement the efficient allocation in the general model with private information about both value and arrival time. Section 6 discusses the limitations imposed by requiring that monetary payments can only accompany actual physical transactions. Section 7 concludes. All proofs are relegated to an Appendix.

1.1 Related Literature

The statistical and operations research literature analyzed continuous time, dynamic allocation problems with long-lived agents and with recall. This strand includes, among others, Zuckerman [1988], Zuckerman [1986], Stadjé [1991], and Boshuizen and Gouweleeuw [1993]. In these models, the agents are non-strategic and hence the planner is perfectly informed about the arrival process. Monetary transfers are therefore not necessary in order to implement the efficient policy.

Bergemann and Valimäki [2010], Cavallo, Parkes and Singh [2010], Parkes and Singh [2003], Said [2012] construct generalizations of the VCG mechanisms for various environments where either the population or the available information changes over time. Athey and Segal [2013] extended the expected externality mechanisms ala Arrow-D'Aspremont-Gerard Varet to dynamic environments which in addition to efficient allocation also satisfies per-period budget balancedness. Roughly speaking, all these papers use various independence assumptions in order to stay within a private values framework, and sometimes need to use payments that are not necessarily connected to physical allocations, or that may depend on events that happen after the physical allocation has been completed ("offline").³ Mierendorff [2013] analyzes an independent, private value model where an

³Lavi and Nisan [2004] provided worst-case analyses of online auctions and compared their outcome with the optimal offline mechanisms.

agent's value for the object may change over time. He shows that the efficient allocation can be implemented by an online mechanism where only the winner of the object pays. 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.

Gershkov and Moldovanu [2010] analyze efficiency in continuous-time optimal stopping frameworks where the agents are short-lived (thus there is no recall) and where the planner has several heterogeneous objects. In a model with discrete time, Gershkov and Moldovanu [2009] and Gershkov and Moldovanu [2012] allow the planner to learn about the distribution of values from past observations. In that model as well agents are short-lived. Learning about values introduces direct informational externalities (i.e., interdependent values), and these papers illustrate that the efficient implementation of the complete information optimal dynamic policy is only possible under strong assumptions about the learning process.

A small strand within the revenue management literature considers the effect of customers that may strategically decide about their arrivals. 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, but restrict the monopolist seller to choosing two prices, independently of the past history of sales. 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 a known arrival processes where the complete information stopping policy is characterized by a constant threshold rule which is independent of the previous arrivals. Thus, strategizing in the time dimension and the ensuing learning - which are the focus of our paper - do not play a role there. Board and Skrzypacz [2010] derive the revenue maximizing mechanism when arrival times are independent, i.e. the arrival process is Poisson, and thus learning does not play a role. In a recent paper Gershkov, Moldovanu and Strack [2014] analyzed a problem of revenue maximizing monopolist that faces a stream of random arrival of consumers. The arrivals are governed by a Markov process (a class of arrival processes that includes Poisson arrival process with unknown rate). Similarly to the present model, in their model recall may be part of the optimal mechanism.

Pai and Vohra [2013] and Mierendorff [2010] analyze revenue maximization in a discrete time, finite horizon framework where the arriving agents are 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.

Besides dynamic pricing, there are other important settings where strategic timing decisions affects learning and implementability of optimal policies. These include: 1) Financial markets where specialist market makers face sequences of traders, some of them better informed than the market. The possibility of strategic timing of trading decisions differentiates two classical models of the market micro-structure literature, Glosten and Milgrom [1985] and Kyle [1985]. 2) Gradual implementation of monetary policy. There investors who anticipate how their behavior changes future monetary policy will strategi-

cally choose investment times, sometimes completely defeating the purpose of the planned policy- see for example the model of Caplin and Leahy [1996]. 3) In another type of models, the agents themselves learn over time by observing others⁴. For example potential investors in an infant industry may want to delay their investment decisions until the market conditions become less uncertain. This kind of behavior may hinder the implementation of a targeted industrial policy. Rob [1991] studies subsidies to early entrants in infant industries and shows that they can restore efficiency by correctly adjusting the informational externality.

2 The Model

A designer endowed with an indivisible object faces a stream of randomly arriving agents in continuous time. The agents' arrivals are described by a 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 ⁵. The time horizon is potentially infinite, but the framework is rich enough to embed the finite horizon case by considering arrival process where after some time \bar{T} no more arrivals occur, i.e., where $\mathcal{N}(t)$ is constant for any $t \geq \bar{T}$.⁶ Since arrivals are described by general counting processes, the designer's beliefs about future arrivals may evolve over time, and may depend on the number of past arrivals and their exact timing.

Each agent has two-dimensional private information: the arrival time $t \geq 0$ and the value $v \geq 0$ he gets if allocated the object. That is, the designer does not observe agents' arrivals. If the agent arrives at time t , gets the object at time $\tau \geq t$ and pays p at time $\tau' \in [t, \tau]$, then her utility is given by $e^{-\delta\tau}v - e^{-\delta\tau'}p$ where $\delta \in (0, 1)$ is the discount factor. We implicitly assume here that all agents disappear after the allocation of the object, i.e., payments cannot be conditioned on information that arrives after the sale. Moreover, we assume that the item cannot be reallocated after an initial assignment.

The agents' values v_i are I.I.D. distributed with common c.d.f. F , on the support $[0, a]$ where $a \leq \infty$. We assume that F admits a continuous density f and has finite mean and variance. 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. Allowing for correlation between the values and the arrival times introduces a correlation between the agents' values, which generates further complications for implementing the efficient allocation, as was illustrated in Gershkov and Moldovanu [2009].

If the object is allocated to the agent with type (t, v) at time $\tau \geq t$, the designer's utility is given by $e^{-\delta\tau}v$. The designer's goal is to implement the efficient allocation (that maximizes his discounted expected utility) in a Bayes-Nash equilibrium.

⁴See Chamley [2004] for a good survey.

⁵Most textbooks on stochastic processes discuss the construction and properties of counting processes. See for example Ross [1983].

⁶Since the designer is interested in implementing the efficient allocation, he never allocates the object strictly after the last arrival in a finite horizon model.

2.1 The Complete Information Case

Let us briefly consider the benchmark case where the designer observes the agents' arrivals and their values for the object, so that agents have no private information. Our environment is then equivalent to the standard continuous-time, infinite-horizon search model with perfect recall. Since the main focus here is on the implementation of the efficient dynamic allocation (or, equivalently, the implementation of the optimal stopping policy), we assume that an optimal stopping time in the complete information model exists, and is almost surely finite.

The optimal stopping policy is deterministic. If the planner allocates the object at some time T , then he will allocate it to the agent with the highest value that arrived until T . Denote by $X(T)$ the highest value observed until time T (if until T no agent arrived, we set $X(T) = 0$), and by $\mathbf{t}_{\mathcal{N}(T)} = (t_1, \dots, t_{\mathcal{N}(T)})$ the agents' arrival times until T . Since values are independent of arrival times, the state of the process at T - on which the stopping policy depends - can be taken to be $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$.

Optimal policies in our framework have the following property. A stopping policy satisfies the *instant reservation price (IRP)* property if for any time T and for any history of arrivals $\mathbf{t}_{\mathcal{N}(T)}$, stopping at state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ implies stopping also at all states $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ with $X'(T) \geq X(T)$.

Lemma 1 *The optimal stopping policy in the complete information case satisfies the IRP property. In particular, for any time T and for any history of arrivals $\mathbf{t}_{\mathcal{N}(T)}$, there exists a cutoff $v_T^*(\mathbf{t}_{\mathcal{N}(T)})$ such that it is optimal to stop search as soon as the highest available value exceeds this cutoff.*

3 An Illustration of the Main Ideas

In our motivating illustration, the arriving process is an unknown counting process where the designer learns the precise distribution of the inter-arrival times after the first arrival⁷. We illustrate below the main difficulty the planner faces when trying to implement the first-best allocation rule.

Example 1 *Assume that the inter-arrival times are I.I.D. The designer believes that all arrivals distribute either uniformly on the interval $[1, 2]$, or uniformly on the interval $[2, 3]$ and assigns equal probabilities to each alternative. The distribution of the agents' values is, in both cases, uniform on the interval $[0, 1]$.*

3.1 The Complete Information Case

In order to characterize the complete information, dynamically efficient policy we use a result by Zuckerman [1988] who analyzed an infinite horizon, complete information model where the arrival process is a *renewal*, i.e., inter-arrival times are I.I.D. random variables with known, common distribution G .⁸ Note that, although inter-arrival times

⁷Even if the designer is not completely informed, the arrival process is still a counting process.

⁸The Poisson process is a special case where G is exponential (see Mortensen [1986] for the analysis of that case)

are independently distributed random variables, arrival times are correlated in a renewal process. In particular, the planner's beliefs about the timing of the next arrival depend on the elapsed time since the last arrival.

Contrary to what one may expect from the discrete case with deterministic arrivals, or from the Poisson process case, the designer may not wish to allocate the object immediately upon arrival, i.e., the recall option may be used by the optimal stopping policy⁹. Nevertheless, Zuckerman identified a large class of inter-arrival distributions G for which the optimal policy never employs the recall option, and is therefore characterized by a reservation value such that the object is allocated to the first arrival whose value is above the reserve¹⁰.

Definition 1 *A non-negative random variable W is called NBU (new better than used) if, for every $y > 0$, W is stochastically larger than the conditional random variable $(W - y/W \geq y)$.*¹¹

Theorem 1 (Zuckerman, [1988]) *Assume that the inter-arrival distribution G satisfies the NBU property, and let ϕ denote its Laplace Transform. Then, the optimal stopping policy allocates the object to the first arrival whose value is above v^* where v^* is the unique solution to*

$$v^* = \frac{\phi(\delta)}{1 - \phi(\delta)} \int_{v^*}^{\infty} (v - v^*) dF(v).$$

In particular, recall is never used by the optimal policy, and all allocations occur upon arrival.

The intuition is as follows: between arrivals, the seller updates her belief about the timing of the next arrival, and about the option value of not allocating the object right now. If the inter-arrival time satisfies the *NBU* property, the seller is most pessimistic about the timing of the next arrival immediately following an arrival, and gets more and more optimistic about it as time passes without arrivals. Thus, if it is optimal not to allocate the object immediately following an arrival - because the current option value of waiting is higher - it will not be optimal to do so until the next arrival.

Since each of the two possible distributions in our illustration satisfies the *NBU* property, Theorem 1 implies that the optimal complete information policy for the case where the designer knows the relevant distribution of inter-arrival times is such that search stops upon the arrival of the first agent whenever the value of that agent exceeds some fixed, time-independent cutoff, denoted by $x_{[1,2]}(\delta)$ and $x_{[2,3]}(\delta)$, respectively. In the Appendix we prove the intuitive assertion that

$$x_{[1,2]}(\delta) > x_{[2,3]}(\delta).$$

⁹For example, consider a process where times between consecutive arrivals can be either ε or Δ where $\varepsilon \ll \Delta$. Assume that a buyer with a moderately high value arrives at t . Then, for not too low discount factors, it may be optimal to wait until $t + \varepsilon$ hoping for a new arrival with a higher value, but then immediately stop search at $t + \varepsilon$ while recalling the previous buyer if no arrival occurred (because now the next arrival is known to be at the much more distant $t + \Delta$).

¹⁰Thus, the efficient policy coincides with the one obtained for renewal processes without recall by Albright [1974].

¹¹Note that any random variables with an increasing hazard (or failure) rate is *NBU*.

Therefore, in the case where the designer observes the agents' types but does not know the inter-arrival distribution, the dynamically efficient policy is given by ¹²:

1. For $T \in [1, 2]$ the cutoff is $x_{[1,2]}(\delta)$
2. For $T \in (2, 3]$ the cutoff is $x_{[2,3]}(\delta)$ if there were no arrivals before time 2, otherwise the cutoff is $x_{[1,2]}(\delta)$
3. For $T > 3$, the cutoff is $x_{[1,2]}(\delta)$ if the first arrival happened during time interval $[1, 2]$, whereas the cutoff is $x_{[2,3]}(\delta)$ if the first arrival happened during $(2, 3]$.

3.2 The Incomplete Information Case

We now show that a standard expected externality payment scheme a la Vickrey-Clarke-Groves (see also the general scheme defined in Proposition 1 below) generates incentives for the agents to misrepresent their arrival times. They do so in order to influence the terms of trade via the designer's beliefs about the arrival process. Therefore, such payments - that only deal with the allocative externality imposed by an agent that obtains the object, but that do not take into account the informational externality - cannot implement the complete information efficient allocation constructed above.

In order to calculate the relevant externality payment note that the object is allocated upon arrival in the complete information optimal policy (i.e., "recall" is not employed). Hence, by the definition of the optimal cutoff, the designer's continuation value at any time T after the first arrival t_1 must equal the optimal relevant cutoff at the time of the first arrival. Thus, the allocative externality payment, which needs to be paid by an agent who arrives at $T \geq t_1$ and obtains the object, is given by

$$P(t_1) = \begin{cases} x_{[1,2]}(\delta) & \text{if } t_1 \in [1, 2] \\ x_{[2,3]}(\delta) & \text{if } t_1 \in (2, 3] \end{cases} . \quad (1)$$

Given such payments, consider a type (t, v) with $t \in (1, 2)$ and $v \in (x_{[2,3]}(\delta), x_{[1,2]}(\delta))$. Truthful reporting yields utility zero since the object is not allocated to this agent. But, a report of arrival at time $t' = t + 1 \in (2, 3)$ together with a truthful report in the valuation dimension yields utility $e^{-\delta t'}(v - x_{[2,3]}(\delta)) > 0$ ¹³. Hence truthful reporting is not optimal.

3.3 A Subsidy Scheme

For any arrival time t' we define now a subsidy that is paid to an agent that arrives at t' , independently of whether this agent obtains the object or not:

$$S(t') = \begin{cases} x_{[1,2]}(\delta) - x_{[2,3]}(\delta) > 0 & \text{if } t' \in [1, 2] \\ 0 & \text{if } t' \geq 2 \end{cases} . \quad (2)$$

The above scheme subsidizes early arrivals occurring in the time interval $[1, 2]$. We now show that a combination of the externality payment made by the winner (defined in (1))

¹²Since no agents should arrive at $T \in [0, 1)$ the cutoff can be specified arbitrarily up to $T = 1$.

¹³Any report $t' \in (2, 3)$ yields a positive utility with a positive probability.

and the subsidy scheme in (2) - which deals with the informational externality- does implement the complete information dynamically efficient allocation.

An agent with type (t, v) where $t > 2$ cannot gain by misrepresenting his type if all other agents report truthfully. Therefore, it is sufficient to show that an agent with type (t, v) where $t < 2$ does not want to misrepresent his type. There are two cases:

1. If an agent with type (t, v) where $t < 2$ and $v \geq x_{[1,2]}(\delta)$ reports an arrival at $t' > 2$, the price of the object is indeed reduced, but the subsidy is also reduced by exactly the same amount, yielding no gain. A report such that the object is never obtained cannot be profitable either.

2. An agent with type (t, v) where $t < 2$ and $v < x_{[1,2]}(\delta)$ cannot gain by misrepresenting his type: getting the object at some time $t' < 2$ requires paying a price above value; reporting a later arrival reduces the price below value, but also reduces the subsidy to zero, which yields an overall decrease in that agent's expected utility.

Finally, consider a designer who seeks to use only payment schemes where agents that do not get the object pay nothing. We show now that there is no payment scheme in this class that implements the efficient allocation. To see this, note that if there were no arrivals until time $t = 2$, then the principal needs to charge a price $x_{[2,3]}(\delta)$ to the agent that gets the object. But this implies that an agent i that arrives at time $t \in [1, 2]$ with value $v \in [x_{[2,3]}(\delta), x_{[1,2]}(\delta))$ should get a strictly positive expected utility, since he can postpone his arrival and get the object for a price below value. Therefore, a mechanism where only the agent that gets the object makes a payment requires that such an agent i gets the object with a positive probability. But this contradicts efficiency.

The above example shows that implementation of the welfare maximizing policy is impossible without subsidizing agents who do not get the object for some arrival process which is a renewal with unknown interarrival distribution. This negative result does, however, not depend on the uncertainty about the arrival process. In section 6.1 we construct an example where no winner-pay mechanism implements the welfare maximizing allocation with unobservable arrivals. In this example the process is a renewal process with known interarrival distribution, but the principal faces a deadline before which he has to allocate the object (see Example 2).

4 Direct Revelation Mechanisms

We now come back to the general, incomplete information model. Without loss of generality (see Myerson [1986])¹⁴ we can restrict attention to mechanisms where the agents do not observe the history, and only know whether the object is still available or not.¹⁵ Since arrivals are unobservable, without loss of generality we can restrict attention to direct 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

¹⁴Although 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].

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

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.¹⁶

Since recall may be employed by the optimal policy, an allocation to an agent can be conditioned also on information that accrues between the arrival of that agent and the allocation time. We denote by $\eta(T)$ a history at time T : this is a list of reported arrivals and the reported values up to time T . Then

$$\mathcal{H}_T = \prod_{\mathcal{N}(T)=0}^{\infty} [0, T]^{\mathcal{N}(T)} \times \mathbb{R}_+^{\mathcal{N}(T)}$$

is the set of all possible histories at time T . We denote by h a history from the beginning of the game (time zero) until infinity, i.e., $h = \lim_{T \rightarrow \infty} \eta(T)$.

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 . An incentive compatible mechanism is (ex-post) individually rational if the utilities of all agents in the truth-telling equilibrium are non-negative.

Since incentive compatibility considers possible deviations by only one agent, it will be helpful to introduce additional notation. Let h_{-i} be the history (from the beginning of game until infinity) formed by agents' reports other than i , and let $\eta_{-i}^h(t)$ denote the derived history up to time t formed by the reports of agents other than i . We denote by μ the measure on histories generated by the counting process, and by $\mu(h_{-i}/t)$ the conditional measure given an arrival of agent i at time t .

Denote by $\tau_v(t, h_{-i})$ the optimal stopping time if agent i arrives at time t and reports value v while the other agents form history h_{-i} . Recall that, by the standard definition of stopping times, if $\tau_v(t, h_{-i}) = T$, then $\tau_v(t, h'_{-i}) = T$ for any h'_{-i} that agrees with h_{-i} up to time T . In other words $\tau_v(t, h_{-i})$ depends on h_{-i} only via $\eta_{-i}^h(\tau)$. We denote by $H_{-i}(t, v)$ the set of histories h_{-i} such that in the efficient allocation agent i gets the object if he reports type (t, v) , and we denote by $H_{-i}^c(t, v)$ its complement.¹⁷

5 A Subsidy Scheme for Early Arrivals

Let us first consider the base-line case where values are private information, but arrival times are observable. Then the dynamically efficient allocation is implementable by a mechanism where each agent pays the expected externality he imposes on the other current and future agents. The mechanism described below is a variant of the standard dynamic Vickrey-Clarke-Groves mechanisms - dynamic Pivot mechanism - that was characterized by Bergemann and Valimaki [2010]. We take its construction and properties - including the properties of the underlying optimal stopping policy - as given primitives for the sequel.

Denote by $V(t, v, \eta_{-i}(t'), t')$ the designer's expected utility at time t' when agent i arrives at time $t \leq t'$ and has value v while the other agents' types form history η_{-i} , and

¹⁶The 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.

¹⁷If arrivals are not private information we define $H_{-i}(v)$ and $H_{-i}^c(v)$ analogously.

when the designer uses the optimal stopping rule. Observe (also for later uses) that

$$V(t, v, \eta_{-i}(T), T) \begin{cases} = X(T) & \text{if } X(T) \geq v_T^*(\mathbf{t}_{\mathcal{N}(T)}) \\ > X(T) & \text{if } X(T) < v_T^*(\mathbf{t}_{\mathcal{N}(T)}) \end{cases} .$$

Let $P(t, v, \eta_{-i}(T), T)$ denote the payment charged at time T to agent i with value v if he arrived at time t and other agents' reports form history $\eta_{-i}(T)$.

Proposition 1 *Assume that values are private information, but arrival times are observable. The payment scheme*

$$P(t, v, \eta_{-i}^h(T), T) = \begin{cases} V(t, 0, \eta_{-i}^h(T), T) & \text{if } T = \tau_v(t, h_{-i}) \text{ and if } v_i = X(T) \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

implements the dynamically efficient allocation policy¹⁸. The resulting mechanism is ex-post individually rational.

5.1 Unobservable Arrivals

Let us now remove the assumption that the designer observes the arrival times of the agents. Note that a mechanism implementing the efficient dynamic allocation needs here to be individually rational. Otherwise, agents may choose never to show up.

Our illustration above suggests that the allocative externality payment scheme has to be adjusted: intuitively, early arrivals need to be subsidized since they create a positive externality by enabling the designer to learn about the nature of the arrival process.

We now construct a subsidy $S(t') \geq 0$ that is paid to all agents upon their reported arrivals, and that depends only on their reported arrival times. In this construction, we use the physical nature of the arriving process: each agent can only deviate in one direction, by claiming an arrival time later than the true one. Let

$$U(t, t', v) = \max_{v'} \int_{H_{-i}(t', v')} e^{-\delta \tau_{v'}(t', h_{-i})} [v - V(t', 0, \eta_{-i}(\tau_{v'}(t', h_{-i})), \tau_{v'}(t', h_{-i}))] d\mu(h_{-i}/t)$$

be the utility of an agent with type (t, v) - net of any payments made whenever he does get the object - when he reports (t', v') such that v' is chosen optimally given the reported arrival time t' and the type (t, v) . This construction is solely based on the properties of the complete information optimal stopping policy.

Under a very mild Lipschitz condition about the variation of $U(t, t', v)$, our next Proposition shows that, for any finite time T , it is possible to implement the efficient allocation if the object is allocated before T . In particular, since the optimal stopping time is assumed to be almost surely finite, efficiency can be approximated by taking T arbitrarily large.

¹⁸We assumed that agents do not observe the history prior to their arrival. It is easy to see that, in the framework with observable arrivals, the payment scheme described above implements the efficient allocation even if the prior history - consisting of arrivals and reported values of the agents that arrived beforehand - is observable to the agents.

Proposition 2 *Assume that there exists $M \geq 0$ such that for any $t \leq t' \leq t''$ in an interval $[0, T]$, and for any v it holds that*

$$U(t, t'', v) - U(t, t', v) \leq M(t'' - t')$$

Then the subsidy $S(t') = e^{\delta t'} M(T - t')$ together with the payment scheme given in Proposition 1 implements the dynamically efficient allocation for any history h where the optimal stopping time is less than T .

Remark 1 1. *If $U(t, t', v)$ is decreasing in t' , then the above condition is satisfied with $M = 0$, and a subsidy $S(t') \equiv 0$ implements the efficient allocation for any history, i.e. $T = \infty$. Proposition 3 below displays (in terms of the model's primitives) an important class of processes where $U(t, t', v)$ is indeed decreasing in t' . Note that such a condition plays here a similar role to the well known single-crossing condition in static allocation problem with interdependent values. Indeed, $U(t, t', v)$ can be also seen as the expectation of the difference between the value of allocating the object to a particular agent (say agent i) and the externality imposed by that agent in the efficient allocation (represented by the option value function V in the definition of U). In the static case, single crossing requires that this difference is monotone in i 's signal.*

2. *Although the above results rely on a Lipschitz condition or on differentiability in reported arrival t' of the expected utility $U(t, t', v)$, an analogous reasoning - choose S such that the function $U(t, t', v) + e^{-\delta t'} S(t')$ decreases in the reported arrival time t' for any t and v - can be used even if the function $U(t, t', v)$ is not even continuous. Recall, for example, the setting of Example 1, where we illustrated such a case.*
3. *If time is discrete and if the interval where values distribute is bounded, then the above result holds **without any** assumptions on the function $U(t, t', v)$. In the Appendix we show how the minimal subsidy for this case can be constructed by a simple backward induction argument.*

The above rather permissive result is in stark contrast with the restricted possibility result obtained by Gershkov and Moldovanu [2012] where there the planner dynamically learns about the valuations of future agents. Although both models of learning - about future values or about future arrivals - create informational externalities and hence obstacles to implementation, the difference lies in the special nature of arrival times where only one-directional deviations are feasible.

6 Winner-Pay Mechanisms

As mentioned in the Introduction, the applicability of the above solution is restricted for at least two reasons: 1) It requires the designer to have "deep pockets". In other words, the implementation of the efficient allocation may be very costly, disproportional to the benefits from the efficient allocation itself. 2) It generates incentives for agents that have no interest in the object to "arrive" just in order to collect the subsidy. Therefore, if it

is not possible to physically identify "fake" arrivals, it makes sense to restrict attention to the mechanisms where a transfer of money takes place only between the planner and the agent that gets the object. Individual rationality requires here that agents who do not get the object do not make any payments.

Definition 2 *A mechanism is called winner-pay mechanism if the transfers to all agents that do not get the object are zero.*

Our next results identify two important classes of stochastic processes where the efficient dynamic allocation can be implemented via winner-pay mechanisms even if arrivals are unobservable. Thus, identifying fake arrivals is not an issue in such frameworks.

6.1 Renewal Processes

We first look at renewals, i.e., at processes where the inter-arrival times are I.I.D. random variables with a known, common distribution G . For this class of problems, the existence of an efficient policy and a characterization in terms of the infinitesimal generator is established in Boshuizen and Gouweleeuw [1993].

Recall that, as we mentioned earlier, although inter-arrival times are independently distributed random variables, arrival times are correlated in a renewal process. In particular, the planner's beliefs about the timing of the next arrival depend on the elapsed time since the last arrival. Thus, past arrivals do create an informational externality which vanishes only in the special case of a homogenous Poisson process. Hence, in such processes the planner continually adjusts her beliefs about the timing of the next arrival based on the elapsed time since the last arrival - thus past arrivals do create an informational externality.

Denote by \mathbb{T} the elapsed time since the arrival of the last agent. Recall that $X(T)$ denotes the highest value observed until time T , and note that in a renewal process the bivariate process $(X(T), \mathbb{T}(T))$ is Markov. Therefore, the efficient cutoffs v_T^* can be characterized only in terms of \mathbb{T} , the time since the last arrival.

Proposition 3 *Assume that arrivals are unobservable, that the arrival process is a renewal with inter-arrival distribution G , and that the horizon is infinite.*

1. *The payment for the object given in (3) implements the dynamically efficient allocation. In other words, a subsidy is not needed for efficient implementation.*
2. *Assume that the inter-arrival distribution G satisfies the NBU property, and let ϕ denote its Laplace Transform. Then, charging for the object the constant price $P(t, v, \eta_{-i}(T), T) = v^*$ where v^* is the unique solution to*

$$v^* = \frac{\phi(\delta)}{1 - \phi(\delta)} \int_{v^*}^{\infty} (v - v^*) dF(v)$$

implements the efficient allocation.

The above Proposition shows that if the arrival process is a renewal, implementation is possible even if the inter-arrival distribution is not NBU, and hence even if recall is actively employed in the optimal stopping rule. This permissive result assumed that there is no deadline before which the object has to be allocated. The following example shows that if the arrival process is a renewal but there is a deadline, implementation without a subsidy might be impossible.

Example 2 *Suppose the arrival time of the first agent a_0 is uniformly distributed on $[0, 2]$. Furthermore, assume that arrivals are governed by a renewal process where the next agent arrives one unit of time later*

$$a_{i+1} - a_i = 1.$$

Assume that the designer has to allocate the object before time 2. Thus, with probability 1 at most two agents arrive. The agents' valuations v_0, v_1 are uniformly drawn from the interval $[0, 1]$.

If the first agent arrives at a time $\mathbf{t}_1 \leq 1$ the welfare maximizing policy allocates the object to him if and only if

$$v_1 \geq \mathbb{E}[e^{-r} \max\{v_1, v_2\}] = e^{-r} \int_0^1 \max\{v_1, z\} dz = e^{-r} [v_1^2 + (1 - v_1) \frac{v_1 + 1}{2}] = \frac{e^{-r}(1 + v_1^2)}{2}.$$

Let $v^ = e^{-r}(1 - \sqrt{1 - e^{-2r}})$ be the smallest valuation such that the first agent gets the object upon arrival. If the first agent does not get the object, the welfare maximizing policy waits for the second agent and allocates the object to the first agent whenever the valuation of the second agent is below the first agent's valuation.*

Restrict attention to winner-pay mechanisms where the agent who does not get the object makes no payment. Let p_i be the price paid by agent i if he gets the object. As agent i receives no payment if he does not get the object, it follows that he faces a price of v_{-i} if agent 1 arrives before time $\mathbf{t}_1 = 1$ and has a value below v^ . As such an agent with valuation v^* needs to be indifferent to reporting a marginally lower valuation, it follows that he needs to face a price p^* that solves:*

$$v^* - p^* = e^{-r} \mathbb{E}[\max\{v^* - v_2, 0\}] = e^{-r} v^* (v^* - \frac{v^*}{2}) = e^{-r} \frac{(v^*)^2}{2} \Rightarrow p^* = v^* - e^{-r} \frac{(v^*)^2}{2}.$$

To implement the efficient allocation in a winner-pay mechanism with observable arrivals the principal thus needs to use the prices

$$p_1 = \begin{cases} p^* & \text{if } \mathbf{t}_1 \leq 1 \text{ and } v_1 \geq v^* \\ v_2 & \text{if } \mathbf{t}_1 \leq 1 \text{ and } v_1 < v^* \\ 0 & \text{if } \mathbf{t}_1 > 1 \end{cases}$$

$$p_2 = v_1.$$

Consider now the situation where arrivals are unobservable, and where the first agent arrives at time one with a valuation $v_1 < v^$. If he reports his arrival at any later time, he gets the object earlier, more often and for a lower price. Thus, the efficient allocation can not be implemented with unobservable arrivals in a winner-pay mechanism.*

6.2 Pure Birth and Non-homogenous Poisson Processes.

A pure birth process is a counting process such that the instantaneous probability that an agent arrives (arrival rate) only depends on the number of past arrivals and on calendar time. A special member of this class of arrival processes is the non-homogenous Poisson process where the arrival rate is a function of calendar time only. Our first Proposition establishes the existence of a welfare maximizing policy.

An optimal policy that stops almost surely in finite time may not generally exist because the designer might get more and more optimistic over time. In order to prove existence we need to bound the expected discounted number of arrivals. The proof of existence is somewhat involved: it uses a combination of concepts and results from majorization theory, and from the theory of order statistics.

Proposition 4 *Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_n(t)$ such that there exists a non-decreasing function $\beta(t)$ with $\lambda_i(t) \leq \beta(t)$ for all i, t , and such that $\int_0^\infty \beta(t)e^{-\delta t} dt < \infty$. Then there exists an optimal stopping rule for the setting with observable arrivals.*

Throughout the rest of the section we are going to maintain the assumptions of Proposition 4 to ensure existence of an optimal policy. Our second proposition shows that, if the arrival process is a pure birth process, no subsidy is necessary to implement the welfare maximizing allocation even if arrivals are unobservable.

Proposition 5 *Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_{N(T)}(T)$. Charging the payment P defined in Equation 3 which implements the efficient allocation with observable arrivals also implements the efficient allocation when arrivals are unobservable.*

The proof of the above result relies on an theorem proven in Gershkov, Moldovanu and Strack [2014] which establishes that if the policy is monotone in arrival times, the process is a pure birth process and agents with a valuation of zero get a transfer of zero, than any policy which is implementable with observable arrivals is implemented by the same payment with unobservable arrivals. We next derive several more properties of the payment implementing the efficient policy.

Proposition 6 *Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_n(T)$ such that λ is non-decreasing (non-increasing) in both elapsed time T and the number of arrivals $n = N(T)$ up to T . Then the price charged to the agent for the object is non-decreasing (non-increasing) over time.*

As the arrival rate increases (decreases) over time it follows that the continuation value of the planner increases (decreases) over time. The result follows since the price charged to the agent equals the continuation value of the planner

Consider now the special case where the arrival rate is increasing over time. In this case, if the welfare maximizing allocation gives an object to an agent at time t it also allocates the object to an agent with the same valuation at any earlier point in time, i.e. no recall is used. As no recall is used the continuation value (and thus the price) equals the lowest valuation with which an agent would get the object, and we obtain the following result.

Corollary 2 *Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_n(T)$ such that λ is non-decreasing in both elapsed time T and in the number of arrivals $n = \mathcal{N}(T)$ up to T . Then the price charged to the agent at time T equals the cut-off $v_T^*(\mathbf{t}_{\mathcal{N}(T)})$, and no recall is used.*

The above result establishes a class of problems where recall is never used. This allows us to construct an explicitly calculated example by using known results about the optimal policy for the case where recall is impossible.

Example 3 *Assume that the distribution of values is exponential $F(v) = 1 - e^{-v}$, and consider a non-homogenous Poisson arrival process with rate $\lambda(t) = \delta(t+2) \ln(t+2) - 1$. Observe that λ is positive, increasing in t , and that*

$$\int_0^\infty [\delta(t+2) \ln(t+2) - 1] e^{-\delta t} dt < \infty.$$

By results in Albright [1974] and Gershkov and Moldovanu [2010], the optimal cutoff in the optimal stopping problem where arriving agents are short lived (no recall) is given by the differential equation:

$$y' - \delta y = -[\delta(t+2) \ln(t+2) - 1] e^{-y}.$$

The solution $y(t) = \ln(t+2)$ increases in t , and satisfies $\lim_{t \rightarrow \infty} (y(t) e^{-\delta t}) = 0$ and $\frac{d(e^{-\delta t} \ln(t+2))}{dt} \leq 0$ for $\delta \geq \frac{1}{2 \ln 2} = 0.72$, and thus by arguments in Albright [1974], it is the optimal solution for such discount factors. By the argument of the above Proposition, charging $P(t) = \ln(t+2)$ implements the efficient dynamic allocation also in the problem with recall and with unobservable arrivals if the discount factor is high enough.

7 A General Negative Result

Contrasting the above special cases, our last result generalizes the insight at the end of Section 3, and has a negative flavor: it shows that the efficient allocation cannot be implemented via winner-pay mechanisms if a later arrival of some agent makes the designer more pessimistic about future arrivals.

This is in sharp contrast with the very permissive result shown in subsection 6.2: While for any Markov process the welfare maximizing allocation can be implemented without subsidies, such an implementation becomes impossible once future arrivals depend only slightly on the exact timing of arrivals.

Proposition 7 *Suppose, the buyers observe the same information as the seller. Consider an arrival process where the arrival rate at time t given that there were k arrivals before t , $t_1 \leq t_2 \leq \dots \leq t_k \leq t$, is given by a differentiable function $\lambda_k(t_1, t_2, \dots, t_k, t) > 0$. Assume that λ_k is strictly decreasing in t_1, t_2, \dots, t_k , k , bounded from below by $\underline{\lambda} > 0$ and from above by $\bar{\lambda}$, and that there is a finite deadline $T < \infty$ after which no agent arrives. Then, there is no winner-pay mechanism that implements the dynamically efficient allocation.*

Proof of Proposition 7. As there is a finite deadline T and as λ is bounded from above there is a strictly positive probability that only finitely many agent arrive. Each agent who at the time of his arrival has the highest valuation has a strictly positive probability of getting the object.

Consider first the policy that, at each period of time, allocates the object to the next agent who arrives, and denote by τ the allocation times of this policy. Let $\mu = \int w f(w) dw$ be the average valuation of a random buyer. As the arrival rate is strictly bounded away from zero, this policy guarantees a strictly positive continuation value. In particular, it is never optimal to allocate the object to an agent at time t if his valuation is below

$$v^*(t) := \mu \int_t^T \underline{\lambda} e^{-\delta s} e^{-\underline{\lambda} s} ds = \mu \frac{-\underline{\lambda}}{\delta + \underline{\lambda}} [e^{-(\delta+\underline{\lambda})s}]_{s=t}^{s=T} = \mu \underline{\lambda} \frac{e^{-(\delta+\underline{\lambda})t} - e^{-(\delta+\underline{\lambda})T}}{\delta + \underline{\lambda}}$$

As $v^*(t)$ is strictly positive, and as the support of valuations goes to zero, there exist histories after which an agent does not get the object immediately, but may get the object with a positive probability in the future. We now show that such an agent i can deviate from truth telling by reporting an arrival at the first period of time where he would get the object immediately.

Note that if the agent reports an arrival at the time $\tau_i(s, v_i)$ when he would have gotten the object were he to arrive at time s , then he gets the object immediately because, due to his later arrival, the continuation value of the principal is lower. The first arrival time a_i where agent i would get the object immediately upon arrival is a stopping time with respect to the information of the principal. Since the agent and the principal observe the same information, a_i is also a stopping time with respect to the information of the agent, and thus a feasible deviation in the arrival time dimension. Since payments cannot be made after allocating the object, revenue equivalence together with the fact that an agent with valuation of zero does not get the object¹⁹ uniquely pin down transfers. Thus, the information rent of an agent with value v_i is given by the integral over the expected discounted allocation times of lower types

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

As the principal gets more pessimistic if the agent arrives later, it follows that his continuation value is lower after every history following the deviation. The principal allocates the object to the agent whenever the agent's valuation exceeds the continuation value, and hence it follows that the agent gets the object strictly earlier whenever he arrives later, i.e. for all $s < a_i, z$, the time τ_i the object is allocated to agent i satisfies

$$\tau_i(a_i, z) \leq \tau_i(s, z).$$

Hence, the agent's information rent when arriving at time a_i satisfies

$$\begin{aligned} \mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(a_i, z)} dz \mid t_i = a_i, \eta_{-i}(a_i) \right] &\geq \mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(s, z)} dz \mid t_i = a_i, \eta_{-i}(a_i) \right] \\ &> \mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(s, z)} dz \mid t_i = s, \eta_{-i}(a_i) \right]. \end{aligned}$$

¹⁹Such an agent cannot get any payment due to the winner-pay restriction

where the second strict inequality follows because it becomes strictly more likely that another agent arrives when agent i arrived at time s before a_i . Using the law of iterated expectations and *Doobs Optional Sampling Theorem* we obtain that the agent's information rent is strictly higher if reports his arrival at time a_i instead of s :

$$\begin{aligned} & \mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(s,z)} dz \mid t_i = s, \eta_{-i}(s) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(s,z)} dz \mid t_i = s, \eta_{-i}(a_i) \right] \mid t_i = s, \eta_{-i}(s) \right] \\ &< \mathbb{E} \left[\mathbb{E} \left[\int_0^{v_i} e^{-r\tau_i(a_i,z)} dz \mid t_i = a_i, \eta_{-i}(a_i) \right] \mid t_i = s, \eta_{-i}(s) \right]. \end{aligned}$$

Consequently, it is a profitable deviation for agent i to report his arrival at time a_i instead of s . ■

8 Conclusion

We have analyzed dynamic allocations in a continuous time, discounted model where arrivals are governed by a general 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. With observable arrivals, the complete information, dynamically efficient policy can be implemented by an individually rational mechanism where only the winner of the object pays a price equal to the expected allocative externality. The same is true even when arrivals are not observable if the arrival process is a renewal, or a combination of pure-birth/non-homogenous Poisson process with increasing arrival rate. In general, controlling the informational externalities induced by the learning process calls for schemes where monetary flows are not tied to physical allocations. Such schemes are expensive and may create incentives for "fake" arrivals. We show that such schemes are indispensable in order to implement the complete information, efficient policy in situations if late arrivals induce pessimism about future arrivals. For practical applications, it is of interest to further study "second-best" policies in such environments.

In addition, our results offers insights for other settings where the designer sequentially learns from agents' actions, and where the optimization by agents (who understand how their actions affect the learning process) may have potentially undesirable consequences.

9 Appendix

Proof for Example 1. For inter-arrival times that distribute uniformly on $[1, 2]$ and on $[2, 3]$ the Laplace transforms are, respectively: $\phi_{[1,2]}(\delta) = \frac{e^{-\delta} - e^{-2\delta}}{\delta}$; $\phi_{[2,3]}(\delta) = \frac{e^{-2\delta} - e^{-3\delta}}{\delta}$. By the strict convexity (in z) of the function $e^{-\delta z}$ we obtain that

$$\forall \delta \in (0, 1), \phi_{[1,2]}(\delta) > \phi_{[2,3]}(\delta).$$

Under complete information, the optimal cutoff is given by a solution to the equation

$$(x^2 + 1) \phi_{[a,b]}(\delta) = 2x \tag{4}$$

where $\phi_{[a,b]}(\delta)$ is relevant Laplace transform. Since for any $\delta \in (0, 1)$ the Laplace transform lies in the interval $(0, 1)$, there exists a unique solution to equation (4) which lies in the relevant interval of agents' values $[0, 1]$. This solution is given by

$$x_{[a,b]}(\delta) = \frac{1}{\phi_{[a,b]}(\delta)} \left(1 - \sqrt{1 - (\phi_{[a,b]}(\delta))^2} \right)$$

which is monotonically increasing in $\phi_{[a,b]}(\delta)$. Hence, we obtain that

$$x_{[1,2]}(\delta) > x_{[2,3]}(\delta)$$

■

Proof of Lemma 1. Assume by contradiction that the optimal policy Υ does not satisfy *IRP*. Then there exist some time T , history of the arrivals $\mathbf{t}_{\mathcal{N}(T)}$ and $X'(T) > X(T)$ such that Υ prescribes to continue search at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$, while it prescribes stopping (and accepting $X(T)$) at state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$.

The expected utility under Υ at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ can be written as $X'(T)\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) + \beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)$ where $\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)$ is the discounted probability that the object will be allocated to the agent with value $X'(T)$ and where $\beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)$ is the discounted expected utility from all continuations where the object is not allocated to the agent with value $X'(T)$. The probability that the object will be allocated to the agent with value $X'(T)$ is less than one (otherwise it would be optimal to stop at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$) and therefore $\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) < 1$. Since Υ prescribes to continue search at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$, it must be the case that

$$X'(T)\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) + \beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) > X'(T)$$

which implies

$$\beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) > (1 - \alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)) X'(T).$$

Change now policy Υ into Υ' where the only difference between the policies is at state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$. Policy Υ' continues search at time T and applies the same continuation policy at state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ as Υ prescribed after the state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ (while allocating the object to the agent with value $X(T)$ if Υ prescribes to stop and to allocate the object to an agent with value $X'(T)$).

The expected utility generated by Υ' if state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ was reached is $X(T)\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) + \beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)$. Since we know that $\beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) > (1 - \alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)) X'(T)$ we obtain

$$\begin{aligned} X(T)\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) + \beta(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) &> \\ X(T)\alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon) + (1 - \alpha(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon)) X'(T) &> X(T). \end{aligned}$$

That is, there exists a continuation policy applied at state $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$ that generates a higher expected utility than $X(T)$. This contradicts the optimality of stopping at $\{X(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$. ■

Proof of Proposition 2. With a subsidy $S(t')$, the expected utility of agent i with type (t, v) who reports type (t', v') where $t' \geq t$ is given by

$$\begin{aligned} &\int_{H_{-i}(t', v')} e^{-\delta\tau_{v'}(t', h_{-i})} [v - V(t', 0, \eta_{-i}^h(\tau_{v'}(t', h_{-i})), \tau_{v'}(t', h_{-i}))] d\mu(h_{-i}/t) \\ &+ e^{-\delta t'} S(t'). \end{aligned}$$

Given truthful reporting in the arrival time dimension, the payment scheme of Proposition 1 provides incentives to report truthfully the value of the object. In other words, $v \in \arg \max_{v'} U(t, t, v)$ and thus $v \in \arg \max_{v'} [U(t, t, v) + e^{-\delta t'} S(t')]$. Thus, a function $S(t')$ such that $U(t, t', v) + S(t')$ is decreasing in t' for any $t \leq t'$ and v induces any agent to report the earliest possible arrival time, which necessarily coincides with the true arrival time. Therefore we can implement the efficient allocation.

With the proposed subsidy, the expected utility of an agent that arrives at time t with value v , but reports arrival time $t' \geq t$ and value v' (optimized given t') is given by

$$U(t, t', v) + M(T - t')$$

The result will be proved by showing that $U(t, t', v) + M(T - t')$ is decreasing in t' for any $v, t \leq t'$. Consider then $t'' \geq t' \geq t$. We obtain that

$$\begin{aligned} U(t, t'', v) + M(T - t'') &\leq U(t, t', v) + M(T - t') \Leftrightarrow \\ U(t, t'', v) - U(t, t', v) &\leq M(t'' - t') \end{aligned}$$

The last inequality holds by assumed Lipschitz condition. Individual rationality follows immediately by Proposition 1, and because $S(t') \geq 0$. ■

Proof of Proposition 3. 1. We show that $U(t, t', v)$ decreases in t' for any v and $t \leq t'$ (see Proposition 2 and Remark 1-1). By definition, for any $v \geq 0$, $U(t, t', v) \geq 0$, since reporting the true value guarantees for any t and t' a non-negative utility.

Consider an agent with true type (v, t) who reports type (v', t') where $t' > t$ and where v' is optimized given t' and (t, v) . Such report is relevant only if the agent gets then the object with positive probability, and if his expected utility is positive.

We claim that a report (v', t'') where $t \leq t'' < t'$ leads to a higher expected utility than a report (v', t') . Indeed, if an agent with report (v', t') gets the object at some time T , then he should get the object with a report (v', t'') , $t'' \leq t'$ either at T , or earlier. This is so because at the time of the allocation the elapsed time from the last arrival must be

the same, independently of the reported arrival time of that agent. The price charged for the object depends on the elapsed time since the last arrival, and on the second highest value reported up to the time of the allocation. Moreover, this price is monotonically increasing in the second highest reported value. Thus, a later arrival may postpone the allocation, which increases the probability of new arrivals, which in turn increases the second highest value and the charged price. Therefore, a report (v', t'') with $t'' < t'$ leads to a possible earlier allocation at a possibly lower price, and is thus a more advantageous deviation than a report (v', t') . Adjusting the reported value to the earlier arrival further increases expected utility, which allows us to conclude that $U(t, t', v) \leq U(t, t'', v)$, as desired.

2. The result follows immediately from point 1 above together with Theorem 1. ■

Construction of Minimal Subsidy. Assume that time is discrete, $t \in \{0, \dots, T\}$. As above, the payment scheme is constructed from two elements: the early arrival subsidy which is independent of the allocation of the object, and the expected externality payment given in (3). We derive the minimal subsidy by backward induction. Agents that report their arrival at the last period, gets a subsidy of zero. Consider an agent that arrives at the penultimate period, $T - 1$. Since the only deviation such an agent might have in the time dimension is reporting an arrival at T , the minimal subsidy must equal the maximal expected utility that can be obtained by reporting an arrival at T , where the maximum is taken over the values from getting the object. This amount is bounded if the interval on which values distribute is bounded. Such a subsidy induces all types of agents that arrive at period $T - 1$ to report truthfully their arrival times. The proof of the above Proposition shows that, given truthful reporting in the time dimension, the expected externality payment provides the correct incentives to report truthfully their values.

We now proceed one step further, to period $T - 2$. Given the above constructed subsidy for period $T - 1$, we need to consider deviations in the time dimension by agents arriving at $T - 2$ who can declare arrivals at either T or $T - 1$. The subsidy is the maximum between the two maximal expected utilities. By construction any lower subsidy will cause at least some value types of agents that arrive at $T - 2$ to misreport their arrival time. Notice that the construction is history independent because the calculations at any period are conditioned on the fact that the agent for whom the subsidy is calculated arrived at that period, and because the agent does not observe any previous arrivals.

The minimal subsidy for the continuous time case can be obtained from the above discrete time case by considering larger and larger sets of possible discrete arrivals. The bounded variation assumptions made in Proposition 2 ensure that the minimal subsidy stays finite in this limit process. ■

Proof of Theorem 5. We first show that an agent who arrives earlier gets the object earlier. As the assumed arrival process is Markov, the continuation value of a welfare maximizing planner who observes arrivals and valuations only depends on the number of prior arrivals and on calendar time. Hence, the threshold at which an object is allocated to an agent does not depend on his arrival time. Consequently, if the object is allocated to an agent after a history where he arrived at time t , it is also allocated to this agent after every history where he arrived earlier.

Second, we show that in the welfare maximizing allocation rule with payments P , an agent who has a valuation of zero does not receive any payment from the principal. Never allocating an object guarantees the principal a welfare of zero, and hence the continuation value of the welfare maximizing principal is always non-negative. The price charged by the principal equals the continuation value, and hence the price paid by the agent is non-negative.

Theorem 12 (page 24) in Gershkov, Moldovanu & Strack [2014] shows that any allocation rule which is monotone in arrival times, implemented by payments P under observable arrivals and makes no payment to agents with a valuation of zero is also implemented by the same payments under unobservable arrivals. Hence, P as defined in Equation 3 implements the welfare maximizing allocation under unobservable arrivals. ■

For the proof of Proposition 2 we first need the following concepts and results from majorization theory:

Definition 3 1. For any n -tuple $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ let $\gamma_{(j)}$ denote the j th largest coordinate (so that $\gamma_{(n)} \leq \gamma_{(n-1)} \leq \dots \leq \gamma_{(1)}$). We say that α is majorized by β and we write $\alpha \prec \beta$ if the following system of $n - 1$ inequalities and one equality is satisfied:

$$\begin{aligned} \alpha_{(1)} &\leq \beta_{(1)} \\ \alpha_{(1)} + \alpha_{(2)} &\leq \beta_{(1)} + \beta_{(2)} \\ &\dots \leq \dots \\ \alpha_{(1)} + \alpha_{(2)} + \dots + \alpha_{(n-1)} &\leq \beta_{(1)} + \beta_{(2)} + \dots + \beta_{(n-1)} \\ \alpha_{(1)} + \alpha_{(2)} + \dots + \alpha_{(n)} &= \beta_{(1)} + \beta_{(2)} + \dots + \beta_{(n)}. \end{aligned}$$

2. A function $\Psi : R^n \rightarrow R$ is called Schur-convex if $\alpha \prec \beta \Rightarrow \Psi(\alpha) \leq \Psi(\beta)$

Theorem 3 (Marschall and Proschan [1965]) Let $X = (X_1, \dots, X_n)$ be an n -dimensional random vector with a permutation invariant joint distribution. Let $\phi : R^n \rightarrow R$ be a continuous, convex function that is permutation invariant in its arguments. Then the function $\mathbb{E}\phi(\alpha_1 X_1, \dots, \alpha_n X_n)$ is Schur convex, i.e.,

$$\mathbb{E}\phi(a_1 X_1, \dots, a_n X_n) \leq \mathbb{E}\phi(\beta_1 X_1, \dots, \beta_n X_n).$$

whenever $\alpha \prec \beta$.

Proof of Proposition 4. 1. It is sufficient to show the existence of an optimal stopping policy for the process with arrival rate $\beta(t)$, since then the cutoff that induces stopping for the process with rate $\lambda_i(t)$ exists, and is bounded by the optimal cutoff corresponding to the arrival rate $\beta(t)$. To see this, observe that if the planner finds it optimal to stop under the arrival process with rate $\beta(t)$, then the value of continuing search is lower than the value of stopping. This implies that, under the same circumstances, stopping is the optimal action also under arrival rate $\lambda_i(t) \leq \beta(t)$.

In order to prove the existence of an optimal stopping rule for the rate $\beta(t)$, we consider two auxiliary problems where the planner also wants to maximize the expected value of the agents from the allocation:

Problem A: All agents arrive simultaneously, and their number is random. The probability that k agents arrive equals the probability that in the original dynamic problem with rate $\beta(t)$ the discounted number of the agents is between $k - 1$ and k .²⁰

Problem B: Arrivals are sequential with rate $\beta(t)$, but the planner is a "prophet" who observes at each point in time all future arrivals and their values.

We show first that the planner's expected utility in Problem A, with a distribution of the agents that is chosen to mimic that of Problem B, is finite. Afterwards, we show that Problem A generates a higher expected utility than Problem B. Obviously, Problem B generates a higher expected utility than the original problem with rate $\beta(t)$ since in the latter problem the designer has no information about the future. This will allow us to conclude that the designer's expected utility in the original dynamic problem is finite.

We now show that the planner's expected utility in Problem A is finite. The assumption $\int_0^\infty \beta(t)e^{-\delta t} dt < \infty$ implies that the expected discounted number of arrivals in the original dynamic problem with rate $\beta(t)$ is finite. Since Problem A is constructed such that the expected number of the agents mimics the expected discounted number of agents in Problem B (up to increasing the realized number of arrivals to the next integer) we obtain that the expected number of agents in Problem A is finite as well. More precisely, the expected number of arrivals in Problem A is given by

$$B = \sum_{k=0}^{\infty} kP(k) = \sum_{k=0}^{\infty} (k-1)P(k) + \sum_{k=0}^{\infty} P(k) \leq \int_0^\infty \beta(t)e^{-\delta t} dt + 1 < \infty$$

where $P(k)$ is the probability of k arrivals in Problem A.

Let $X_{(k)}$ be the highest order statistic out of k I.I.D random variables X_1, \dots, X_k with mean μ , representing the agents' values. Then the designer's expected utility in Problem A is given by $\sum_{k=0}^{\infty} P(k)\mathbb{E}[X_{(k)}]$.²¹ Since $\mathbb{E}[\sum_{i=1}^k X_i] = k\mu$, we know that $\mathbb{E}[X_{(k)}] \leq k\mu$. Hence

$$\sum_{k=0}^{\infty} P(k)\mathbb{E}[X_{(k)}] \leq \sum_{k=0}^{\infty} P(k)k\mu = B\mu < \infty$$

and the designer's expected utility in Problem A is finite.

b). We now show that for any realizations of arrival times in Problem B, the expected utility in this problem is lower than the expected utility in Problem A with a corresponding number of agents. That is, if in Problem B the arrival times are (t_1, t_2, \dots) , then the planner's expected utility is lower than that in Problem A with a number of agents given by $\lceil \sum_{i=1}^{\infty} e^{-\delta t_i} \rceil$ where $\lceil a \rceil$ denotes the lowest integer greater or equal to a .

For a given realization of arrival times, the planner's expected utility in Problem B is given by $\mathbb{E}[\max\{\alpha_1 X_1, \alpha_2 X_2, \dots\}]$ where $\alpha_i = e^{-\delta t_i}$ and where t_i is the arrival time of the i -th agent²². If $\lceil \sum_{i=1}^{\infty} e^{-\delta t_i} \rceil$ is finite then $\mathbb{E}[\max\{\alpha_1 X_1, \alpha_2 X_2, \dots\}]$ is finite as well.

²⁰The calculation is done as follows: define a reward $R(t) = e^{-\delta t}$ for an arrival at time t . The probability assigned to k arrivals is the mass of the set of histories such that the total accumulated reward is between $k - 1$ and k . We can restrict attention to histories up to a finite time T since the reward for later arrivals is negligible.

²¹We set $\mathbb{E}[X_{(0)}] = 0$.

²²For a realization of arrival times where the sum $\sum_{i=1}^{\infty} e^{-\delta t_i}$ does not exist, we consider Problem A with a number of the agents that goes to infinity.

To see that observe that

$$\mathbb{E}[\max \{\alpha_1 X_1, \alpha_2 X_2, \dots\}] \leq \mathbb{E} \sum_{i=1}^{\infty} \alpha_i X_i = \left[\sum_{i=1}^{\infty} e^{-\delta t_i} \right] \mu.$$

For any history where $\sum_{i=1}^{\infty} e^{-\delta t_i}$ exists, we only need to consider the first K arrivals (where K may be arbitrarily large) since the effect of further arrivals is negligible. Therefore, the designer's expected utility in the corresponding Problem A is given by $\mathbb{E}X_{(l)}$ where $l = \left\lceil \sum_{i=1}^K e^{-\delta t_i} \right\rceil$.

Since $\mathbb{E}[\max \{\alpha_1 X_1, \alpha_2 X_2, \dots, \alpha_K X_K\}]$ is monotone in the α_i 's we have

$$\mathbb{E}[\max \{\alpha_1 X_1, \alpha_2 X_2, \dots, \alpha_K X_K\}] \leq \mathbb{E}[\max \{\tilde{\alpha}_1 X_1, \tilde{\alpha}_2 X_2, \dots, \tilde{\alpha}_K X_K\}] \quad (5)$$

where $1 \geq \tilde{\alpha}_i \geq \alpha_i$ for $i \in \{1, \dots, K\}$ and $\sum_{i=1}^K \tilde{\alpha}_i = \left\lceil \sum_{i=1}^K \alpha_i \right\rceil$.²³

Consider now $\sum_{i=1}^K \tilde{\alpha}_i = l$ as above. Since $0 \leq \tilde{\alpha}_i \leq 1$ for any i , the vector $(\tilde{\alpha}_1, \dots, \tilde{\alpha}_K)$ is majorized by the vector $(\underbrace{1, \dots, 1}_l, \underbrace{0, \dots, 0}_{K-l})$. Since the X_i 's are I.I.D., and since the

maximum is a continuous, permutation invariant convex function, we can apply Theorem 3 to obtain

$$\mathbb{E}[\max \{\tilde{\alpha}_1 X_1, \tilde{\alpha}_2 X_2, \dots, \tilde{\alpha}_K X_K\}] \leq \mathbb{E}[\max \{X_1, \dots, X_l\}] = \mathbb{E}X_{(l)}. \quad (6)$$

Inequalities (5) and (6) allow us to conclude that for any realization of the agents' arrivals, the designer's expected utility is higher in Problem A than in Problem B with a corresponding number of the agents.

We have showed that the designer's expected utility in Problem A is finite, and hence that the expected utility in the original dynamic problem with rate $\beta(t)$ is also finite. Therefore, at each point in time, and for any history, there exists a cutoff such that in the original problem the planner finds it optimal to stop if the current agent has a value that exceeds this cutoff. ■

Proof of Proposition 6. Proposition 4 proved the existence of an optimal stopping policy in the problem with arrival rate $\lambda_i(t)$. The optimal stopping cutoff in this problem, $v_T^*(t_{\mathcal{N}(T)})$, only depends on T and on $\mathcal{N}(T)$, and it is non-decreasing in both these variables. The monotonicity follows from the fact that the more agents arrived and the greater t is, the higher is the arrival rate and hence, the higher is the option of continuing. Therefore, the more agents arrived and the greater t is, the higher is the cutoff that will induce the planner to stop. ■

Proof of Corollary 2. If an agent with value v who arrives at t does not get the object at time t , he can never get the object at a later time. In particular, even if recall is allowed, it will not be used by the optimal stopping policy. Hence, setting a price $P(t, v, \eta_{-i}(T), T) = v_T^*(\mathcal{N}(T))$ implements the efficient dynamic policy also under the setting with unobservable arrivals and recall, since postponing an arrival by a single agent necessarily leads to either the object being already sold, or to an increase in its price. ■

²³Since $\alpha_i \leq 1$ and $\sum_{i=1}^K \alpha_i \leq K$ such $\tilde{\alpha}_i$ s always exist.

References

- [1974] 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.
- [2010] Bergemann, D. and Välimäki, J.: "Efficient Dynamic Auctions", *Econometrica* **78**(2), 771-789.
- [2010] Board, S. and Skrzypacz, A.: "Revenue Management with Forward Looking Buyers", discussion paper, UCLA.
- [1993] Boshuizen, F. and Gouweleeuw, J.: "General Optimal Stopping Theorems for Semi-Markov Processes", *Advances in Applied Probability* **24**, 825-846.
- [1996] Caplin, A. and Leahy, J.: "Monetary Policy as a Process of Search", *American Economics Review* **86**(4), 689-702.
- [2010] Cavallo, R., Parkes, D. and Singh, S.: "Efficient Mechanisms with Dynamic Populations and Dynamic Types," discussion paper, Harvard University.
- [2004] Chamley, C.: *Rational Herds: Economic Models of Social Learning*, Cambridge University Press.
- [1978] Dolan, R.J.: "Incentive Mechanisms for Priority Queueing Problems." *Bell Journal of Economics* **9** (2), 421-436.
- [2006] Gallien, J.: "Dynamic Mechanism Design for Online Commerce", *Operations Research* **54** (2), 291-310.
- [2009] Gershkov, A. and Moldovanu, B.: "Learning About The Future and Dynamic Efficiency", *American Economic Review* **99**(4), 1576-1588.
- [2010] 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 Strack, P.: "Revenue Maximizing Mechanisms with Strategic Customers: Name-Your-Own-Price", *working paper*.
- [1985] Glosten, L. and Milgrom, P.R.: "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders", *Journal of Financial Economics* **17**, 71-100.

- [1986] Green, J. and Laffont, J.J.: "Partially Verifiable Information and Mechanism Design", *Review of Economic Studies* **53**, 447-456.
- [2005] Kittsteiner, T., and Moldovanu, B.: "Priority Auctions and Queue Disciplines that Depend on Processing Time", *Management Science* **51**(2), 236-248.
- [1985] Kyle, A.: "Continuous Auctions and Insider Trading", *Econometrica* **53**, 1315-1335.
- [2004] Lavi, R. and Nisan, N.: "Competitive Analysis of Incentive Compatible On-line Auctions," *Theoretical Computer Science* **310**(1), 159-180.
- [1965] Marschall, A. and Proschan, F.: "An Inequality for Convex Functions Involving Majorization," in *Journal of Mathematical Analysis and Applications* **12**, 87-90.
- [2010] Mierendorff, K.: "Optimal Dynamic Mechanism Design with Deadlines", discussion paper, University of Bonn
- [2013] Mierendorff, K.: "The Dynamic Vickrey Auction", *Games and Economic Behavior* **82**, 192-204.
- [1986] Mortensen, D.T.: "Job Search and Labor Market" in *Handbook of Labor Economics*, Ed: O. Ashenfelter and R. Layard, Volume II, Elsevier, 849-919
- [1986] Myerson, R.: "Multistage Games with Communication," *Econometrica* **54** (2), 323-358
- [2013] Pai, M. and Vohra, R.: "Optimal Dynamic Auctions and Simple Index Rules", *Mathematics of Operations Research* **38**, 682-697.
- [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)
- [1995] Pratt, J.W., Raiffa, H. and Schlaifer, R.: *Statistical Decision Theory*, Cambridge: The MIT Press
- [1991] Rob, R.: "Learning and Capacity Expansion under Demand Uncertainty", *Review of Economic Studies* **58**, 655-675.
- [1983] Ross, S.M.: *Stochastic Processes*, Wiley: New York.
- [2012] Said, M.: "Auctions with Dynamic Populations: Efficiency and Revenue Maximization", *Journal of Economic Theory* **147** (6), 2419-2438.
- [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,

- [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.